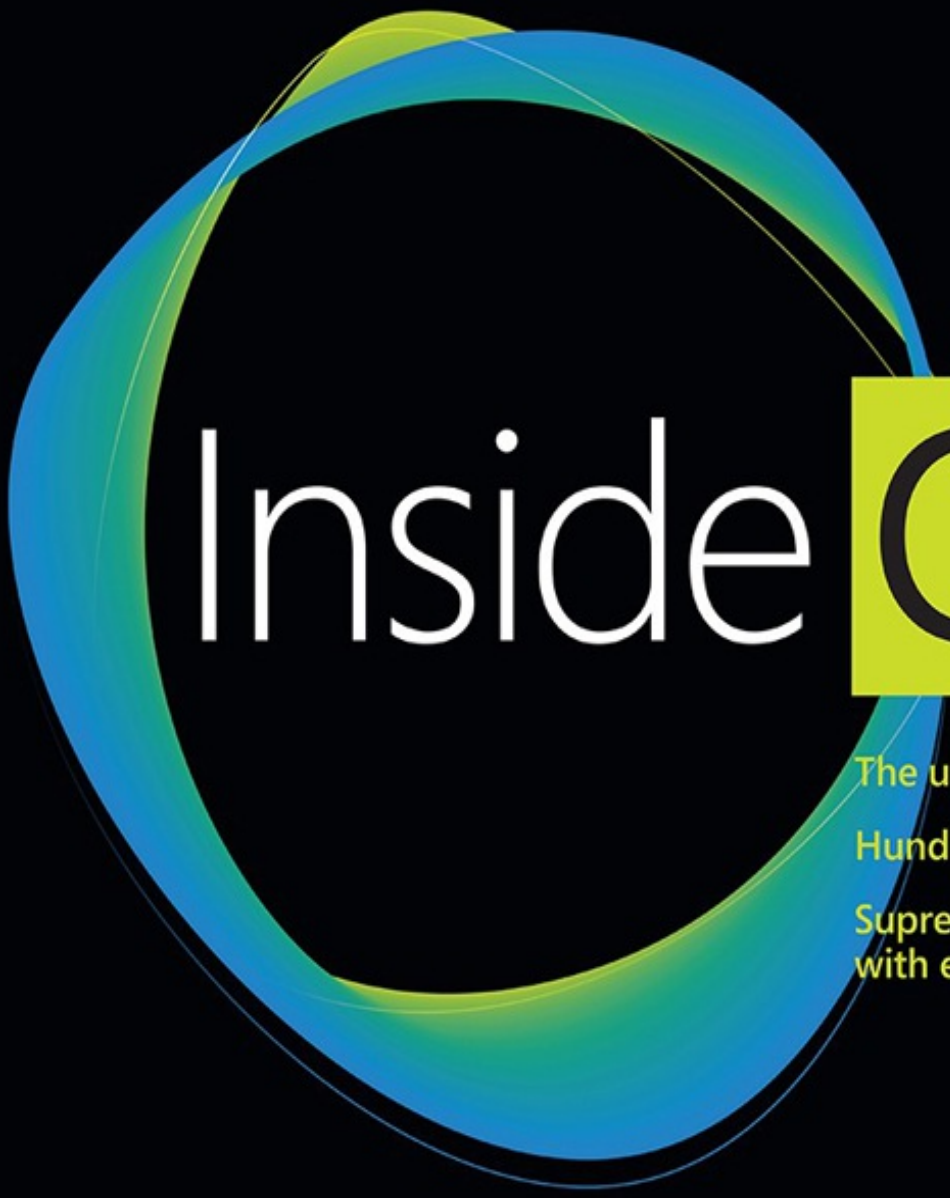
Microsoft



Inside OUT

The ultimate in-depth reference

Hundreds of timesaving solutions

Supremely organized, packed with expert advice

Windows Server 2019

Orin Thomas

Microsoft MVP, RD, Cloud and Datacenter Specialist

Contents

1. Cover Page
2. Title Page
3. Copyright Page
4. Contents at a Glance
5. Table of Contents
6. About the Author
7. Introduction
    1. Changes since Windows Server 2016 Inside Out
    2. Acknowledgments
    3. Errata, updates, and book support
8. Chapter 1. Administration tools
    1. Remote not local
    2. Privileged Access Workstations
    3. Windows Admin Center
    4. Remote Server Administration Tools
    5. PowerShell
    6. Remote Desktop
    7. SSH
9. Chapter 2. Installation options
    1. Windows Server 2019 editions
    2. Windows Server servicing branches
    3. Server Core
    4. Server with Desktop Experience
    5. Roles and features
10. Chapter 3. Deployment and configuration

1. Bare metal versus virtualized
2. Windows images
3. Answer files
4. Windows Deployment Services
5. Virtual Machine Manager
6. Infrastructure configuration as code
7. Desired State Configuration
8. Chef Infra Server
9. Puppet
10. Package-management utilities

11. Chapter 4. Active Directory
    1. Managing Active Directory
    2. Domain controllers
    3. AD DS structure
    4. Accounts
    5. Group policy
    6. Restoring deleted items
    7. Managing AD DS with PowerShell
12. Chapter 5. DNS, DHCP, and IPAM
    1. DNS
    2. DHCP
    3. IPAM
13. Chapter 6. Hyper-V
    1. Dynamic memory
    2. Smart paging
    3. Resource metering
    4. Guest integration services
    5. Generation 2 VMs
    6. Enhanced Session Mode
    7. Discrete Device Assignment
    8. Nested virtualization

9. PowerShell Direct
10. HVC for Linux
11. Virtual hard disks
12. Managing checkpoints
13. Virtual Fibre Channel adapters
14. Storage QoS
15. Hyper-V storage optimization
16. Hyper-V virtual switches
17. Virtual machine network adapters
18. Optimizing network performance
19. Virtual machine MAC addresses
20. Network isolation
21. Hyper-V replica
22. Hyper-V failover clusters
23. Hyper-V guest clusters
24. Live migration
25. Storage migration
26. Exporting, importing, and copying VMs
27. VM Network Health Detection
28. VM drain on shutdown
29. Domain controller cloning
30. Shielded virtual machines
31. Managing Hyper-V using PowerShell

14. Chapter 7. Storage
    1. Storage spaces and storage pools
    2. Storage Replica
    3. SMB 3.1.1
    4. iSCSI
    5. iSNS server
    6. Scale-Out File Servers
    7. Server for NFS

8. Deduplication
9. Storage Quality of Service
10. ReFS
11. Storage-related PowerShell cmdlets

15. Chapter 8. File servers
    1. Shared folder permissions
    2. File Server Resource Manager
    3. Distributed File System
    4. BranchCache
    5. PowerShell commands

16. Chapter 9. Internet Information Services
    1. Managing sites
    2. Application pools
    3. IIS users and delegation
    4. Managing FTP
    5. Managing IIS using PowerShell

17. Chapter 10. Containers
    1. Container concepts
    2. Isolation modes
    3. Managing containers with Docker
    4. Managing containers
    5. Applying updates
    6. Container networking
    7. Linux containers on Windows
    8. Container orchestration

18. Chapter 11. Clustering and high availability
    1. Failover clustering
    2. Network Load Balancing

19. Chapter 12. Active Directory Certificate Services
    1. CA types

2. Certificate revocation lists
3. Certificate Services role services
4. Certificate templates
5. Certificate autoenrollment and renewal
6. CA management

20. Chapter 13. Active Directory Federation Services
    1. AD FS components
    2. Claims, claim rules, and attribute stores
    3. Claims provider
    4. Relying party
    5. Relying party trust
    6. Claims provider trust
    7. Configuring certificate relationship
    8. Attribute stores
    9. Claim rules
    10. Configure Web Application Proxy
    11. Workplace Join
    12. Multifactor authentication
    13. Managing AD FS with PowerShell
    14. Managing Web Application Proxy with PowerShell
21. Chapter 14. Dynamic Access Control and Active Directory Rights Management Services
    1. Dynamic Access Control
    2. Configuring Group Policy to support DAC
    3. Configuring User and Device Claims
    4. Configuring Resource Properties
    5. Central access rules
    6. Central access policies
    7. Staging
    8. Access Denied Assistance
    9. Installing AD RMS

10. AD RMS certificates and licenses
11. AD RMS Templates
12. AD RMS Administrators and Super Users
13. Trusted User and Publishing Domains
14. Exclusion policies
15. Managing AD RMS with Windows PowerShell

22. Chapter 15. Routing and Remote Access
    1. Remote Desktop Gateway
    2. Virtual private networks
    3. LAN routing
    4. Network Address Translation (NAT)
    5. DirectAccess
    6. Managing Remote Access using PowerShell
23. Chapter 16. Remote Desktop Services
    1. Deployment
    2. Remote Desktop Connection Broker
    3. Deployment properties
    4. Remote Desktop Session Host
    5. Remote Desktop Virtualization Host
    6. Remote Desktop Web Access
    7. Remote Desktop licensing
    8. Managing Remote Desktop Services using PowerShell
24. Chapter 17. Azure IaaS and hybrid services
    1. Windows Server IaaS VMs
    2. Azure Active Directory
    3. Azure hybrid cloud services
25. Chapter 18. Windows Subsystem for Linux
    1. Linux on Windows Server
    2. Installing WSL
    3. WSL 2.0

26. Chapter 19. Hardening Windows Server and Active Directory
    1. Hardening Active Directory
    2. Hardening Windows Server
    3. Shielded VMs
27. Chapter 20. Security systems and services
    1. Security Compliance Toolkit
    2. Attack Surface Analyzer
    3. Credential Guard
    4. Windows Defender Application Control
    5. Virtualization-based security
    6. Controlled Folder Access
    7. Exploit Protection
    8. Windows Defender
    9. Windows Defender SmartScreen
28. Chapter 21. Maintenance and monitoring
    1. Data collector sets
    2. Alerts
    3. Event Viewer
    4. Network monitoring
    5. Azure Monitor
    6. Windows Server Backup
    7. Azure Backup
    8. Vssadmin
    9. Windows Server Update Services
    10. Azure Update Management
    11. Monitoring and maintenance related PowerShell cmdlets
    12. WSUS related PowerShell cmdlets
29. Chapter 22. Upgrade and migration
    1. Supported upgrade and migration paths
    2. Active Directory
    3. Active Directory Certificate Services

4. DNS
5. DHCP
6. File and storage servers

30. Chapter 23. Troubleshooting
    1. Troubleshooting methodology
    2. Command-line tools
    3. Sysinternals tools

31. Index
32. Code Snippets

1. i
2. ii
3. iii
4. iv
5. v
6. vi
7. vii
8. viii
9. ix
10. x
11. xi
12. xii
13. xiii
14. xiv
15. xv
16. xvi
17. xvii
18. xviii
19. xix
20. xx
21. xxi
22. xxii

23. 1
24. 2
25. 3
26. 4
27. 5
28. 6
29. 7
30. 8
31. 9
32. 10
33. 11
34. 12
35. 13
36. 14
37. 15
38. 16
39. 17
40. 18
41. 19
42. 20
43. 21
44. 22
45. 23
46. 24
47. 25
48. 26
49. 27
50. 28
51. 29
52. 30
53. 31
54. 32

55. 33
56. 34
57. 35
58. 36
59. 37
60. 38
61. 39
62. 40
63. 41
64. 42
65. 43
66. 44
67. 45
68. 46
69. 47
70. 48
71. 49
72. 50
73. 51
74. 52
75. 53
76. 54
77. 55
78. 56
79. 57
80. 58
81. 59
82. 60
83. 61
84. 62
85. 63
86. 64

87. 65
88. 66
89. 67
90. 68
91. 69
92. 70
93. 71
94. 72
95. 73
96. 74
97. 75
98. 76
99. 77
100. 78
101. 79
102. 80
103. 81
104. 82
105. 83
106. 84
107. 85
108. 86
109. 87
110. 88
111. 89
112. 90
113. 91
114. 92
115. 93
116. 94
117. 95
118. 96

119. 97
120. 98
121. 99
122. 100
123. 101
124. 102
125. 103
126. 104
127. 105
128. 106
129. 107
130. 108
131. 109
132. 110
133. 111
134. 112
135. 113
136. 114
137. 115
138. 116
139. 117
140. 118
141. 119
142. 120
143. 121
144. 122
145. 123
146. 124
147. 125
148. 126
149. 127
150. 128

151. 129
152. 130
153. 131
154. 132
155. 133
156. 134
157. 135
158. 136
159. 137
160. 138
161. 139
162. 140
163. 141
164. 142
165. 143
166. 144
167. 145
168. 146
169. 147
170. 148
171. 149
172. 150
173. 151
174. 152
175. 153
176. 154
177. 155
178. 156
179. 157
180. 158
181. 159
182. 160

183. 161
184. 162
185. 163
186. 164
187. 165
188. 166
189. 167
190. 168
191. 169
192. 170
193. 171
194. 172
195. 173
196. 174
197. 175
198. 176
199. 177
200. 178
201. 179
202. 180
203. 181
204. 182
205. 183
206. 184
207. 185
208. 186
209. 187
210. 188
211. 189
212. 190
213. 191
214. 192

215. 193
216. 194
217. 195
218. 196
219. 197
220. 198
221. 199
222. 200
223. 201
224. 202
225. 203
226. 204
227. 205
228. 206
229. 207
230. 208
231. 209
232. 210
233. 211
234. 212
235. 213
236. 214
237. 215
238. 216
239. 217
240. 218
241. 219
242. 220
243. 221
244. 222
245. 223
246. 224

247. 225
248. 226
249. 227
250. 228
251. 229
252. 230
253. 231
254. 232
255. 233
256. 234
257. 235
258. 236
259. 237
260. 238
261. 239
262. 240
263. 241
264. 242
265. 243
266. 244
267. 245
268. 246
269. 247
270. 248
271. 249
272. 250
273. 251
274. 252
275. 253
276. 254
277. 255
278. 256

279. 257
280. 258
281. 259
282. 260
283. 261
284. 262
285. 263
286. 264
287. 265
288. 266
289. 267
290. 268
291. 269
292. 270
293. 271
294. 272
295. 273
296. 274
297. 275
298. 276
299. 277
300. 278
301. 279
302. 280
303. 281
304. 282
305. 283
306. 284
307. 285
308. 286
309. 287
310. 288

311. 289
312. 290
313. 291
314. 292
315. 293
316. 294
317. 295
318. 296
319. 297
320. 298
321. 299
322. 300
323. 301
324. 302
325. 303
326. 304
327. 305
328. 306
329. 307
330. 308
331. 309
332. 310
333. 311
334. 312
335. 313
336. 314
337. 315
338. 316
339. 317
340. 318
341. 319
342. 320

343. 321
344. 322
345. 323
346. 324
347. 325
348. 326
349. 327
350. 328
351. 329
352. 330
353. 331
354. 332
355. 333
356. 334
357. 335
358. 336
359. 337
360. 338
361. 339
362. 340
363. 341
364. 342
365. 343
366. 344
367. 345
368. 346
369. 347
370. 348
371. 349
372. 350
373. 351
374. 352

375. 353
376. 354
377. 355
378. 356
379. 357
380. 358
381. 359
382. 360
383. 361
384. 362
385. 363
386. 364
387. 365
388. 366
389. 367
390. 368
391. 369
392. 370
393. 371
394. 372
395. 373
396. 374
397. 375
398. 376
399. 377
400. 378
401. 379
402. 380
403. 381
404. 382
405. 383
406. 384

407. 385
408. 386
409. 387
410. 388
411. 389
412. 390
413. 391
414. 392
415. 393
416. 394
417. 395
418. 396
419. 397
420. 398
421. 399
422. 400
423. 401
424. 402
425. 403
426. 404
427. 405
428. 406
429. 407
430. 408
431. 409
432. 410
433. 411
434. 412
435. 413
436. 414
437. 415
438. 416

439. 417
440. 418
441. 419
442. 420
443. 421
444. 422
445. 423
446. 424
447. 425
448. 426
449. 427
450. 428
451. 429
452. 430
453. 431
454. 432
455. 433
456. 434
457. 435
458. 436
459. 437
460. 438
461. 439
462. 440
463. 441
464. 442
465. 443
466. 444
467. 445
468. 446
469. 447
470. 448

471. 449
472. 450
473. 451
474. 452
475. 453
476. 454
477. 455
478. 456
479. 457
480. 458
481. 459
482. 460
483. 461
484. 462
485. 463
486. 464
487. 465
488. 466
489. 467
490. 468
491. 469
492. 470
493. 471
494. 472
495. 473
496. 474
497. 475
498. 476
499. 477
500. 478
501. 479
502. 480

503. 481
504. 482
505. 483
506. 484
507. 485
508. 486
509. 487
510. 488
511. 489
512. 490
513. 491
514. 492
515. 493
516. 494
517. 495
518. 496
519. 497
520. 498
521. 499
522. 500
523. 501
524. 502
525. 503
526. 504
527. 505
528. 506
529. 507
530. 508
531. 509
532. 510
533. 511
534. 512

535. 513
536. 514
537. 515
538. 516
539. 517
540. 518
541. 519
542. 520
543. 521
544. 522
545. 523
546. 524
547. 525
548. 526
549. 527
550. 528
551. 529
552. 530
553. 531
554. 532
555. 533
556. 534
557. 535
558. 536
559. 537
560. 538
561. 539
562. 540
563. 541
564. 542
565. 543
566. 544

567. 545
568. 546
569. 547
570. 548
571. 549
572. 550
573. 551
574. 552
575. 553
576. 554
577. 555
578. 556
579. 557
580. 558
581. 559
582. 560
583. 561
584. 562
585. 563
586. 564
587. 565
588. 566
589. 567
590. 568
591. 569
592. 570
593. 571
594. 572
595. 573
596. 574
597. 575
598. 576

599. 577
600. 578
601. 579
602. 580
603. 581
604. 582
605. 583
606. 584
607. 585
608. 586
609. 587
610. 588
611. 589
612. 590
613. 591
614. 592
615. 593
616. 594
617. 595
618. 596
619. 597
620. 598
621. 599
622. 600
623. 601
624. 602
625. 603
626. 604
627. 605
628. 606
629. 607
630. 608

631. 609
632. 610
633. 611
634. 612
635. 613
636. 614
637. 615
638. 616
639. 617
640. 618
641. 619
642. 620
643. 621
644. 622
645. 623
646. 624
647. 625
648. 626
649. 627
650. 628
651. 629
652. 630
653. 631
654. 632
655. 633
656. 634
657. 635
658. 636
659. 637
660. 638
661. 639
662. 640

663. 641
664. 642
665. 643
666. 644
667. 645
668. 646
669. 647
670. 648
671. 649
672. 650
673. 651
674. 652
675. 653
676. 654
677. 655
678. 656
679. 657
680. 658
681. 659
682. 660
683. 661
684. 662
685. 663
686. 664
687. 665
688. 666
689. 667
690. 668
691. 669
692. 670
693. 671
694. 672

695. 673
696. 674
697. 675
698. 676
699. 677
700. 678
701. 679
702. 680
703. 681
704. 682
705. 683
706. 684
707. 685
708. 686
709. 687
710. 688
711. 689
712. 690
713. 691
714. 692
715. 693
716. 694
717. 695
718. 696
719. 697
720. 698
721. 699
722. 700
723. 701
724. 702
725. 703
726. 704

727. 705
728. 706
729. 707
730. 708
731. 709
732. 710
733. 711
734. 712
735. 713
736. 714
737. 715
738. 716
739. 717
740. 718
741. 719
742. 720
743. 721
744. 722
745. 723
746. 724
747. 725
748. 726
749. 727
750. 728
751. 729
752. 730
753. 731
754. 732
755. 733
756. 734
757. 735
758. 736

759. 737
760. 738
761. 739
762. 740
763. 741
764. 742
765. 743
766. 744
767. 745
768. 746
769. 747
770. 748
771. 749
772. 750
773. 751
774. 752
775. 753
776. 754
777. 755
778. 756
779. 757
780. 758
781. 759
782. 760
783. 761
784. 762
785. 763
786. 764
787. 765
788. 766
789. 767
790. 768

791. 769
792. 770
793. 771
794. 772
795. 773
796. 774
797. 775
798. 776
799. 777
800. 778

Windows Server 2019 Inside Out

Orin Thomas



Windows Server 2019 Inside Out

Published with the authorization of Microsoft Corporation by:
Pearson Education, Inc.

Copyright © 2020 by Orin Thomas
All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, request forms, and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit www.pearson.com/permissions
No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

ISBN-13: 978-0-13-549227-7
ISBN-10: 0-13-549227-0

Technet24

Library of Congress Control Number: 2020935953

`ScoutAutomatedPrintCode`

Trademarks

Microsoft and the trademarks listed at http://www.microsoft.com on the "Trademarks" webpage are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an "as is" basis. The author, the publisher, and Microsoft Corporation shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the programs accompanying it.

Special Sales

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact

our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

Editor-in-Chief: Brett Bartow
Executive Editor: Loretta Yates
Sponsoring Editor: Charvi Arora
Development Editor: Rick Kughen
Managing Editor: Sandra Schroeder
Senior Project Editor: Tracey Croom
Copy Editor: Charlotte Kughen
Indexer: Valerie Haynes Perry
Proofreader: Dan Foster
Technical Editor: Vince Averello
Editorial Assistant: Cindy Teeters
Cover Designer: Twist Creative, Seattle
Compositor: Danielle Foster
Graphics: Vived Graphics

Technet24

Contents at a glance

Chapter 1
Administration tools

Chapter 2
Installation options

Chapter 3
Deployment and configuration

Chapter 4
Active Directory

Chapter 5
DNS, DHCP, and IPAM

Chapter 6
Hyper-V

Chapter 7

Storage

Chapter 8
File servers

Chapter 9
Internet Information Services

Chapter 10
Containers

Chapter 11
Clustering and high availability

Chapter 12
Active Directory Certificate Services

Chapter 13
Active Directory Federation Services

Chapter 14
Dynamic Access Control and Active Directory Rights Management Services

Chapter 15
Routing and Remote Access

Chapter 16
Remote Desktop Services

Chapter 17
Azure IaaS and hybrid services

Chapter 18
Windows Subsystem for Linux

Chapter 19
Hardening Windows Server and Active Directory

Chapter 20
Security systems and services

Chapter 21
Maintenance and monitoring

Chapter 22
Upgrade and migration

Chapter 23
Troubleshooting

Index

Table of Contents

Introduction

Changes since ***Windows Server 2016 Inside Out***

Acknowledgments

Errata, updates, and book support

Chapter 1 Administration tools

Remote not local

Privileged Access Workstations

Windows Admin Center

Installing Windows Admin Center

Windows Admin Center extensions

Show script

Remote Server Administration Tools

RSAT consoles

Server Manager console

PowerShell
    Modules
    PowerShell Gallery
    Remoting
    One-to-many remoting
    PowerShell ISE
    PowerShell Direct
Remote Desktop
SSH

Chapter 2 Installation options

Windows Server 2019 editions
Windows Server servicing branches
    Long Term Servicing Channel
    Semi Annual Channel
    Insider Preview Builds
Server Core
    Server Core interface
    Server Core roles
    Server Core App Compatibility Features on Demand
    When to deploy Server Core
Server with Desktop Experience

Technet24

Roles and features

Chapter 3 Deployment and configuration

Bare metal versus virtualized

Windows images

Modifying Windows images

Servicing Windows images

Mounting images

Adding drivers and updates to images

Adding roles and features

Committing an image

Build and capture

Answer files

Windows Deployment Services

WDS requirements

Managing images

Configuring WDS

Configuring transmissions

Driver groups and packages

Virtual Machine Manager

Virtual machine templates

VMM storage

VMM networking

- Adding a WDS to VMM
- VMM host groups

Infrastructure configuration as code

Desired State Configuration

- DSC configuration files
- Local Configuration Manager
- DSC resources
- DSC push model
- DSC pull server

Chef Infra Server

- Chef servers
- Chef Development Kit
- Deploying Chef agents
- Deploying Chef cookbooks and recipes

Puppet

- Puppet Master Server
- Deploying Puppet agent to Windows Server
- Managing Windows Server configuration
- Puppet Windows Module Pack

Package-management utilities

- PowerShell Gallery
- Chocolatey

Technet24

Chapter 4 Active Directory

- Managing Active Directory
  - Remote rather than local administration
  - Active Directory Administrative Center
  - Active Directory Users and Computers console
  - Active Directory Sites and Services console
  - Active Directory Domains and Trusts console
- Domain controllers
  - Deployment
  - Server Core
  - Global catalog servers
  - Read only domain controllers
  - Virtual domain controller cloning
- AD DS structure
  - Domains
  - Domain functional levels
  - Forests
  - Account and resource forests
  - Organizational units
  - Flexible Single Master Operations roles
- Accounts

- User accounts
- Computer accounts
- Group accounts
- Default groups
- Service accounts

Group policy

- GPO management
- Policy processing
- Group Policy preferences
- Administrative templates

Restoring deleted items

- Active Directory Recycle Bin
- Authoritative restore
- Active Directory snapshots

Managing AD DS with PowerShell

- Active Directory module
- Group Policy module
- ADDSDeployment module

Chapter 5 DNS, DHCP, and IPAM

DNS

- DNS zone types
- Zone delegation
- Forwarders and conditional forwarders

Technet24

- Stub zones
- GlobalNames zones
- Peer Name Resolution Protocol
- Resource records
- Zone aging and scavenging
- DNSSEC
- DNS event logs
- DNS options
- Delegated administration
- Managing DNS with PowerShell

DHCP

- Scopes
- Server and scope options
- Reservations
- DHCP filtering
- Superscopes
- Multicast scopes
- Split scopes
- Name protection
- DHCP failover
- Administration

IPAM

- Deploy IPAM
- Configure server discovery

IPAM Administration
Managing IPAM with PowerShell

Chapter 6 Hyper-V

Dynamic memory
Smart paging
Resource metering
Guest integration services
Generation 2 VMs
Enhanced Session Mode
Discrete Device Assignment
Nested virtualization
Nested virtualization dynamic memory
Nested virtualization networking
PowerShell Direct
HVC for Linux
Virtual hard disks
Fixed-sized disks
Dynamically expanding disks
Differencing disks
Modifying virtual hard disks
Pass-through disks
Managing checkpoints

Technet24

- Virtual Fibre Channel adapters
- Storage QoS
- Hyper-V storage optimization
    - Deduplication
    - Storage tiering
- Hyper-V virtual switches
    - External switches
    - Internal switches
    - Private switches
- Virtual machine network adapters
- Optimizing network performance
    - Bandwidth management
    - SR-IOV
    - Dynamic virtual machine queue
    - Virtual machine NIC teaming
- Virtual machine MAC addresses
- Network isolation
- Hyper-V replica
    - Configuring Hyper-V replica servers
    - Configuring VM replicas
    - Replica failover
    - Hyper-V replica broker
- Hyper-V failover clusters

Hyper-V host cluster storage
Cluster quorum
Cluster networking
Force Quorum Resiliency
Cluster Shared Volumes
Active Directory detached clusters
Preferred owner and failover settings
Hyper-V guest clusters
Hyper-V guest cluster storage
Shared virtual hard disk
Hyper-V VHD Sets
Live migration
Storage migration
Exporting, importing, and copying VMs
VM Network Health Detection
VM drain on shutdown
Domain controller cloning
Shielded virtual machines
Managing Hyper-V using PowerShell

Chapter 7 Storage

Storage spaces and storage pools
Storage pools

Technet24

- Storage space resiliency
- Storage space tiering
- Thin provisioning and trim
- Creating virtual disks
- Storage Spaces Direct

Storage Replica

- Supported configurations
- Configuring replication

SMB 3.1.1

iSCSI

iSNS server

Scale-Out File Servers

Server for NFS

Deduplication

Storage Quality of Service

ReFS

Storage-related PowerShell cmdlets

- Deduplication
- iSCSI
- iSCSITarget
- NFS
- Storage
- Storage Replica

Chapter 8 File servers

- Shared folder permissions
  - Using File Explorer
  - Windows Admin Center
  - Server Manager
- File Server Resource Manager
  - Folder level quotas
  - File screens
  - Storage reports
  - File classification
  - File management tasks
  - Access-Denied Assistance
- Distributed File System
  - DFS namespace
  - DFS replication
- BranchCache
- PowerShell commands
  - Shared Folder cmdlets
  - File Server Resource Manager cmdlets
  - BranchCache Cmdlets
  - DFS Cmdlets

Chapter 9 Internet Information Services

- Managing sites
  - Adding websites
  - Virtual directories
  - Modifying site settings
  - Adding web applications
  - Configuring TLS certificates
  - Site authentication
  - Modifying custom error response
  - Adding or disabling the default document
  - Directory browsing
  - IP address and domain name filtering
  - URL authorization rules
  - Request filters
- Application pools
  - Creating application pools
  - Configuring application pool recycling settings
- IIS users and delegation
  - IIS user accounts
  - Delegating administrative permissions
- Managing FTP
- Managing IIS using PowerShell

Chapter 10 Containers

- Container concepts
- Isolation modes
  - Process Isolation mode
  - Hyper-V Isolation mode
- Managing containers with Docker
  - Installing Docker
  - Demon.json
  - Retrieving container OS image
  - Container registries and images
- Managing containers
  - Starting a container
  - Modifying a running container
  - Creating a new image from a container
  - Using Dockerfiles
  - Managing container images
  - Service accounts for Windows containers
- Applying updates
- Container networking
  - NAT
  - Transparent
  - Overlay

Layer 2 Bridge
Linux containers on Windows
Container orchestration
Kubernetes
Docker Swarm

Chapter 11 Clustering and high availability

Failover clustering
Cluster quorum modes
Cluster storage and cluster shared volumes
Cluster networks
MPIO
Cluster Aware Updating
Failover and preference settings
Multisite clusters
Cloud witness
Virtual machine failover clustering
Rolling upgrades
Workgroup clusters
Cluster sets
Managing failover clustering with PowerShell
Network Load Balancing

Network Load Balancing prerequisites
NLB cluster operation modes
Managing cluster hosts
Port rules
Filtering and affinity
Managing NLB with PowerShell

Chapter 12 Active Directory Certificate Services

CA types
Enterprise CA
Standalone CAs
Certificate revocation lists
CRL distribution points
Authority Information Access
Revoking a certificate
Publishing CRLs and delta CRLs
Certificate Services role services
Certificate templates
Template properties
Adding and editing templates
Certificate autoenrollment and renewal
CA management
Handling certificate requests

Technet24

- CA backup and recovery
- Key archiving and recovery
- CAPolicy.inf
- Managing Certificate Services using PowerShell
- Managing Certificate Services using Certutil.exe and Certreq.exe

Chapter 13 Active Directory Federation Services

- AD FS components
- Claims, claim rules, and attribute stores
- Claims provider
- Relying party
- Relying party trust
- Claims provider trust
- Configuring certificate relationship
- Attribute stores
- Claim rules
  - Relying party trust claim rules
  - Claims provider trust claim rules
- Configure Web Application Proxy
- Workplace Join

Multifactor authentication

Managing AD FS with PowerShell

Managing Web Application Proxy with PowerShell

Chapter 14 Dynamic Access Control and Active Directory Rights Management Services

Dynamic Access Control

Configuring Group Policy to support DAC

Configuring User and Device Claims

Configuring Resource Properties

Central access rules

Central access policies

Staging

Access Denied Assistance

Installing AD RMS

AD RMS certificates and licenses

AD RMS Templates

AD RMS Administrators and Super Users

Trusted User and Publishing Domains

Exclusion policies

Apply AD RMS templates automatically

Technet24

Managing AD RMS with Windows PowerShell

Dynamic Access Control cmdlets

Chapter 15 Routing and Remote Access

Remote Desktop Gateway

RD Gateway connection and resource policies

Configuring server settings

Configuring clients to use RD Gateway

Virtual private networks

IKEv2 Always On VPN protocol

SSTP VPN protocol

L2TP/IPsec protocols

PPTP VPN protocol

VPN authentication

Deploying a VPN server

Disable VPN protocols

Granting access to a VPN server

LAN routing

Network Address Translation (NAT)

DirectAccess

DirectAccess topologies

DirectAccess server

Network Location Server
Configuring DirectAccess
Managing Remote Access using PowerShell

Chapter 16 Remote Desktop Services

Deployment
Remote Desktop Connection Broker
Deployment properties
Remote Desktop Session Host
Session collection settings
Personal session desktops
RemoteApp
Group Policy configuration
Remote Desktop Virtualization Host
Virtual machine preparation
Virtual desktop collections
Pooled virtual desktops
Personal virtual desktops
DDA and RemoteFX
Remote Desktop Web Access
Remote Desktop licensing
Installing RDS CALs
Activating a License Server

Technet24

Managing Remote Desktop Services using PowerShell

Chapter 17 Azure IaaS and hybrid services

Windows Server IaaS VMs

Creating Azure IaaS VMs

IaaS VM networking

IaaS VM administration

Azure Active Directory

Azure Active Directory Connect

Azure AD Connect server requirements

Installing Azure AD Connect

Using UPN suffixes and non-routable domains

Monitor Azure AD Connect Health

Forcing synchronization

Configure object filters

Implement and manage Azure AD self-service password reset

Azure AD Password Protection

Azure AD DS

Azure hybrid cloud services

Connect Windows Admin Center

Creating Azure IaaS VMs from Windows

Admin Center
Azure File Sync
Azure Arc
Azure Site Recovery
Azure Network Adapter

Chapter 18 Windows Subsystem for Linux

Linux on Windows Server
Installing WSL
WSL 2.0

Chapter 19 Hardening Windows Server and Active Directory

Hardening Active Directory
Hardening domain controllers
Least privilege
Role-Based Access Control
Password policies
Account security options
Protected accounts
Authentication policies silos
Disable NTLM
Block server operators from scheduling tasks

Technet24

- Enable Local Security Authority protection
- KRBTGT account password
- Enhanced Security Administrative Environment forest

Hardening Windows Server

- User rights
- Service accounts
- Just Enough Administration
- Privileged Access Management
- Local Administrator Password Solution
- Advanced auditing
- Windows Firewall with Advanced Security

Shielded VMs

- Guarded fabric

Chapter 20 Security systems and services

Security Compliance Toolkit

- Policy Analyzer tool
- Local Group Policy Object tool

Attack Surface Analyzer

Credential Guard

Windows Defender Application Control

Virtualization-based security

Controlled Folder Access

Exploit Protection

Windows Defender

Windows Defender SmartScreen

Chapter 21 Maintenance and monitoring

Data collector sets

Alerts

Event Viewer

- Event log filters
- Event log views
- Event subscriptions
- Event-driven tasks

Network monitoring

- Resource Monitor
- Message Analyzer

Azure Monitor

Windows Server Backup

- Backup locations
- Backing up data
- Role- and application-specific backups
- Restore from backups

Technet24

Restore to an alternative location
Azure Backup
Preparing Azure Backup
Backing up data to Azure Backup Agent
Restore from Azure Backup
Vssadmin
Windows Server Update Services
Products, security classifications, and languages
Autonomous and replica modes
Update files
WSUS security roles
WSUS groups
WSUS policies
Deploying updates
Automatic approval rules
Azure Update Management
Monitoring and maintenance related PowerShell cmdlets
WSUS related PowerShell cmdlets

Chapter 22 Upgrade and migration

Supported upgrade and migration paths
Upgrading roles and features

Converting evaluation version to licensed version
Upgrading editions
Windows Server Migration Tools
Active Directory
FRS to DFSR migration
Migrating to a new forest
Active Directory Certificate Services
Preparation
Migration
Verification and post migration tasks
DNS
DHCP
Preparing to migrate DHCP
Migration
Verification and post migration tasks
File and storage servers
Migrate file servers using Storage Migration Service
Migrate file and storage servers using WSMT

Chapter 23 Troubleshooting

Troubleshooting methodology

Technet24

- Redeployment
- Symptoms and diagnosis
- Dependencies
- Ranking hypothetical solutions
- Applying solutions

Command-line tools

Sysinternals tools

- Process Explorer
- Process Monitor
- ProcDump
- PsTools
- VMMap
- SigCheck
- AccessChk
- Sysmon
- AccessEnum
- ShellRunAs
- LogonSessions
- Active Directory Explorer
- Insight for Active Directory
- PsPing
- RAMMap

Index

About the author



Orin Thomas is a principal cloud operations advocate at Microsoft and has written more than three dozen books for Microsoft Press on topics including Windows Server, Windows Client, Azure, Office 365, System Center, Exchange Server, Security, and SQL Server. He has authored Azure Architecture courses at Pluralsight, has authored multiple Microsoft Official Curriculum and EdX courses on a variety of IT Pro topics, and is completing his Doctorate in Information Technology on cloud computing security and compliance at Charles Sturt University. You can follow him on twitter at *http://twitter.com/orinthomas*.

Introduction

This book is primarily written for IT Professionals who work with Windows Server operating systems on a regular basis. As such, it's likely that Windows Server 2019 isn't the first version of Windows Server that you've been responsible for managing. This is because the majority of Windows Server administrators have been working with some version of the operating system for more than a decade, with a good percentage having experience going back to the days of Windows NT 4. With that in mind, this book doesn't spend a great amount of time on introductory concepts or techniques; instead, it aims to provide intermediate to advanced coverage of the most important roles and features available with Windows Server 2019, including its hybrid cloud capabilities.

This book is also written under the assumption that as an experienced IT professional, you know how to use a search engine to find relevant technical information. This leads to an obvious question, "Why would I buy a book if I can find relevant technical information with a search engine?" The answer is that even though you

may be good at tracking down technical information and have experience filtering useful knowledge from wildly inaccurate guesses, you can only search for something if you have an idea about it in the first place.

When presenting at conferences and user groups on Windows Server topics, I regularly encounter IT Professionals who have worked with Windows Server for many years who are unaware of specific functionality or techniques related to the product, even if that functionality or technique has been available for many years. This is because Windows Server 2019 includes so many roles, features, and moving parts, you are simply unlikely to know everything about the operating system. My aim in writing this book is to give you comprehensive coverage so that you'll learn things that you didn't know or simply missed, because when you've been solving a critical problem, you've been focused on the specifics of that problem and haven't had time to explore every facet of what the Windows Server operating system is capable of.

CHANGES SINCE WINDOWS SERVER 2016 INSIDE OUT

This book includes several new chapters as well as revisions and updates—from moderate to substantive—of chapters that were present in the Windows Server

Technet24

2016 version of this book. Some of the substantive changes include removing content related to the Nano Server deployment option, which is no longer supported; coverage of Windows Admin Center; a new chapter on Azure IaaS and hybrid services, as well as Windows Subsystem for Linux; and splitting an extension of the security chapter from the 2016 edition into two separate chapters in this text. There is also coverage of new roles and features including Storage Migration Services, Azure File Sync, and Azure Update Management. While there are some chapters where only cosmetic changes were made, from the perspective of total word count, this book is about 15 percent longer than its predecessor, *Windows Server 2016 Inside Out*.

ACKNOWLEDGMENTS

I'd like to thank Rick Kughen, Vince Averello, Dan Foster, Charvi Arora, and Loretta Yates for the assistance they provided in bringing this text to print. I would also like to thank Thomas Maurer for his advice on revisions for this new edition of the text.

ERRATA, UPDATES, AND BOOK SUPPORT

We've made every effort to ensure the accuracy of this book and its companion content. You can access updates

to this book—in the form of a list of submitted errata and their related corrections—at:

MicrosoftPressStore.com/WindowsServer2019InsideOut/errata

If you discover an error that is not already listed, please submit it to us at the same page.

For additional book support and information, please visit:

MicrosoftPressStore.com/Support

Please note that product support for Microsoft software and hardware is not offered through the previous addresses. For help with Microsoft software or hardware, go to *http://support.microsoft.com*.

Technet24

Chapter 1

Administration tools

Remote not local

Privileged Access Workstations

Windows Admin Center

Remote Server Administration Tools

PowerShell

Remote Desktop

SSH

You can use a variety of tools to manage Windows Server 2019. Some, such as PowerShell, the Microsoft Management Consoles, and Server Manager, are built in to the operating systems. You'll need to download others, such as Windows Admin Center, for free from Microsoft's website.

Microsoft's general systems administration philosophy is that while you can do almost everything with a

graphical console such as Windows Admin Center, Active Directory Administrative Center, or the Server Manager console, any task that you do repeatedly should be automated using Windows PowerShell. Microsoft best practice is that almost all administration tasks should be performed remotely rather than by signing in directly to the server and performing them locally.

In this chapter, we'll look at how you can perform administration tasks remotely, what to consider when putting together your remote administration toolkit, and the various tools that you can use to remotely administer Windows Server 2019.

REMOTE NOT LOCAL

Windows Server is designed to be administered remotely rather than locally. This "remote first" philosophy shouldn't come as a surprise to experienced administrators. The vast majority of Windows Server instances are running as virtual machines, either in datacenters or in the cloud, and we are long past the days where your primary method of switching between different servers that you were working on was by selecting different options on a KVM switch.

You need to be familiar with how to use your tools remotely. You should avoid signing in to each server

individually using Remote Desktop and firing up the console that is relevant to the role or feature that you want to manage. You should also avoid using Remote Desktop to connect to a server just to run a PowerShell script.

The main reason you should avoid Remote Desktop is that it only allows you to interact with one server at a time. Consider the amount of time that it takes to perform a task, such as resetting the password of an unprivileged local user account on 100 different servers. If you were doing it using Remote Desktop, you'd have to connect to each server individually, change the password, sign off, and move on to the next server. If you were doing it by leveraging PowerShell's remote functionality, you might be able to accomplish the same goal with a single script that you could write and execute in a fraction of the time.

Inside OUT

Automate where possible

Automate your tasks where possible and when it makes sense. Understand that you won't be able to automate everything. Concentrate on automating what you can. In the long run, automation will reduce the amount of time that you spend on tasks that you know how to do. It will also allow you to spend more time figuring out how to perform those complex tasks you don't know how to do or that resist automation.

PRIVILEGED ACCESS WORKSTATIONS

Servers are only as secure as the computers that you use to manage them. An increasing number of security incidents have occurred because a privileged user's computer was infected with malware and that computer was then used to perform server administration tasks. Privileged Access Workstations (PAW) are specially configured computers that you use to perform remote administration tasks. The idea of a PAW is that you have a computer with a locked-down configuration that you only use to perform server administration tasks. You don't use this computer to read your email or browse the Internet; you just use it to perform server administration tasks.

Consider configuring a PAW in the following way:

- Configure Windows Defender Application Control (Device Guard) to allow only specifically authorized and digitally signed software to run on the computer.
- Configure Credential Guard to protect credentials stored on the computer.
- Use BitLocker to encrypt the computer's storage and protect the boot environment.
- The computer should not be used to browse the Internet or to check email. Server administrators should have completely separate computers to perform their other daily job tasks. Block Internet browsing on the PAW both locally and on the perimeter network firewall.
- Block the PAW from accessing the Internet. Software updates should be obtained from a dedicated secure update server on the local network. External tools should be obtained from another

computer and transferred to the PAW.

- Server administrators should not sign in to the PAW using an account that has administrative privileges on the PAW. In Chapter 19, "Hardening Windows Server and Active Directory," you'll learn about using secure account forests to improve account security.
- Only specific user accounts used by server administrators should be able to sign on to the PAW. Consider additional restrictions such as sign-in hours. Block privileged accounts from signing in to computers that are not PAWs or servers to be managed, such as the IT staff's everyday work computers.
- Configure servers to only accept administrator connections from PAWs. This can be done through Windows Defender Firewall with Advanced Security.
- Use configuration-management tools to monitor the configuration of the PAW. Some organizations rebuild PAWs entirely every 24 hours to ensure that configurations are not altered. Use these tools to restrict local group membership and ensure that the PAW has all appropriate recent software updates applied.
- Ensure that audit logs from PAWs are forwarded to a separate secure location.
- Disable the use of unauthorized storage devices. For example, you can configure policies so that only USB storage devices that have a specific BitLocker organizational ID can be used with the computer.
- Block unsolicited inbound network traffic to the PAW using Windows Defender Firewall.

Inside OUT

Jump Servers

Jump servers are another security procedure that can be used in conjunction with Privileged Access Workstations. Jump servers only allow servers to accept administrative connections from specific hosts. For example, you only allow domain controllers to be administered from computers that have a specific IP address and a computer certificate issued by a specific certification authority. You can configure

jump servers to only accept connections from PAWs and servers to be administered to only accept connections from jump servers. Some organizations that use jump servers have them rebuilt and redeployed every 24 hours to ensure that their configuration does not drift from the approved configuration.

WINDOWS ADMIN CENTER

Windows Admin Center is a web-based console that allows you to remotely manage Windows Server through a web browser. You can connect to and use Windows Admin Center using Edge, Firefox, and Chrome (though not through Internet Explorer). You can install Windows Admin Center on computers running Windows 10, Windows Server 2016, or Windows Server 2019. Windows Admin Center can be installed on Windows Server deployed using the Server Core installation option.

You can use Windows Admin Center, shown in Figure 1-1, to manage computers running Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019. You can use Windows Admin Center to perform a limited number of tasks on computers running Windows Server 2008 R2, but as Windows Server 2008 R2 fell out of support in January 2020, it's reasonable to assume that no additional administrative functionality for this older operating system will be added to future versions of Windows Admin Center.

Technet24

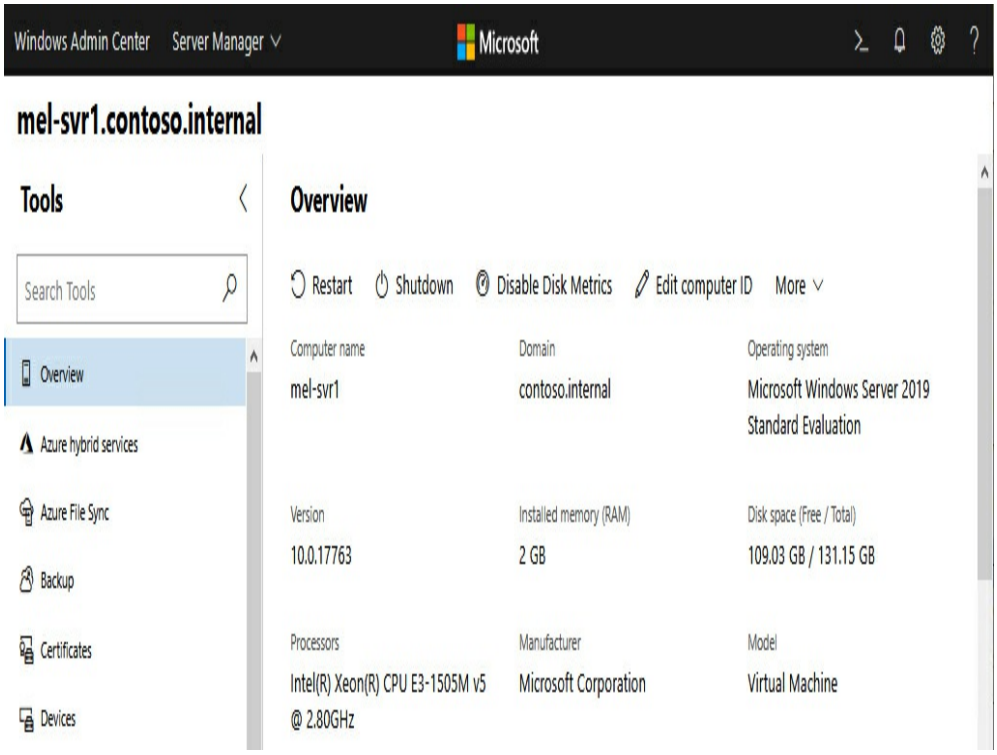


Figure 1-1 Windows Admin Center

When connected to a gateway server, click Add on the All Connections page, shown in Figure 1-2, to add additional servers that you want to manage. A dialog box will prompt you to enter the credentials of an account that has local Administrative privileges on the server that you wish to add.

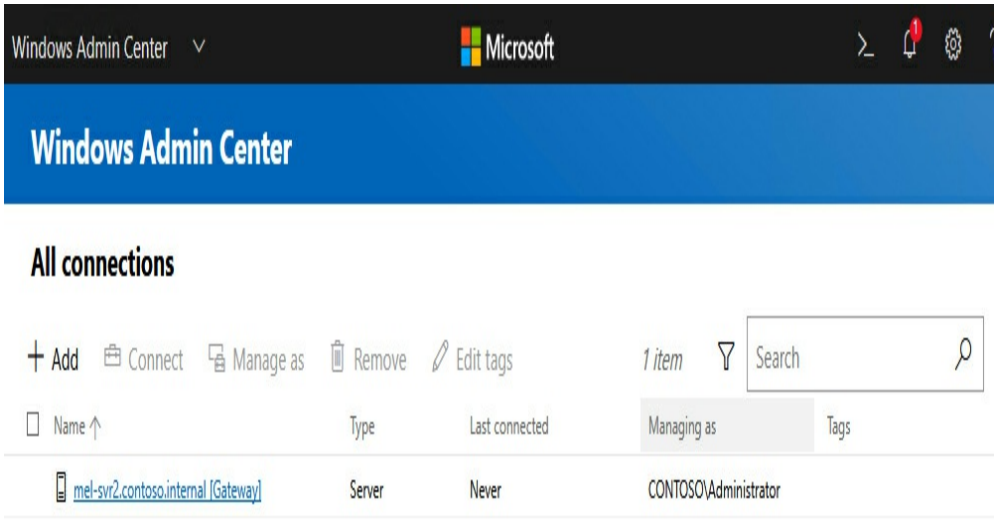


Figure 1-2 Windows Admin Center All Connections

You can use Windows Admin Center (WAC) to perform the on-premises server management tasks listed in Table 1-1. Note that many tools within WAC only allow you to perform some—not all—of the tasks that you can currently perform using one of the older Microsoft Management Consoles. Eventually, these tools will evolve until they reach parity.

Table 1-1 WAC tools and functionality

WAC Tool	Function
Overview	Server details and control server state
Active Directory	Perform some Active Directory management tasks
Backup	Manage Azure Backup

Technet24

Certificates	View and modify certificates, including certificate expiration
Containers	Manage containers
Devices	Manage devices (similar to the Device Manager console)
DHCP	Perform some DHCP administration tasks
DNS	Perform some DNS administration tasks
Events	View events. Similar to the Event Viewer console
Files	Browse files and folders
Firewall	Manage firewall rules
Installed Apps	Add and remove programs
Local Users and Groups	Manage local users and groups
Network	Manage network devices
PowerShell	Interact with the server through a web-based PowerShell session
Processes	Manage running processes
Registry	Remotely manage the registry

Remote Desktop	Connect to the server using a web-based remote desktop client
Roles and Features	Manage server roles and features
Scheduled Tasks	Manage scheduled tasks
Services	Manage services
Storage	Manage server storage devices
Storage Migration Service	Migrate file servers to Azure or Windows Server 2019
Storage Replica	Manage storage replication
System Insights	View trend analysis data about CPU, storage, and network utilization
Updates	Manage software updates
Virtual Machines	Manage Hyper-V virtual machines
Virtual Switches	View and manage virtual switches

Windows Admin Center can also be used to manage Azure hybrid services, such as Azure Backup, Azure Software Update, Azure Site Recovery, Azure Network Adapter, and Azure Monitor. You'll learn more about the

Azure hybrid services available through Windows Admin Center in Chapter 17, "Azure IaaS and hybrid services."

Windows Admin Center (WAC) is complementary to—rather than a complete replacement for—existing RSAT tools. It might appear to be the WAC product team's intention to replicate all tasks that can be performed using the consoles that are present in the Remote Server Administration Tools (RSAT). However, at the current pace—and given the need to surface any new Windows Server functionality in WAC—it's unlikely that WAC will serve as a complete replacement for the RSAT tools any time before the next Long Term Servicing Channel (LTSC) release (which might or might not be Windows Server 2022).

Inside OUT

The Future is WAC

Going forward, graphical management of new features in Windows Server will be possible only through WAC. While WAC eventually will supplant existing GUI tools such as Microsoft Management Consoles and the Server Management Console, there will be a period of some years when you'll have to use WAC and the traditional consoles, such as Active Directory Users and Computers and Active Directory Administrative Center, to get tasks done (because no one tool yet covers all functionality). Generally speaking, if it's something new to Windows Server 2019, you'll probably be doing it through WAC or PowerShell. If it's something that you've been doing with Windows Server for some time, then there's a good chance you'll continue to use your original toolkit for the next few years.

Installing Windows Admin Center

You can install Windows Admin Center (WAC) by downloading it from Microsoft's website. There are four Windows Admin Center deployment options: Local client, gateway server, managed server, and failover cluster.

- **Local client.** When you choose this installation option, you install Windows Admin Center on your workstation. You connect to the WAC instance locally, which is similar to installing the RSAT tools on a local workstation. When you install WAC locally, a shortcut to the WAC console is placed on your desktop.
- **Gateway server.** When you install Windows Admin Center in the gateway server configuration, you install WAC on a computer running Windows Server 2016 or Windows Server 2019 and then make remote connections to the WAC instance hosted on that computer through your preferred browser. Once connected to the WAC instance, you can add servers that you want to manage to the web-based console. When you perform an administrative task, the instructions to perform that task are issued from the gateway server and are run against the target server.
- **Managed server.** The managed server deployment is a version of WAC in gateway server configuration deployed on a cluster node to manage the cluster.
- **Failover cluster.** The gateway server is deployed as a highly available service. This requires the configuration of a Cluster Shared Volume to store persistent data used by WAC. A script is available from Microsoft's website that simplifies the process of performing a high availability deployment.

Windows Admin Center comes in two versions. There's the preview version, which will have the most recent feature updates, and there's the "broadly validated" version that is likely to be more reliable. You should

only use the broadly validated version in production environments; the preview version is useful for experimental or test environments.

By default, when you install WAC on a Windows 10 computer, the service will be available on local port 6516. Installing WAC on Windows 10 places a WAC icon on the desktop with a shortcut to the local WAC instance.

When you install on Windows Server configured with the desktop experience GUI installation option, you get the option of configuring which port will be used, as shown in Figure 1-3. Also, you can choose between using a self-signed SSL (TLS) certificate or an SSL (TLS) certificate that is already installed on the computer. If you're deploying a gateway server, things will be a lot simpler if you deploy a TLS certificate from a trusted certificate authority (CA) because it won't be necessary to go through the hassle of responding to dialog boxes about whether to trust the certificate.

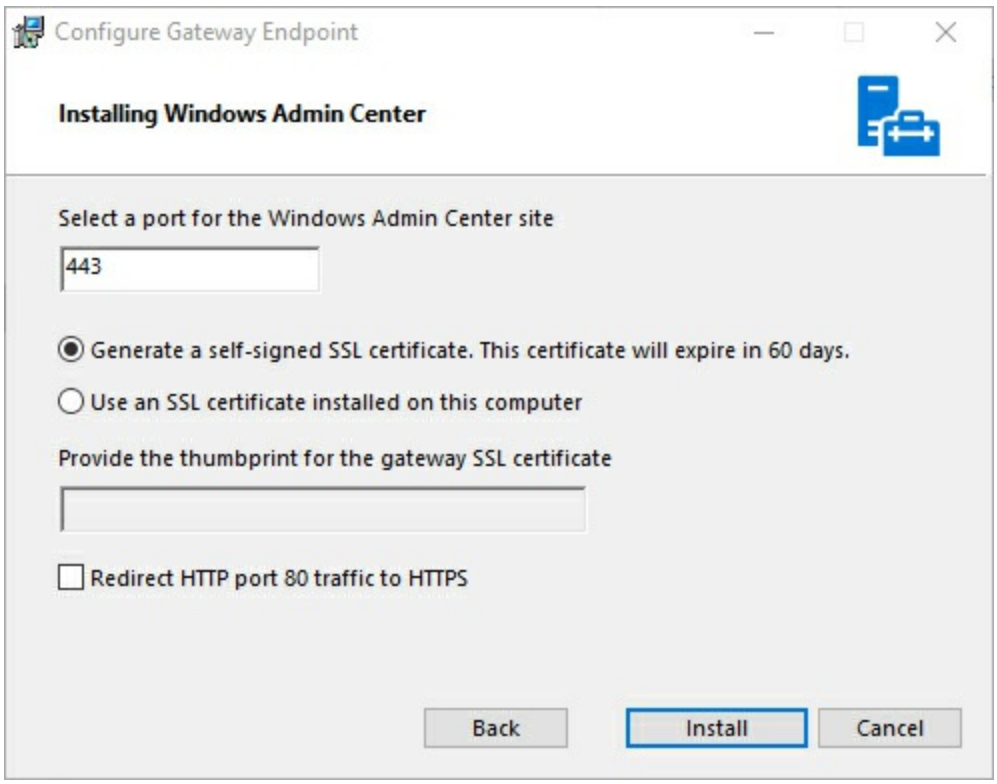


Figure 1-3 Select Windows Admin Center port

You can install Windows Admin Center on a Server Core instance of Windows Server using msiexec and by specifying the management port and SSL certificate option. (It should be TLS certificate as the SSL protocol has been phased out.) The syntax of the command line installation where a trusted certificate is used is:

Click here to view code image

```
msiexec /i <WindowsAdminCenterInstallerName>.msi /qn /L*v log.txt
SME_PORT=<port>
SME_THUMBPRINT=<thumbprint> SSL_CERTIFICATE_OPTION=installed
```

The SME_PORT is the port you want to use, and the

SME_THUMBPRINT is the thumbprint of the installed SSL (TLS) certificate.

By default, installing WAC updates the computer's trusted host files. When you deploy WAC, you can configure it to update automatically or manually. When you configure Windows Admin Center to update automatically, as shown in Figure 1-4, new versions will be installed as they become available through Microsoft Update. If you don't configure this option, you'll need to manually install newer versions of Windows Admin Center as they become available.

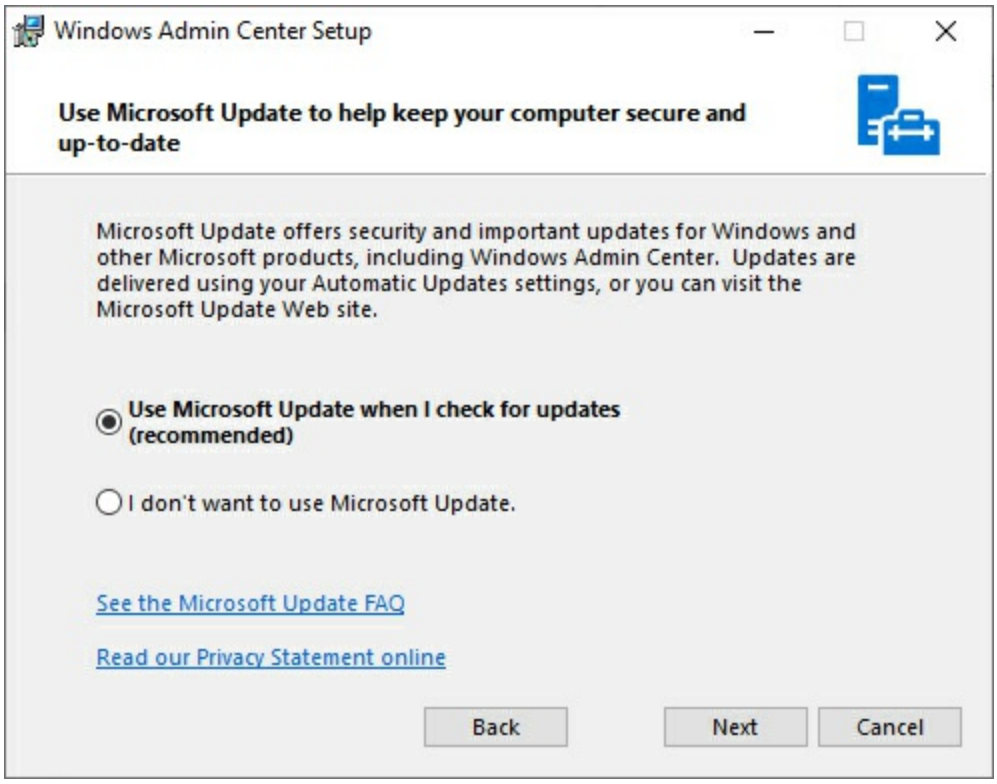


Figure 1-4 Configure Windows Admin Center Setup Updates

More Info

Installing Windows Admin Center

You can learn more about installing Windows Admin Center at https://docs.microsoft.com/en-us/windows-server/manage/windows-admin-center/plan/installation-options.

Inside OUT

WAC Web Server

Windows Admin Center doesn't use Internet Information Services (IIS); instead, it uses its own web server. While it's possible to deploy Windows Admin Center on a computer that has IIS installed, generally you'll be better off deploying WAC in a gateway configuration and using the gateway server to remotely manage a computer that has IIS installed on it.

Windows Admin Center extensions

Windows Admin Center extensions allow for the extension of Windows Admin Center functionality. Windows Admin Center includes extensions for roles built in to Windows Server, such as Storage Migration Services and third-party extensions. Microsoft is encouraging third parties to add extensions to Windows Admin Center as an alternative to requiring systems administrators to use product-specific consoles.

By default, Windows Admin Center will display extensions published to Microsoft's official NuGet feed. This feed includes extensions published and updated by Microsoft as well as those published by trusted third-party vendors. Also, you can configure Windows Admin Center to display extensions or installations from any

NuGet feed that supports the NuGet V2 APIs or a specially configured file share accessible to the computer hosting Windows Admin Center.

Extensions are available in Windows Admin Center by selecting the settings icon and then selecting Extensions. The Available Extension pane, shown in Figure 1-5, displays all extensions that are available but not installed from the currently configured feed. You can update currently installed extensions if new versions of those extensions are available through the Installed Extensions pane.

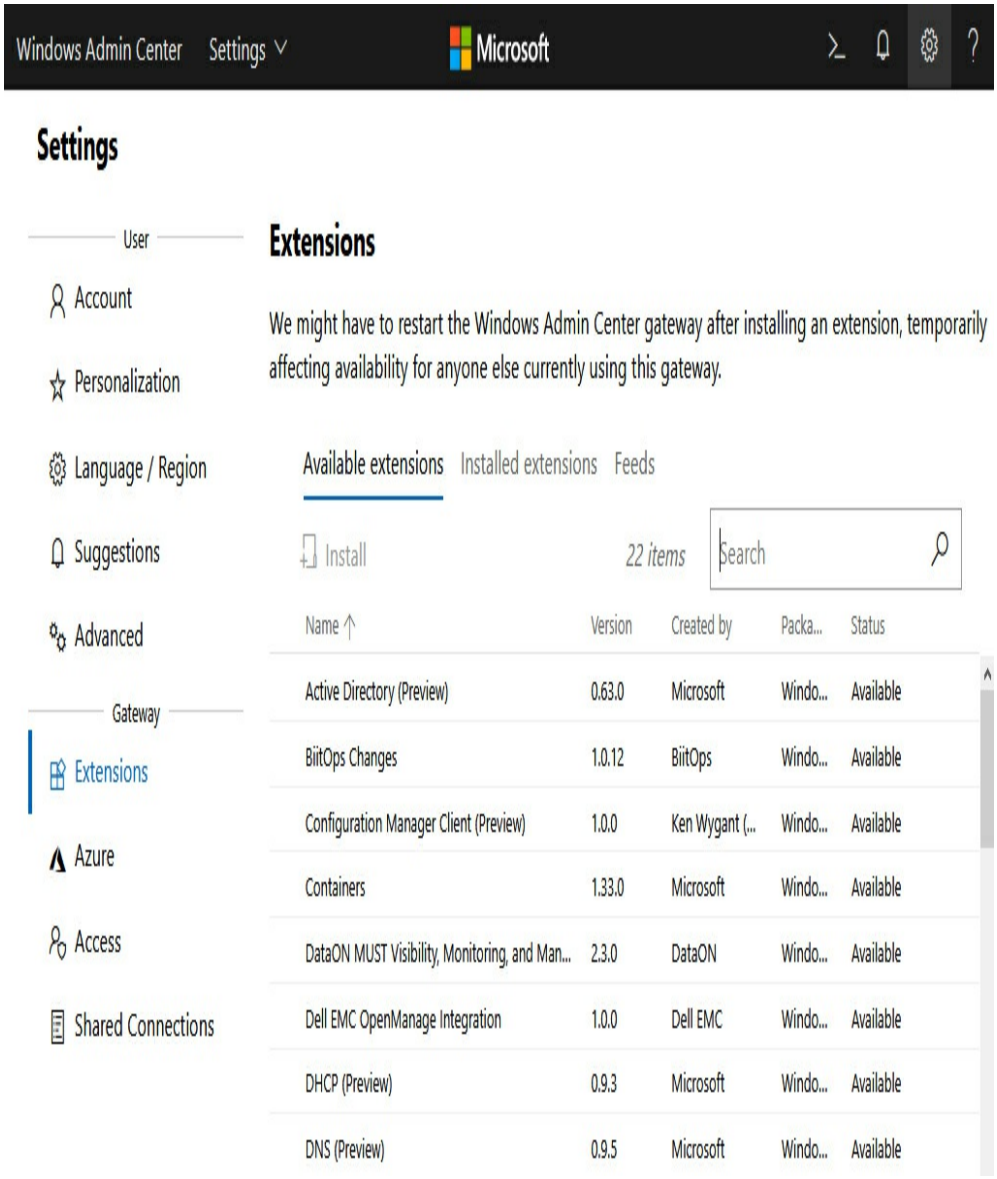


Figure 1-5 Windows Admin Center Extensions

More Info

WAC extensions

You can learn more about WAC extensions at https://docs.microsoft.com/en-us/windows-server/manage/windows-admin-center/configure/using-extensions.

Show script

When you perform a task in Windows Admin Center, you can click the PowerShell icon in the Windows Admin Center title bar to view PowerShell source code relevant to the tasks, as shown in Figure 1-6. This allows you to copy and save useful PowerShell code for reuse later.

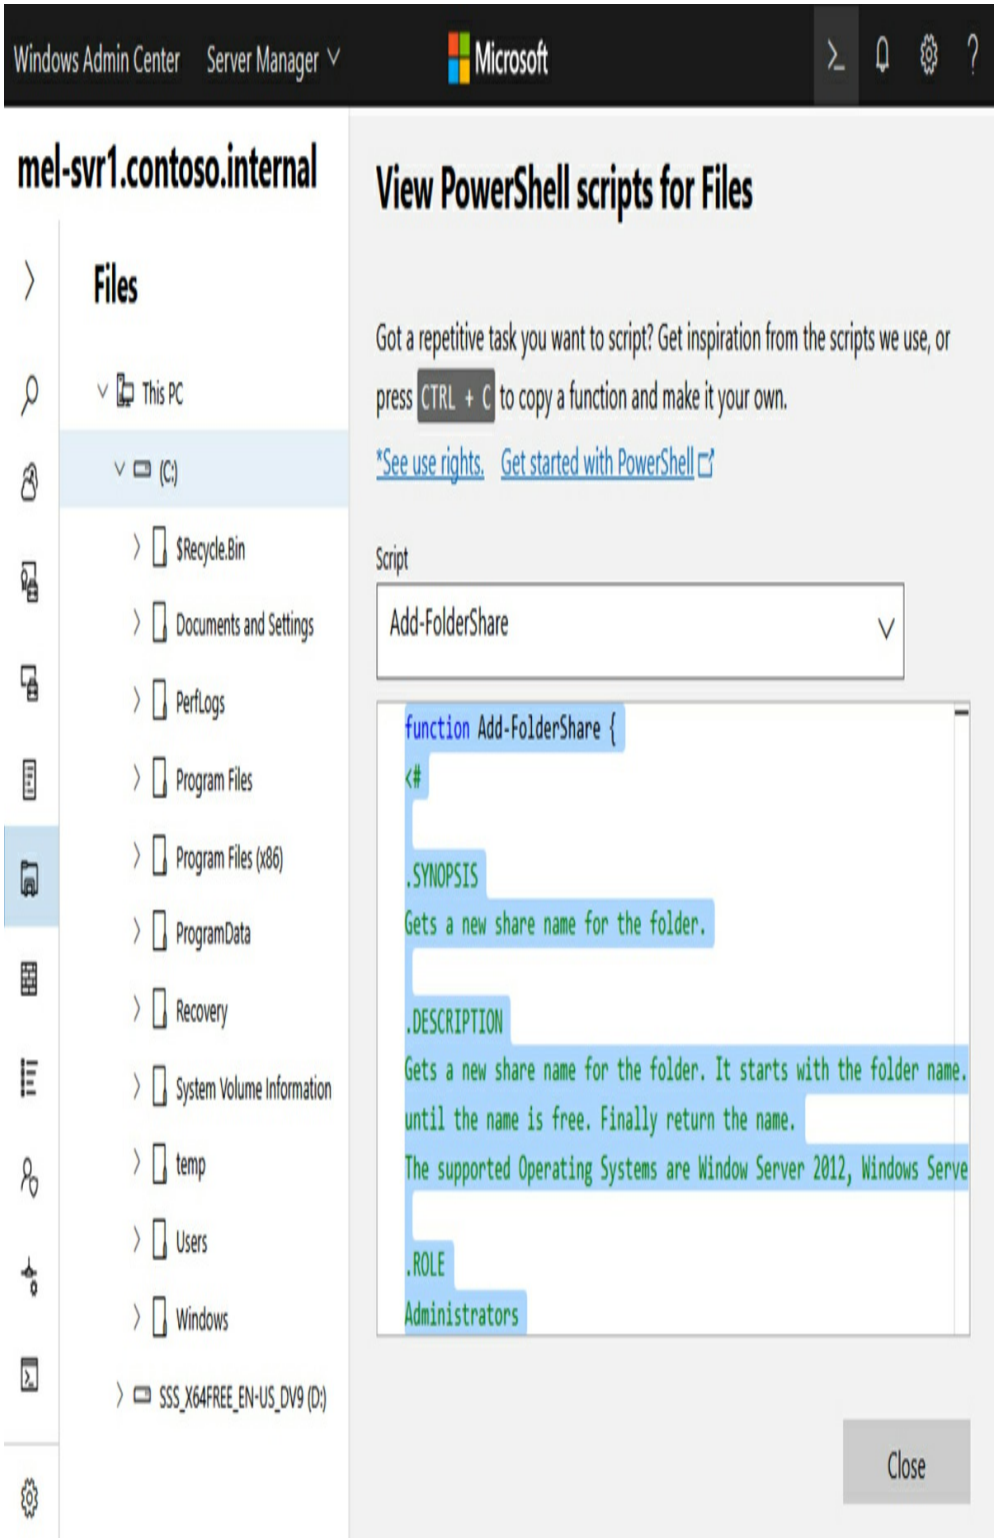

Windows Admin Center
Server Manager
Microsoft
mel-svr1.contoso.internal
Files
This PC
(C:)
$Recycle.Bin
Documents and Settings
PerfLogs
Program Files
Program Files (x86)
ProgramData
Recovery
System Volume Information
temp
Users
Windows
SSS_X64FREE_EN-US_DV9 (D:)
View PowerShell scripts for Files
Got a repetitive task you want to script? Get inspiration from the scripts we use, or press CTRL + C to copy a function and make it your own.
*See use rights.
Get started with PowerShell
Script
Add-FolderShare
function Add-FolderShare {
<#
.SYNOPSIS
Gets a new share name for the folder.
.DESCRIPTION
Gets a new share name for the folder. It starts with the folder name.
until the name is free. Finally return the name.
The supported Operating Systems are Window Server 2012, Windows Serve
.ROLE
Administrators
Close

Figure 1-6 Show script

REMOTE SERVER ADMINISTRATION TOOLS

The Remote Server Administration Tools (RSAT) are a set of consoles that you can install on a computer running Windows 10 that allow you to manage computers running Windows Server 2019. The RSAT consoles can also be installed on a computer running Windows Server 2016 and Windows Server 2019. If you want to install every RSAT console, issue this command:

Click here to view code image

```
Install-WindowsFeature -IncludeAllSubFeature RSAT
```

Installing the RSAT tools on a computer running Windows 10 involves downloading the latest version of the RSAT tools from Microsoft's website. You can determine the location of this download using your favorite search engine. A simpler method of installing the RSAT tools on Windows 10 involves running the following command from an elevated command prompt if Chocolatey is installed:

```
choco install -y rsat
```

Chocolatey is a third-party package manager for Windows operating systems. You will learn more about Chocolatey in Chapter 3, "Deployment and

configuration.”

Inside OUT

Windows 10 for Server 2019

While RSAT tools for Windows 10 can manage Windows Server 2012 R2, Windows Server 2012, and Windows Server 2008 R2, the RSAT tools packages for Windows 8.1 can only be used to manage servers running Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008; Windows Server 2016 cannot be run using the RSAT tools packages for Windows 8.1. Similarly, the RSAT tools for Windows 7 can only be used to manage servers running Windows Server 2008 R2 and Windows Server 2008, but they can’t be used to manage Windows Server 2012, Windows Server 2012 R2, and Windows Server 2016. Put another way, only the most recent version of the client operating system gets a version of the RSAT tools to manage the most recent version of the server operating system.

RSAT consoles

Table 1-2 provides a list of the RSAT tools that are available on Windows Server 2019. Most of these tools are available from the Tools menu of the Server Manager console.

Table 1-2 RSAT consoles

Console	Function
Active Directory Administrative Center	This is an advanced console for managing Active Directory users, computers, domains, and forests, as well as Dynamic Access Control and Authentication policies. It performs most of the functions from the traditional Active Directory management consoles available on earlier versions of the Windows Server operating system.

Technet24

Active Directory Domains and Trust	This is a console for configuring and managing trust relationships between domains and forests.
Active Directory Rights Management Services	This allows you to manage and configure Active Directory Rights Management Services.
Active Directory Sites and Services	This allows you to manage Active Directory Site configuration, including Global Catalog servers and universal group-membership caching.
ADSI Edit	ADSI Edit is an Active Directory object and attribute editor.
Certification Authority	Certificate Authority is used to manage Active Directory Certificate Services.
Cluster-	This is used to manage cluster-aware updating.

Aware Updating	
Component Services	This allows you to manage component services, view the event log, and manage services.
Computer Management	The Computer Management console is similar to the Computer Management console in Windows Server 2008 R2. It allows you to manage Tasks, Event Viewer, Shared Folders, Local Users and Groups, Performance, Device Manager, Disk Management, and Services.
Connection Manager Administration Kit	This allows you to create and deploy remote access connections.
Defragment and Optimize Drives	This allows you to manage the defragmentation and optimization of drives.
DFS Management	This allows you to manage the Distributed File System.

DHCP	This allows you to manage Dynamic Host Configuration Protocol servers.
Disk Cleanup	This console allows you to remove files and folders that are no longer necessary, such as old update and temporary files.
DNS	This console is where you configure and manage DNS servers.
Event Viewer	This console allows you to view and manage event logs.
Failover Cluster Manager	This allows you to configure and manage failover clusters.
Fax Service Manager	This allows you to configure and manage the fax service.
File Server Resource Manager	This console is for managing file servers, including file classification, storage reports, and quotas.
Group Policy Manag	Use this console for management of group policies, including running the resultant set of policy reports.

ement	
Hyper-V Manager	Use this console to configure and manage Hyper-V virtualization.
Internet Information Services (IIS) 6.0 Manager	This allows you to configure and manage version 6 of IIS.
Internet Information Services (IIS) Manager	This is where you configure and manage IIS, version 7 and later.
iSCSI Initiator	This console allows you to configure iSCSI initiator settings.
Local Securit	Here is where you configure and manage local security policy settings.

y Policy	
Network Load Balancing Manager	This is where you can configure and manage network load balancing.
Network Policy Server	Use this console to configure Network Policies.
ODBC Data Sources (32-bit)	This is where you manage 32-bit ODBC data sources.
ODBC Data Sources (64-bit)	Use this console to manage 64-bit ODBC data sources.
Online Responder Management	Online Responder Management is where you configure and manage Online Certificate Status Protocol arrays.

Performance Monitor	This console allows you to view performance information.
Print Management	You can configure and manage print servers from this console.
Remote Access Management	You can configure and manage remote access from this console.
Remote Desktop Services	Use this console to configure and manage Remote Desktop Services.
Resource Monitor	The Resource Monitor provides real-time monitoring of RAM, CPU, Disk, and Network utilization.
Routing and Remote Access	This allows you to configure and manage routing and remote access, including DirectAccess.

Services	This console is where you configure and manage services.
Services for Network File System (NFS)	From here, you can configure and manage the Network File System.
Shielding Data File Wizard	This console is where you manage shielded data files and policies.
System Configuration	System Configuration allows you to manage the system configuration.
System Information	Use this console to view system information.
Task Scheduler	Use the Task Scheduler to configure and manage scheduled tasks.

Template Disk Wizard	Use this wizard to configure and manage virtual hard disk templates for shielded VMs.
Volume Activation Tools	Use this console to configure and manage volume licensing.
Windows Deployment Services	Use this console to configure and manage Windows Deployment Services.
Windows Firewall with Advanced Security	This console allows you to configure and manage Windows Firewall with Advanced Security.
Windows Memory Diagnostics	The Windows Memory Diagnostics console is where you perform memory diagnostics to check whether the server memory is corrupted. A reboot may be required.

Windows Server Backup	This is where you configure and manage Windows Server Backup.
Windows Server Update Services	This allows you to configure and manage the Windows Server Update Services server.
WINS	The WINS console allows you to configure and manage a Windows Internet Naming Services server.

Server Manager console

Perhaps the most important RSAT console is the Server Manager console. This console opens automatically on a computer running Server 2019 if you install the server with the Desktop Experience option. Also, it will automatically prompt you to try out Windows Admin Center and provide you with a helpful link to download that tool.

A number of roles that had separate consoles in previous versions of Windows Server have that functionality included in the Windows Server 2019 Server Manager console. For example, many storage-

management tasks and IPAM tasks are completed in Server Manager rather than in individual role-related consoles.

Inside OUT

Server Manager and Windows Admin Center

When Windows Server 2016 was released, new GUI functionality was exposed through Server Manager rather than through new Microsoft Management Consoles. With the release of Windows Admin Center, Server Manager is unlikely to see any additional functionality added. It might even be that Windows Admin Center will replace the functionality of the Server Manager console before the WAC team gets around to replicating the functionality of older Microsoft Management Consoles, such as the Active Directory Users and Computers (ADUC).

When you launch Server Manager console, it will query the servers that you have added to the All Servers group to determine which roles are installed on those servers. Then it will provide a list of categories based on the roles detected on the servers. Figure 1-7 shows an All Servers group where the AD DS, DHCP, and DNS roles have been detected on servers MEL-DC, MEL-SERVERCORE, MEL-SVR1, and MEL-SVR2.

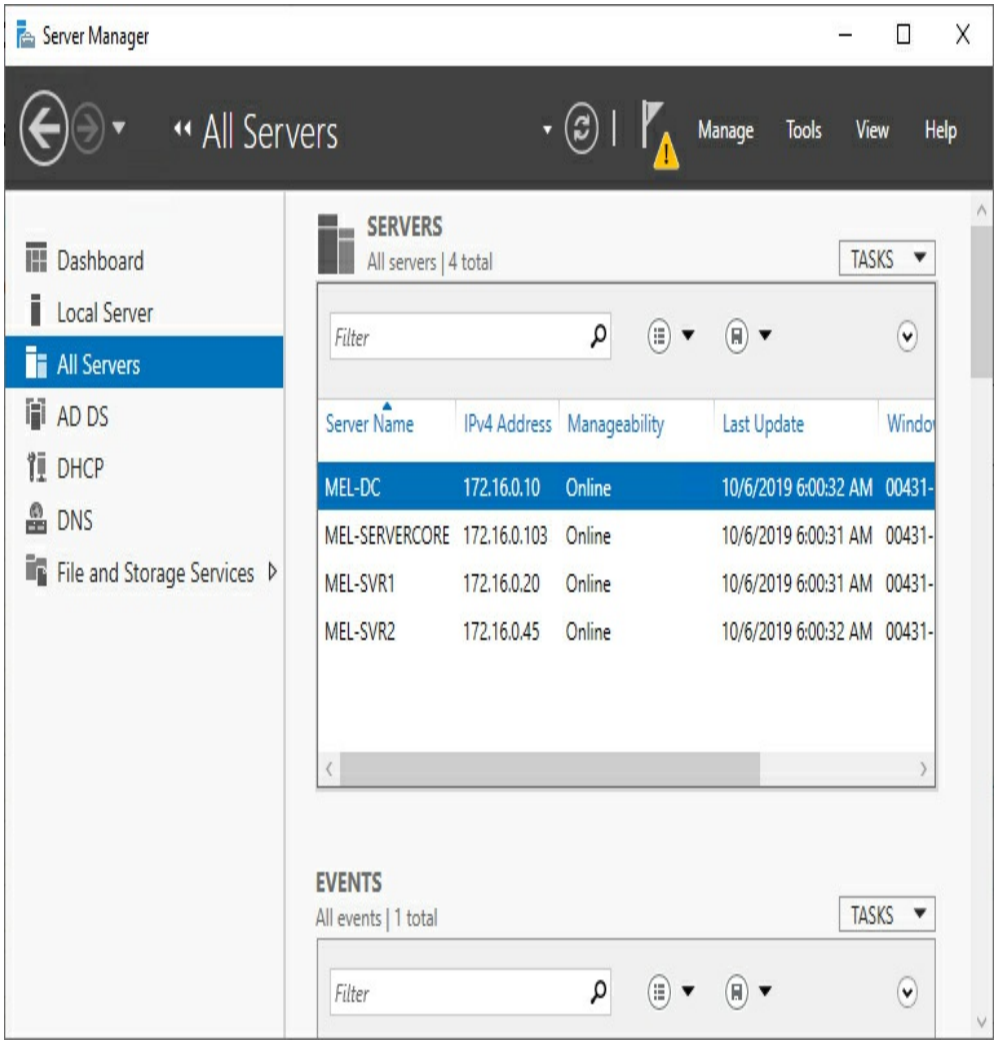


Figure 1-7 All Servers list

To add servers to the All Servers, group, right-click the All Servers node of the Server Manager console and click Add Servers. On the Add Servers dialog box, shown in Figure 1-8, select the servers using one of three methods: by querying Active Directory, searching on the basis of DNS name, or importing servers from a list.

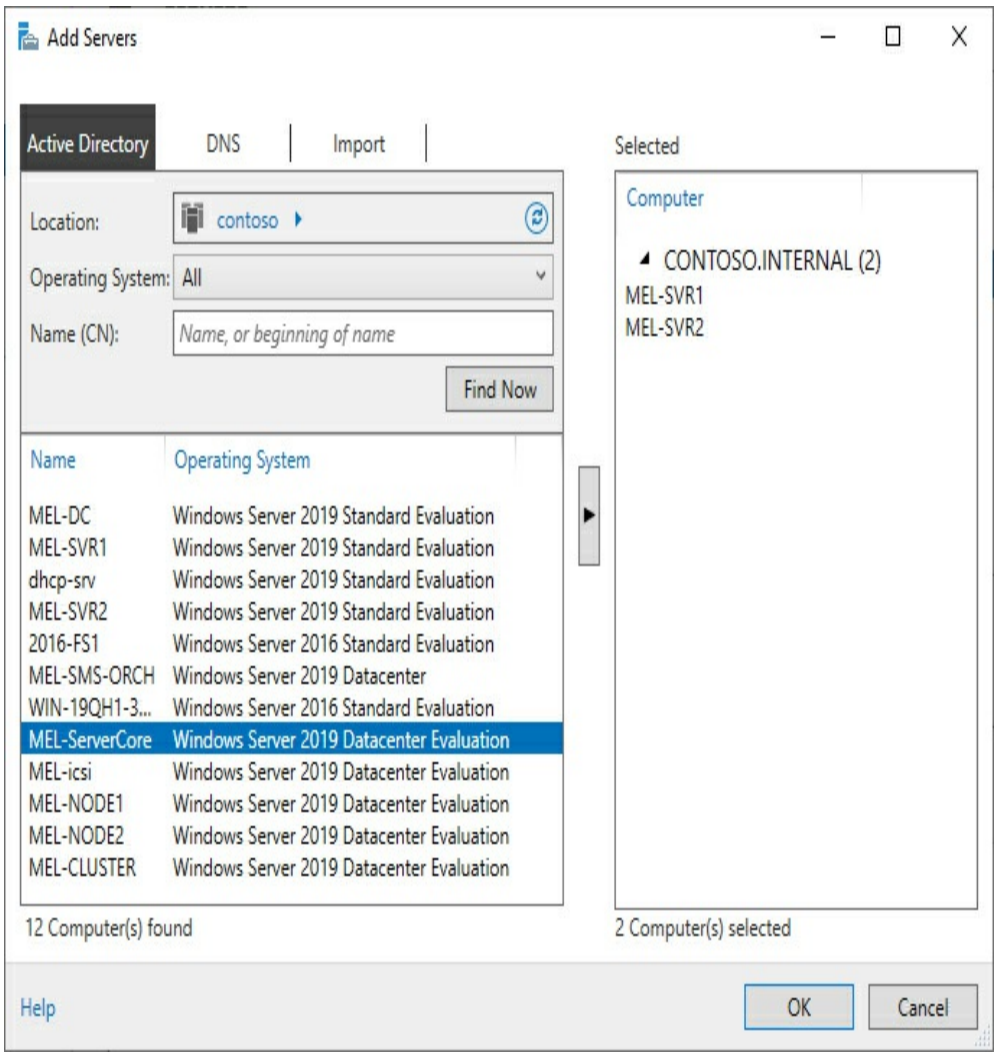


Figure 1-8 Add Servers

Once you have a list of roles in Server Manager, you can perform role-specific administration on each server by right-clicking that server. When you do this, the appropriate RSAT console will be launched with the target server in focus. Figure 1-9 shows MEL-DC selected in the list of DNS servers, which includes an option to launch the DNS Manager console. You can

also use the context menu—available by right-clicking—to do the following (where supported):

- Add roles and features to the target server
- Restart the target server
- Open the Computer Management console
- Open a Remote Desktop Connection session to the target server
- Open a Windows PowerShell session on the target server
- Configure NIC Teaming on the target server
- Manage the server using an alternate set of credentials

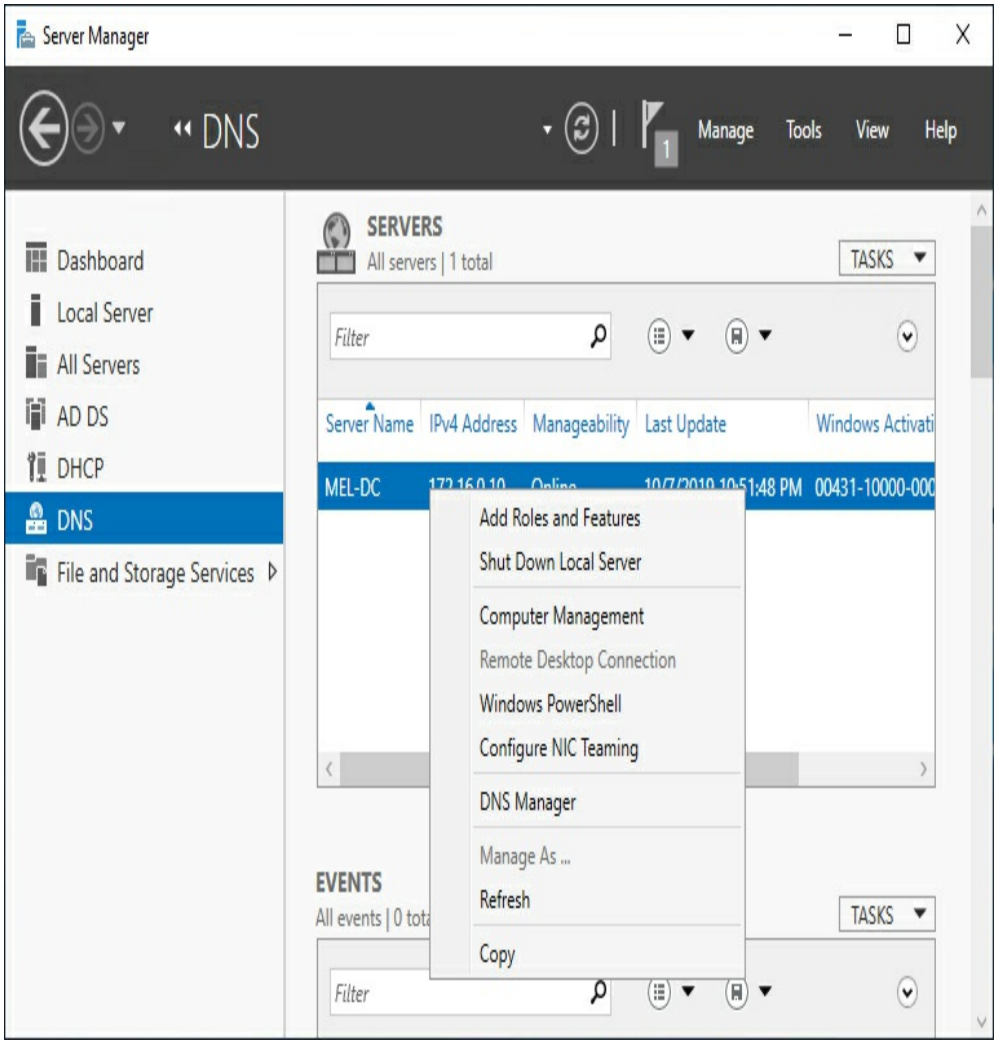


Figure 1-9 Server Manager individual server context menu

POWERSHELL

PowerShell is Microsoft's primary scripting, automation, and management tool. It's not hyperbole to suggest that the most important skill for Windows Server administrators going forward is using PowerShell. In

almost all cases, you can access greater functionality and settings through PowerShell than you can through Windows Admin Center or the RSAT consoles.

PowerShell includes a substantial amount of documentation explaining what each cmdlet can do and how you can do it. While you'll learn the names of the PowerShell cmdlets that are relevant to each role covered by each chapter of this book, once you know the name of the command you want to use to perform a task, you can use PowerShell's built-in help to teach you the precise details of how to use that cmdlet to perform that task. You can get help for each cmdlet by typing **help <cmdletname>**. For example, to get help with the *get-service* cmdlet, type **help get-service** into a PowerShell session.

> **Inside OUT**
>
> **Latest documentation**
>
> When you need help with a PowerShell cmdlet, the first thing you should do is to run the update-help cmdlet on an Internet-connected server. This will pull the most up-to-date PowerShell documentation from Microsoft's servers on the Internet down to your local server. This documentation is almost always the same as the cmdlet's documentation available on Microsoft's website, so it saves you a search with your favorite search engine.

Some important basic commands to get you started with PowerShell are listed in Table 1-3.

Table 1-3 PowerShell help commands

Cmdlet or command	Functionality
Update-Help	This command updates the help documentation for each cmdlet to the most recent version available from Microsoft.
Get-Command	This lists all PowerShell commands available on the computer.
Get-Command -Module <ModuleName>	This command allows you to get all cmdlets in a specific PowerShell module. For example, to view all the DNS server cmdlets, run the command *Get-Command -Module DNSServer*.
Get-Command -Noun <NounName>	Each PowerShell cmdlet is made up of a verb-noun combination. You can use this command to view all cmdlets associated with a particular noun. For example, to view all the cmdlets associated with the DNSServer noun, run this command: *Get-command -noun DNSserver*.
Help <cmdletname>	This command allows you to view the summary help documentation for a specific PowerShell cmdlet.

Help -detailed <cmdletname>	Use this command to view the detailed help documentation for a specific PowerShell cmdlet.
Help -examples <cmdletname>	This command allows you to view examples of how to use a specific PowerShell cmdlet to perform tasks.

Modules

Modules are collections of PowerShell cmdlets. In older versions of PowerShell, you needed to manually load a module each time you wanted to use one of its associated cmdlets. In Windows Server 2016 and Windows Server 2019, any module that is installed will load automatically when you try to run an associated cmdlet. Viewing cmdlets by module using the *Get-Command -Module <modulename>* cmdlet allows you to view just those cmdlets associated with a specific role or feature.

PowerShell Gallery

The PowerShell Gallery is a collection of modules published by the community that extend the functionality of PowerShell beyond what is available

with a default installation of Windows Server. Table 1-4 lists the commands that you can use to get started with the PowerShell Gallery.

Table 1-4 PowerShell Gallery basics

Command	Functionality
Find-Module -Repository PSGallery | out-host -paging	This will list the available modules in the PowerShell Gallery in a paged format. You'll be prompted to install the NuGetProvider to interact with the PowerShell Gallery.
Find-Module -Repository PSGallery -Name <ModuleName>	This will list the modules with a specific name. You can use wildcards. For example, to view all modules that start with the name *AzureRM*, run the command *Find-Module -Repository PSGallery -Name AzureRM**.
Install-Module -Repository PSGallery -Name <ModuleName>	This will install the *Modulename* module. For example, to install the *AzureRM* module, run the command *Install-Module -Repository PSGallery -Name AzureRM*.
Update-Module	This will update any module that you've installed using *Install-Module*.

Technet24

Get-InstalledModule	Use this command to view all modules installed from the PowerShell Gallery.

Remoting

By default, PowerShell remoting is enabled on Windows Server 2019, but also by default it requires a connection from a private network and an account that is a member of the local Administrators group. A remote session allows you to run commands on a remote computer in a manner similar to how you'd remotely enact commands when connected using an SSH session.

You can enter a remote session to a Windows Server 2019 computer that is a member of the same Active Directory forest by using the following commands, which will prompt you for the credentials that you will use to connect to the remote computer, as shown in Figure 1-10.

Click here to view code image

```
$cred = Get-Credential
Enter-PSSession -computername <computername> -Credential $cred
```



Figure 1-10 Get-Credential

Inside OUT

Get-Credential

The PowerShell command $variablename = Get-Credential will open a dialog box, which allows you to enter a username and password that will then be stored in an encrypted variable that you can use in future PowerShell commands within the same session.

You can enable PowerShell remoting using the *Enable-PSRemoting* cmdlet if it has been disabled. PowerShell remoting relies on WSMan (Web Server Management). WSMan uses port 5985 and can be configured to support TLS over port 5986.

Remote PowerShell connections are made to a defined endpoint, usually the default endpoint. The account used to make a connection to this remote endpoint must have local Administrator privileges on the target computer. With Just Enough Administration, which

you'll learn about in Chapter 19, "Hardening Windows Server and Active Directory," you'll learn about creating additional endpoints that allow connections to limited privilege restricted sessions.

To enable PowerShell remoting to computers that are not domain-joined, you need to configure the trusted hosts list on the client computer from which you want to establish the remote session. You do this on the client computer using the *set-item* cmdlet. For example, to trust the computer at IP address 192.168.3.200, you'll need to run this command:

Click here to view code image

```
Set-Item wsman:\localhost\Client\TrustedHosts -Value
192.168.3.200 -Concatenate
```

Once you've run the command to configure the client, you'll be able to establish a PowerShell remote session using the *Enter-PSSession* cmdlet. If you want more information about remoting, you can run the following command to bring up help text on the subject:

```
Help about_Remote_faq -ShowWindow
```

One-to-many remoting

PowerShell allows you to run one command against many machines, which is known as one-to-many remoting, or fan-out administration. You can use one-

to-many remoting to run the same command against any number of computers. Rather than signing in to each computer to check whether a particular service is running, you can use PowerShell remoting to run the same command that checks the status of the service against each computer within the scope of the command.

For example, you could use the following command to read a list of computers from a text file named *computers.txt*:

```
$Computers = Get-Content c:\Computers.txt
```

You could then use the following command to get the properties of the Windows Update service:

Click here to view code image

```
Invoke-Command -ScriptBlock { get-service wuauserv } -
computername $Computers
```

You can also use the *Invoke-Command* cmdlet to run a script from the local computer against a number of remote computers. For example, to run the script *FixStuff.ps1* against the computers in the file *computers.txt*, run this command:

Click here to view code image

```
$Computers = Get-Content c:\Computers.txt
Invoke-Command -FilePath c:\FixStuff.ps1
```

PowerShell ISE

PowerShell ISE (Integrated Scripting Environment) is a tool that is built in to the operating system on computers running Windows 10, and it is the graphical version of Windows Server 2019 that you can use to create, manage, and run PowerShell scripts and commands. If you install the Windows Server App Compatibility Features on Demand feature in Windows Server 2019, you can also deploy PowerShell ISE on a Server Core deployment.

PowerShell ISE includes a script pane and a command window, and it also provides you with the option to list the available cmdlets sorted by module as shown in Figure 1-11.

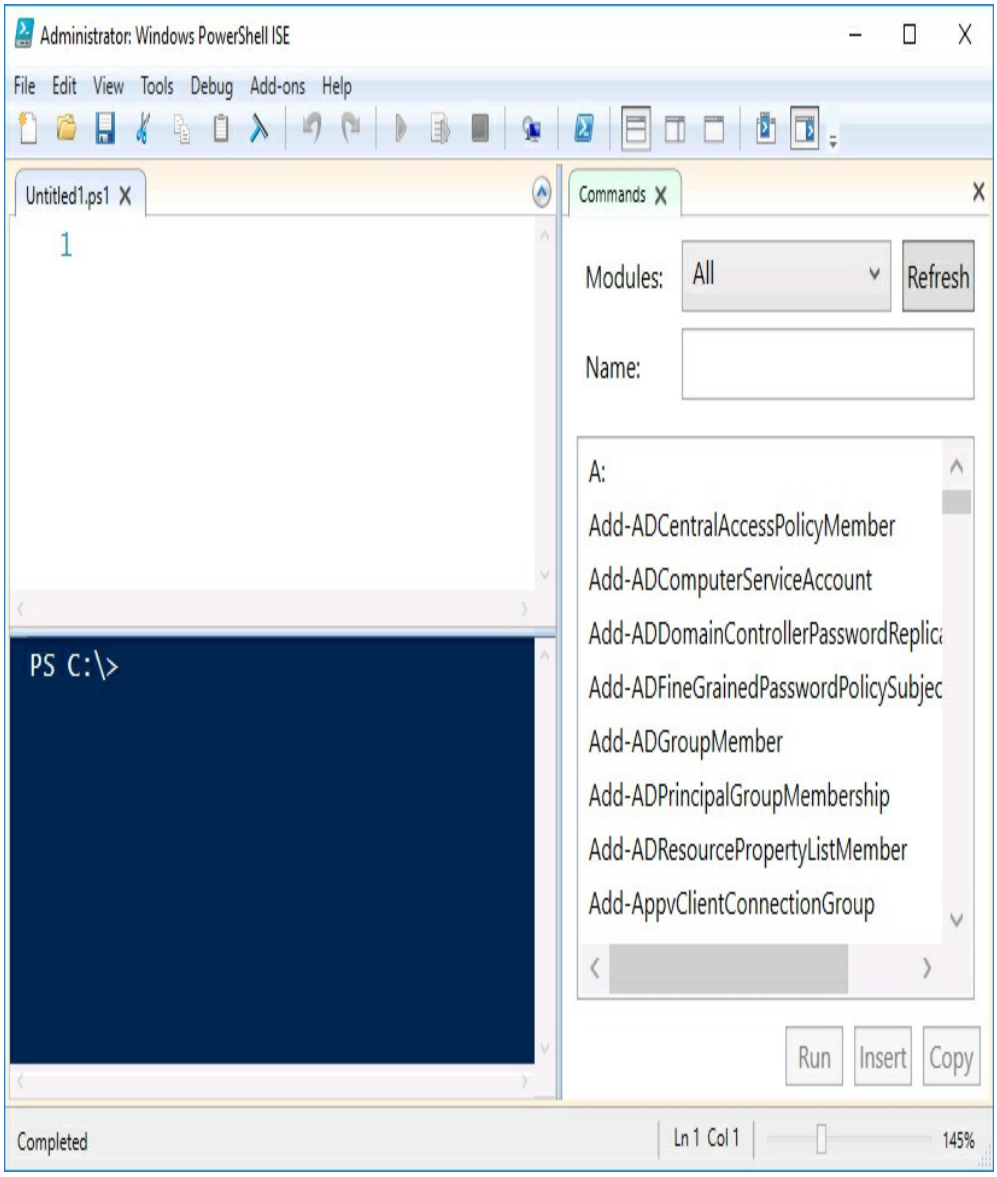


Figure 1-11 PowerShell ISE

PowerShell ISE includes the following functionality:

- **IntelliSense.** IntelliSense provides code completion and displays possible cmdlets, parameters, parameter values, files, and folders. Code completion functionality extends to the command window.
- **Color coding.** PowerShell code is color coded to assist code

Technet24

readability. This color coding extends to the command window.

- **Visual debugging.** Visual debugging allows you to step through PowerShell code and configure and manage breakpoints.
- **Brace Matching.** Brace matching allows you to locate matching braces to ensure that all opened braces are matched by closing braces.
- **Context-sensitive help.** Context-sensitive help allows you to view information about cmdlets, parameters, and values.
- **Run all code or just a selection.** Rather than running all the PowerShell code in the script pane, you can instead just highlight and run discrete lines of PowerShell code.

Remote tabs

Remote tabs allow you to establish remote PowerShell sessions using PowerShell ISE rather than using the *Enter-PSSession* PowerShell cmdlet. To initiate a remote session, click the New Remote PowerShell Tab item in the PowerShell ISE's File menu. You'll be prompted for the address of the computer to which you want to connect, as well as the name of the user account with which you want to make the connection, as shown in Figure 1-12.

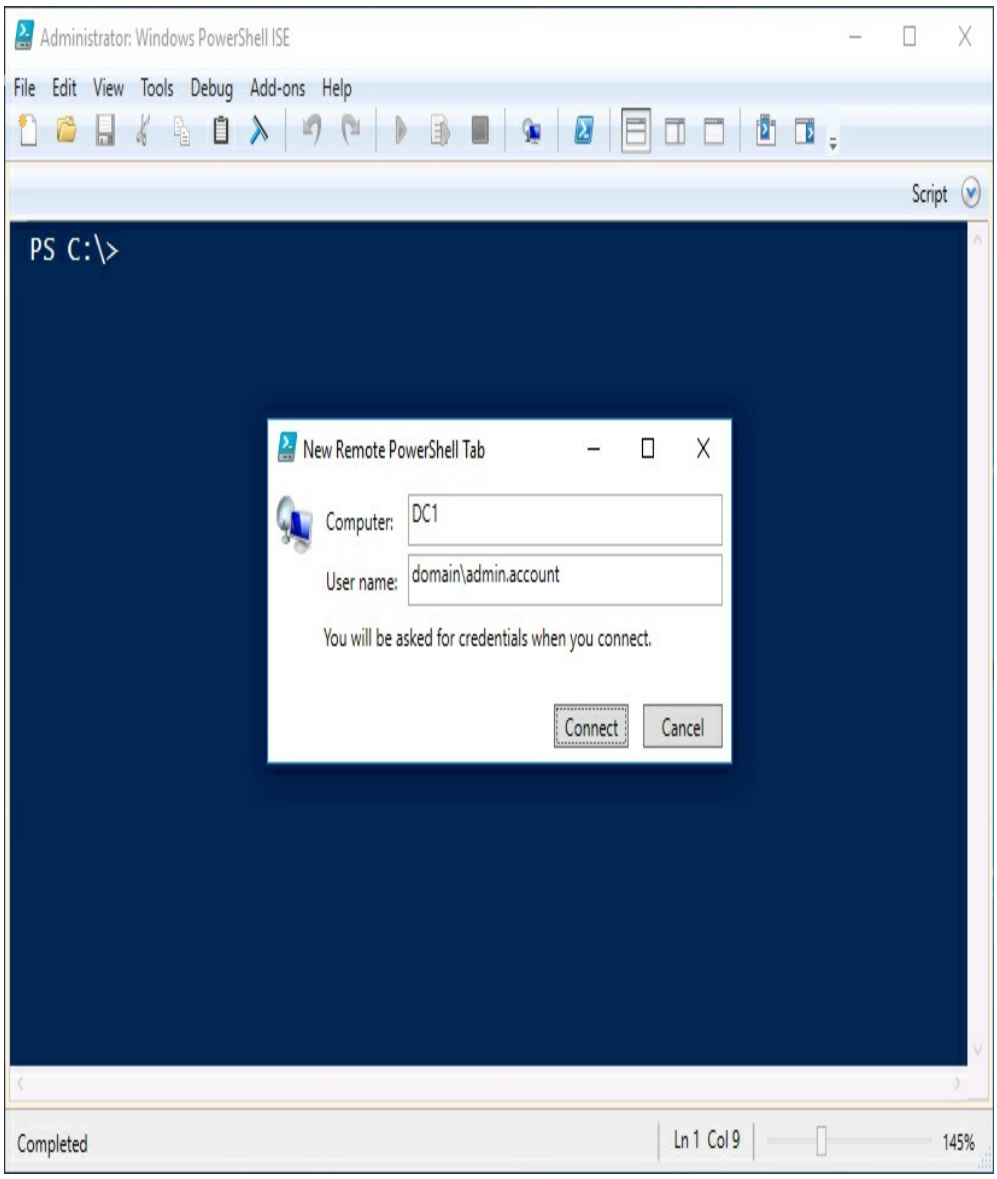


Figure 1-12 New Remote PowerShell Tab

Once a session is established, you can use the PowerShell ISE script pane to run scripts or sections of PowerShell code in the same manner that you would if you were using PowerShell ISE on the local computer.

Snippets

Snippets are commonly-used small sections of code that you can insert into your PowerShell scripts. Figure 1-13 shows an *If-Else* snippet already in the script pane, with the properties of a *do-while* snippet displayed. You can add snippets to the PowerShell ISE using the *New-Snippet* cmdlet.

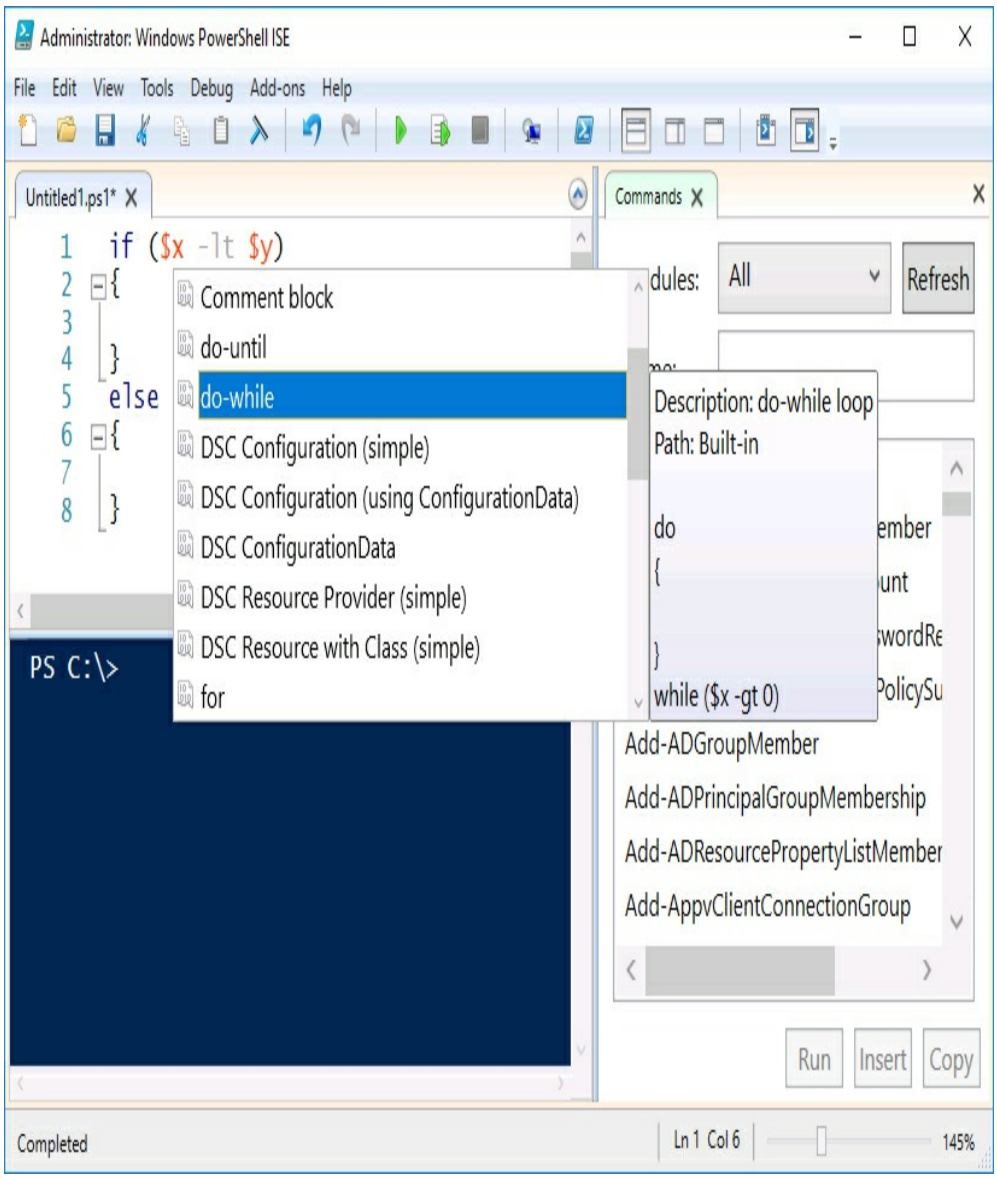


Figure 1-13 Snippets

Inside OUT

VSCode

Like the Server Manager console, PowerShell ISE is a bit of a dinosaur compared to newer options that can be downloaded from Microsoft's website. Visual Studio Code (VSCode) is a free Integrated Development Environment (IDE) that has extensions that can greatly assist in the creation and debugging of PowerShell scripts. Like Windows Admin Center, VSCode enjoys a rapid update cycle, which is

one of the reasons it isn't included with the operating system. Generally speaking, you should write your PowerShell scripts on your workstation using a tool like VSCode with all its cool PowerShell extensions. If you have to debug a PowerShell script on the server on which the script is running, then you should use PowerShell ISE.

PowerShell Direct

PowerShell Direct allows you to create a connection from a virtualization host to a VM running on that host. To use PowerShell Direct, you must have Hyper-V administrator privileges on the Hyper-V host computer and administrator credentials on the virtual machine to which you wish to connect. PowerShell Direct can allow you to make a connection to a virtual machine that may not otherwise be configured to support remote connections. It can also be used to make a connection to the VM when there is a network disruption.

PowerShell direct only works with Hyper-V hosts and guests running Windows Server 2019, Windows Server 2016, Windows Server Semi Annual Channel (SAC), and Windows 10. To get a list of the virtual machines on the host Hyper-V server, use the command *Get-VM*. To connect to the virtual machine using PowerShell Direct, use the *Enter-PSSession* cmdlet with the *-VMName* parameter.

To create a direct connection from the Hyper-V virtualization host to a computer running a supported version of Linux, ensure that the distribution's

virtualization extensions are installed and that an SSH server is installed. Once these steps have been taken, you can use the hvc.exe command to connect. For example, the following command allows you to make an SSH connection using the prime user account to a VM named Ubuntu-Server from an elevated PowerShell prompt on the virtualization host:

```
Hvc.exe ssh prime@ubuntu-server -v
```

REMOTE DESKTOP

Remote Desktop is the way that many administrators are likely to remotely perform one-off tasks on servers running the GUI version of Windows Server. While best practice is to use PowerShell or Windows Admin Center for remote administration, sometimes it is quicker to just establish a Remote Desktop session. This is because using Remote Desktop allows you to perform tasks on the remote server in a manner that appears similar to being directly signed in at the console.

Inside OUT

Remote Desktop and Azure

Remote Desktop remains the default method for connecting to Azure Infrastructure as a Service (IaaS) VMs running the Windows Server 2016 and Windows Server 2019 operating system when those virtual machines are deployed from the PowerShell Gallery. You'll learn more about running Windows Server on Azure in Chapter 17, "Azure IaaS and hybrid services."

By default, Remote Desktop is disabled on computers running Windows Server 2019. You enable it either through the Remote tab of the System Properties dialog box or by running the following PowerShell command:

Click here to view code image

```
Set-ItemProperty -Path
"HKLM:\System\CurrentControlSet\Control\Terminal Server" -Name
"fDenyTSConnections" –Value 0
```

You can make Remote Desktop connections to computers running the Server Core installation option if Remote Desktop is enabled.

By default, Remote Desktop Connection connects to Remote Desktop services on port 3389. When you enable Remote Desktop using the GUI, a remote desktop–related firewall is automatically enabled. If you enable Remote Desktop using PowerShell, you'll also need to manually enable a firewall rule to allow connections. You can do this using the following PowerShell command:

Click here to view code image

```
Enable-NetFirewallRule -DisplayGroup "Remote Desktop"
```

By default, the Allow Connections Only From Computers Running Remote Desktop With Network Level Authentication option is selected, as shown in Figure 1-14. Network Level Authentication requires that

a user be authenticated prior to the Remote Desktop session being established. Network Level Authentication is supported by the Remote Desktop Connection client, which is available on all Windows operating systems, but it might not be supported by third-party Remote Desktop clients.

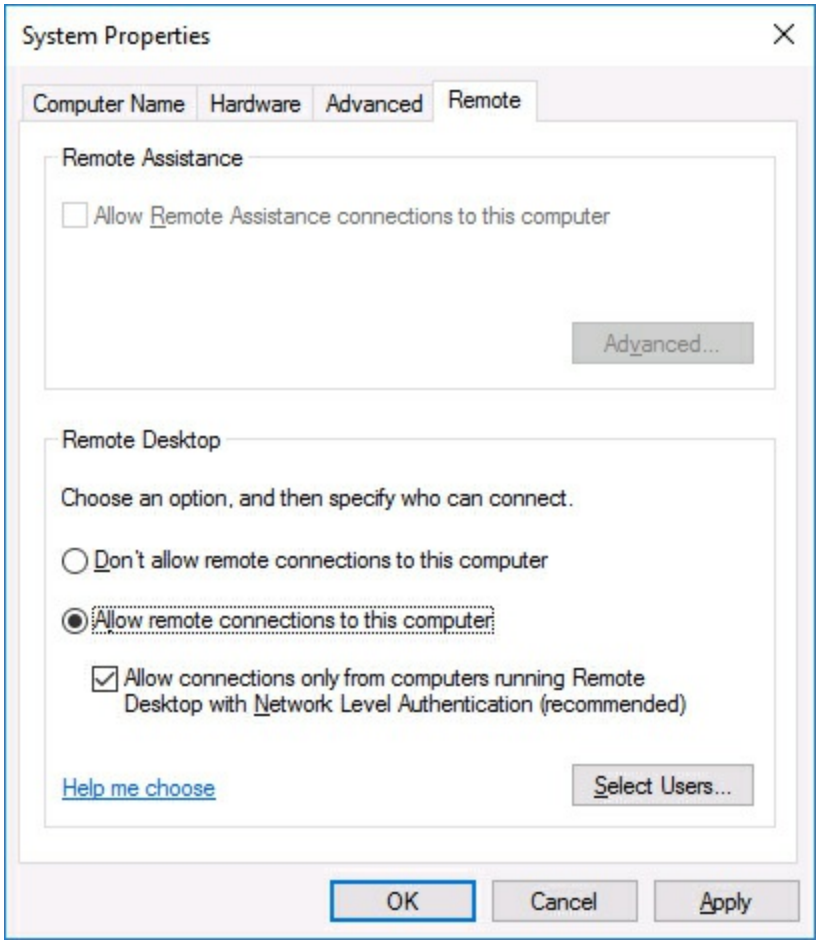


Figure 1-14 Allow Remote Desktop Connections To This Computer option

Only users who are members of the local Administrators

group and members of the local Remote Desktop Users group can make connections via Remote Desktop. If you want to grant someone permission to access the server without granting him or her full administrative privileges, add the user to the local Remote Desktop Users group.

Inside OUT

Just Enough Administration

Just Enough Administration (JEA) allows you to grant access—through a PowerShell endpoint—to only specific PowerShell cmdlets and parameters. This is a more effective way of giving a user limited administrative access than providing the user with Remote Desktop access to a server. You'll learn more about JEA in Chapter 19, "Hardening Windows Server and Active Directory."

You can map local volumes to a remote host in an active Remote Desktop Connection session by configuring the Local Resources And Devices setting on the Local Resources tab of the Remote Desktop Connection dialog box, as shown in Figure 1-15. While it is less effective over low-bandwidth connections, it can provide a simple way to transfer files from your client computer to a remote server instead of setting up FTP or another file transfer method.

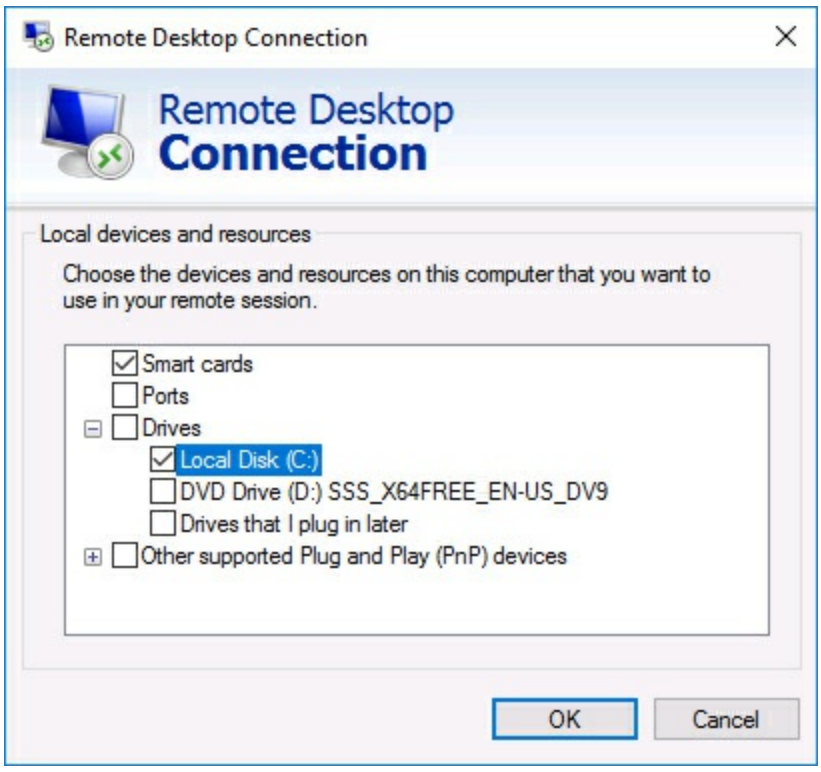


Figure 1-15 Remote Desktop Connection dialog box

SSH

Windows Server 2019 and the most recent version of Windows 10 include the ability to install both an SSH client and an SSH server. Although PowerShell now runs on Linux and Mac OS operating systems, many administrators in hybrid environments prefer to use SSH to make remote administration connections when using command-line interfaces.

To add an SSH client and SSH server to Windows Server 2019, use the following PowerShell commands:

Click here to view code image

```
Add-WindowsCapability -Online -Name OpenSSH.Server~~~~0.0.1.0
Add-WindowsCapability -Online -Name OpenSSH.Client~~~~0.0.1.0
```

Even though the version appears from the issued command to be version 0.0.1.0, the most recent supported version will be installed. You can check which version of SSH is installed using the command *SSH -V* from the command prompt.

Once you've added the capability, you need to do a few things to get the SSH server working before you're ready to go. If you're intending to use key-based—rather than password-based—authentication, you should also run the following command on the server to install an OpenSSH-related PowerShell module that includes tools to help you configure that functionality:

Click here to view code image

```
Install-Module -Force OpenSSHUtils -Scope AllUsers
```

Once you have run these commands, it is necessary to configure the disabled *ssh-agent* service to automatically start, and you also need to configure the sshd service to automatically start. You can do this by running the following PowerShell commands:

Click here to view code image

```
Set-Service -Name ssh-agent -StartupType 'Automatic'
Set-Service -Name sshd -StartupType 'Automatic'
```

The final step to get SSH running on Windows Server

2019 is to run the following commands to start the relevant services:

```
Start-Service ssh-agent
Start-Service sshd
```

If you do all this, you'll be able to connect using password-based authentication from an SSH client using the syntax:

```
ssh username@hostname_or_IP_address
```

If you're connecting to a domain account, use this format:

```
ssh username@domain@hostname_or_IP_address
```

The vast majority of people who use SSH use key-based authentication rather than password-based authentication. To get key-based authentication working with Windows Server 2019's SSH server, you'll need to perform additional steps. On the client, change to the *.ssh* directory, and run the *ssh-keygen* command, and accept the defaults. Doing this with the default values will create a public and private key. The private key will be *id_rsa*, and the public key will be *id_rsa.pub*.

The next thing you should do is add your private key to your Windows security context. You can do this by running the following three commands:

Click here to view code image

```
Set-Service ssh-agent -StartupType 'Automatic'
Start-service ssh-agent
Ssh-add ~\.ssh\id_rsa
```

Once you've done this, you'll want to deploy your public key to the Windows Server 2019 server with which you want to use SSH key-based authentication. To do this, perform the following steps (where *chancellor* is the name of the user account for which you're configuring SSH key-based authentication, and the IP address is the *IP address* of the server that you are configuring):

Click here to view code image

```
Ssh chancellor@172.16.0.15 mkdir c:\users\chancellor\.ssh\
Scp c:\users\chancellor\.ssh\id_rsa.pub
chancellor@172.16.0.15:C:\Users\Administrator\
.ssh\authorized_keys
```

To configure some of the permissions for the authorized keys file, you'll then need to run the following PowerShell command, located in that *OpenSSHUtils* PowerShell module mentioned earlier. You might even want to SSH to the server using password-based authentication to do this:

Click here to view code image

```
Repair-AuthorizedKeyPermission
C:\users\Chancellor\.ssh\authorized_keys
```

Because the PowerShell cmdlet doesn't entirely work as it should, you'll also need to run *Icacls authorized_keys /remove "NT SERVICE\sshd"* because *"NT*

SERVICE\sshd" should not have any permissions to the *authorized_keys* file. (If it does, key-based authentication doesn't seem to work.)

The final step requires you to edit the *C:\ProgramData\ssh\sshd_config* file and comment out the following lines (which are at the end of the file):

Click here to view code image

```
# Match Group administrators
#       AuthorizedKeysFile
__PROGRAMDATA__/ssh/administrators_authorized_keys
```

You might need to do this locally on the Windows Server 2019 box. Once this is done, you can restart the sshd service using the PowerShell command *restart-service sshd*. Then you will be able to connect from your client using key-based authentication.

More Info

OpenSSH on Windows Server

You can learn more about OpenSSH on Windows Server at https://docs.microsoft.com/en-us/windows-server/administration/openssh/openssh_server_configuration.

Chapter 2

Installation options

Windows Server 2019 editions

Windows Server servicing branches

Server Core

Server with Desktop Experience

Roles and features

Prior to installing Windows Server 2019, you need to make a number of decisions. The choices that you make will be determined by the role that you want the server to play in hosting your organization's workloads. You'll need to decide which edition of Windows Server to install. You'll also need to decide whether you'll install the Server with Desktop Experience or the Server Core version. You'll also need to determine whether you want to deploy a Windows Server version from the Long Term Servicing Channel or from the Semi Annual Channel. In this chapter, you'll learn about the different editions,

channels, and versions of Windows Server 2019.

WINDOWS SERVER 2019 EDITIONS

There are two primary editions of Windows Server 2019 and several minor editions of Windows Server 2019. The primary editions are:

- **Windows Server 2019 Standard.** Edition of Windows Server 2019 suitable for environments where most servers are deployed physically rather than as virtual machines. Licensed on a per processor core basis and requires Windows Server CALs.
- **Windows Server 2019 Datacenter.** Edition of Windows Server 2019 suitable for datacenters that include a high degree of virtualization. Licensed on a per processor core basis and requires Windows Server CALs.

The minor editions of Windows Server 2019 are:

- **Windows Server 2019 Essentials.** Edition of Windows Server for small businesses. Supports up to 25 users and 50 devices. Licensed on a per processor basis. Does not require CALs because of the user and device restriction.
- **Microsoft Hyper-V Server 2019.** A freely available hypervisor that does not require a license or CALs.

Table 2-1 lists the different limitations between Standard and Datacenter editions of Windows Server 2019.

Table 2-1 Edition limitations

Technet24

Item	Standard edition	Datacenter edition
Maximum users	Determined by available client access licenses	Determined by available client access licenses
Maximum SMB connections	16777216	16777216
Maximum routing and remote access connections	No limit	No limit
Maximum Internet Authentication Service connections	2147483647	2147483647
Maximum Remote Desktop Services connections	65535	65535
Maximum x64 processor sockets	64	64
Maximum number of cores	No limit	No limit
Maximum RAM	24 TB	24 TB
Included virtualization guest licenses	2 included VM licenses as well as Hyper-V host	Unlimited VM licenses as well as Hyper-V host

Included Windows container licenses	Unlimited	Unlimited
Included Hyper-V container licenses	2	Unlimited
Storage Replica instances	Single volume, up to 2 TB	Unlimited replica size and unlimited number of replicas.

For most organizations, the primary difference between the two servers will be the way that virtual machines are licensed. Datacenter edition allows you to host unlimited virtual machines, and Standard edition allows you to host two virtual machines before you'll have to purchase a license for each additional machine. Generally speaking, while you should check the costs yourself, if you're hosting more than a half dozen virtual machines running a Windows Server operating system, or you're using a half dozen Hyper-V containers, you'll be better off purchasing Datacenter edition.

Another difference in the roles supported between Windows Server 2019 Standard edition and Windows Server 2019 Datacenter edition is that Datacenter edition supports hosting shielded virtual machines, something that cannot be done with the Standard edition. Shielded virtual machines are BitLocker-encrypted virtual machines that can't be tampered with

by a virtualization administrator.

Another difference is that Datacenter edition supports inherited activation as a host or guest virtual machine. Standard edition only supports inherited activation when hosted as a guest virtual machine on a Hyper-V host running Datacenter edition. Inherited activation means that activation occurs automatically based on the Windows Server edition. This can save quite a bit of time if you are hosting a large number of Windows Server 2019 virtual machines, as you don't need to muck around with KMS servers.

WINDOWS SERVER SERVICING BRANCHES

Windows Server 2019 supports two separate servicing branches: the Long Term Servicing Channel (LTSC) and the Semi Annual Channel (SAC). These require separate installation media, and you won't be able to install a SAC release from LTSC media or an LTSC release from SAC media.

Long Term Servicing Channel

LTSC versions are the versions of Windows Server that people usually think of as the main versions of Windows Server. Windows Server 2016 and Windows Server 2019 are LTSC versions, with the next LTSC version of Windows Server likely to be named either

Windows Server 2021 or Windows Server 2022. LTSC versions of Windows Server can be installed either as a Server with Desktop Experience or in the Server Core configuration.

LTSC provides five years of mainstream support with five years of extended support. This is similar to the support options provided for previous versions of Windows Server, such as Windows Server 2012 R2. Most organizations are likely to deploy this servicing branch. For Windows Server 2019, extended support will end January 9, 2029. At present, Microsoft is offering an additional three years of extended security updates to customers willing to pay an annual fee for Windows Server 2008 and Windows Server 2008 R2, which will fall out of support in January 2020. Something similar may be available for Windows Server 2019 to extend the availability of security updates past 2029, but that's pure speculation.

LTSC is the only servicing branch that is available to servers installed using Server with Desktop Experience installation options.

Semi Annual Channel

SAC releases are Server Core only and are generally released twice a year. SAC releases are named with the year and month of release. For example, Windows Server version 1809 was released in September 2018,

and Windows Server 1903 was released in March 2019. There may even be a Windows Server 2003 SAC release in March 2020 and a Windows Server 2009 release in September 2020, which might confuse some people who don't understand the new naming scheme and suspect that Microsoft is rereleasing the "classics."

SAC releases include new features and capabilities that have been developed since the most recent LTSC release, as well as the features and capabilities that were present in that last LTSC release. Essentially, the SAC release Windows Server 1903 is a Server Core version of Windows Server 2019 with a few new features and capabilities, but which will only be supported by Microsoft for 18 months rather than the usual 10 years.

SAC releases are appropriate for organizations that have automated operating system deployment pipelines and that are comfortable with treating servers as "cattle" and not "pets." In theory, if your organization is using an SAC release, you're comfortable with the 18-month support window because you've implemented enough automation that the current instances of Windows Server are automatically replaced whenever the next SAC release is released by Microsoft. If you work for an organization that still has Windows Server 2003 on its subnet, then you probably want to give SAC releases of Windows Server a pass.

For the most part, features and capabilities that are available within an SAC release will be integrated into the next LTSC release. Sometimes, features within an SAC release don't make it into the next LTSC release, but that's probably because they weren't quite ready to be put into a version of the operating system that will be supported for the next decade. Occasionally, new features will turn up in the LTSC release; for example, Samba migrations for Storage Migration Services appeared with the August 2019 cumulative update, but this will be the exception rather than the rule.

More Info

Servicing Channels

You can learn more about Windows Server servicing channels at https://docs.microsoft.com/en-us/windows-server/get-started-19/servicing-channels-19.

Insider Preview Builds

Windows Server Insider Preview builds are prerelease versions of the software that are similar to the old Release Candidate or Beta versions. Insider Preview Builds are for testing purposes only and should not be used in production environments. New features appear in Insider Preview Builds before appearing in the SAC releases. Features that appear in the SAC releases will generally then appear in the LTSC releases.

Unlike SAC releases, whether Server with Desktop

Experience is included with Insider Preview Builds depends on the Preview Build. Before the release of the next LTSC build, there likely will be Insider Preview Builds that include the option to install with the GUI. So, the Insider Preview Build of Windows Server 2021/2022 will allow Server with Desktop Experience deployment, but a hypothetical Insider Preview build of Windows Server 2203 will not.

Inside OUT

At your own risk

Insider Preview Builds are very much a “run this at your own risk” proposition. Some builds will work perfectly well, and some will be less reliable than a teenager’s promises to do the dishes.

More Info

Insider Previews

You can learn more about Windows Server Insider Preview Builds and sign up for the program at https://www.microsoft.com/en-us/software-download/windowsinsiderpreviewserver.

SERVER CORE

Server Core is the default installation option for Windows Server 2019. Because Server Core has fewer components, it doesn’t need to have software updates applied as often as the Server with GUI installation option does. Because components such as the built-in web browser have been removed, Server Core is less

vulnerable to malware than the Server with Desktop Experience installation option. Additionally, because it doesn't require all the components of the version that includes the GUI, it has a lower resource footprint.

The nature of Server Core is not explicitly called out by the Windows Server installation wizard shown in Figure 2-1. The version listed as Desktop Experience is the Server with the Desktop Experience option. The Standard version is Server Core. Unless you are paying attention during a manual deployment, it's possible that you'll accidentally deploy Server Core rather than Server with Desktop Experience. I've done that more times than I can count, so you've been warned.

Technet24

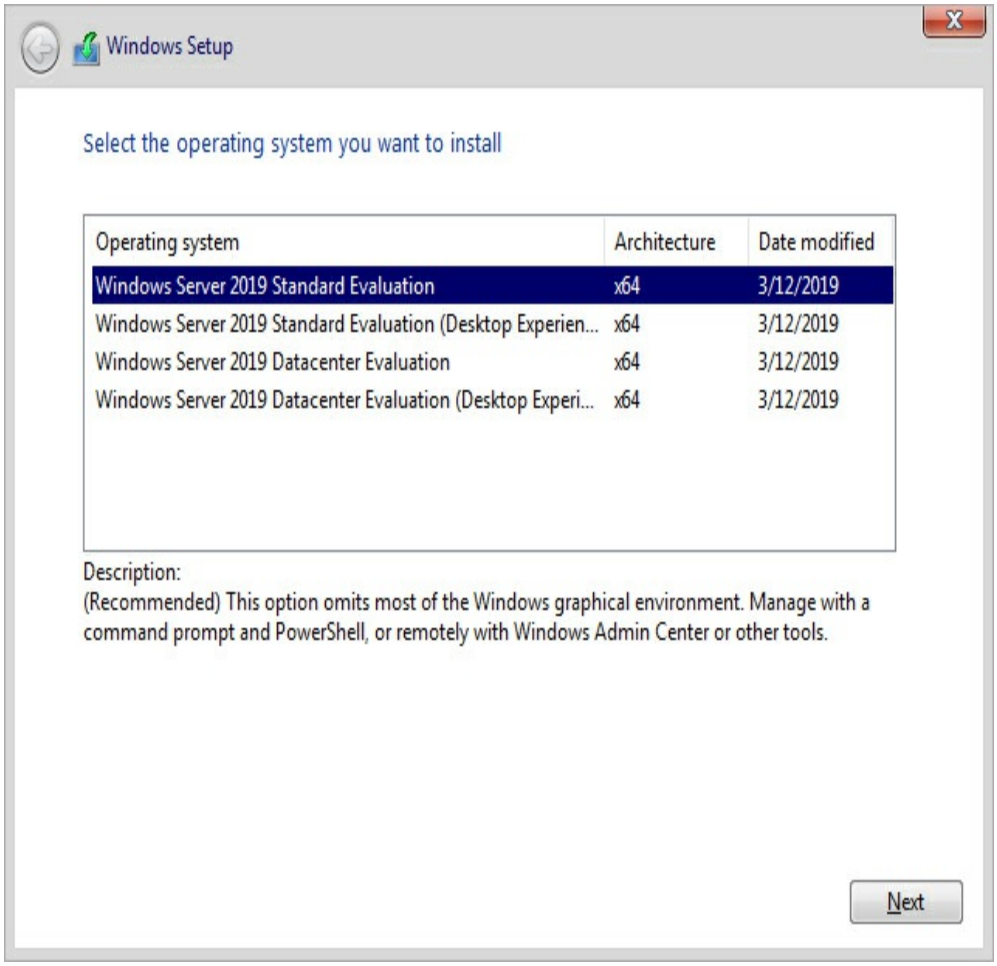


Figure 2-1 Windows Server 2019 installation options

Server Core interface

The Server Core interface is a command prompt. As Figure 2-2 shows, to interact with the command prompt, you must press Ctrl+Alt+Del to unlock it. You can then sign in with a domain-based or local administrator account.

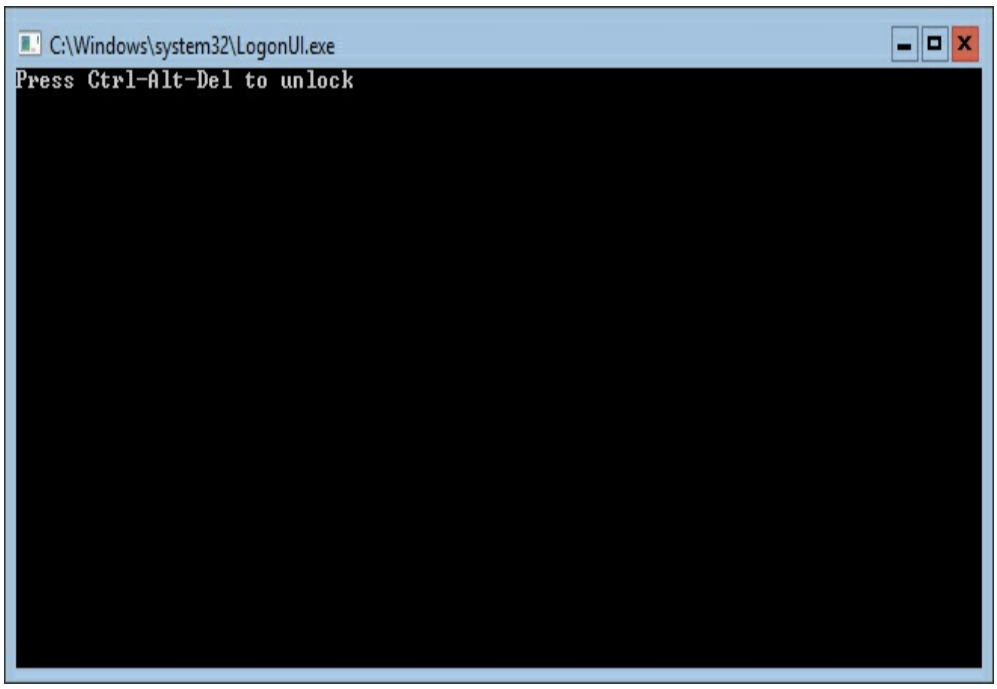


Figure 2-2 Server Core interface

Once you have signed on, you can enter a local PowerShell session by typing **PowerShell.exe**. While Server Core is primarily a command-line environment, there are some graphical tools that can be launched from the command prompt or Task Manager. These include:

- **Task Manager.** Task Manager functions in the same way as Task Manager on Server with Desktop Experience or Windows 10, and it can use Task Manager to run tasks by selecting the Run New Task item from the File menu.
- **Notepad.** You can launch Notepad.exe to edit and view the contents of text files.
- **MSInfo32.exe.** This allows you to view information about the installation.

Technet24

- **Regedit.exe and Regedt32.exe.** These allow you to edit the registry on a Server Core computer.
- **TimeDate.cpl.** This opens the Time And Date control panel.
- **Intl.cpl.** This opens the Regional Settings control panel.
- **Iscsicpl.exe.** This opens the iSCSI Initiator Properties control panel, and allows you to connect to shared storage over iSCSI.

Installing Server Core App Compatibility Features on Demand, which you will learn more about later in this chapter, allows you to access even more GUI tools when signed on directly to the Server Core desktop.

> **Inside OUT**
>
> **Oops**
>
> If you type **Exit** at the Server Core command prompt, the command prompt closes. While I've seen some administrators reboot their Server Core computers to get it back, there is an easier way. If you do accidentally close your command prompt, press Ctrl-Alt-Del, select Task Manager, click More Details, click File, click Run New Task, and then type **cmd.exe**. This will open a command prompt without requiring you to log out or restart the computer.

Server Core roles

Server Core supports the following roles, which you can install using the *Add-WindowsFeature* PowerShell cmdlet, by using the Add Roles And Features Wizard, available in the Server Manager console, or by using Windows Admin Center from a remote computer, as shown in Figure 2-3.

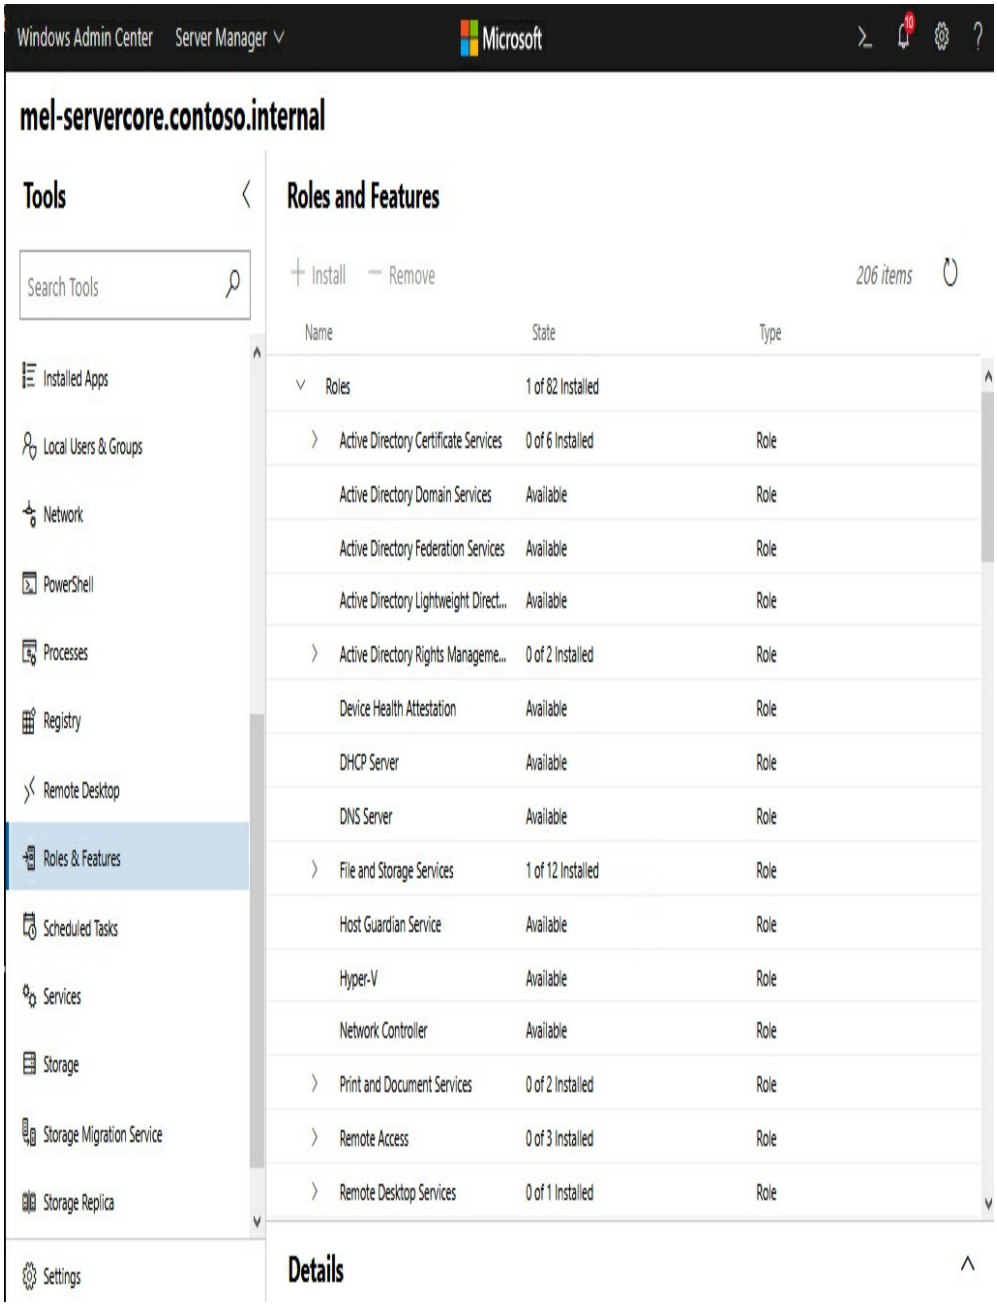


Figure 2-3 Server Core Roles

You can add the following features using this blade of the console:

- Active Directory Certificate Services
- Active Directory Domain Services
- Active Directory Federation Services
- Active Directory Lightweight Directory Services (AD LDS)
- Active Directory Rights Management Services
- Device Health Attestation
- DHCP Server
- DNS Server
- File and Storage Services
- Host Guardian Service
- Hyper-V
- Network Controller
- Print and Document Services
- Remote Access
- Remote Desktop Services
- Volume Activation Services
- Web Server (IIS)
- Windows Deployment Services
- Windows Server Update Services

Server Core also supports the following features, as shown in Figure 2-4.

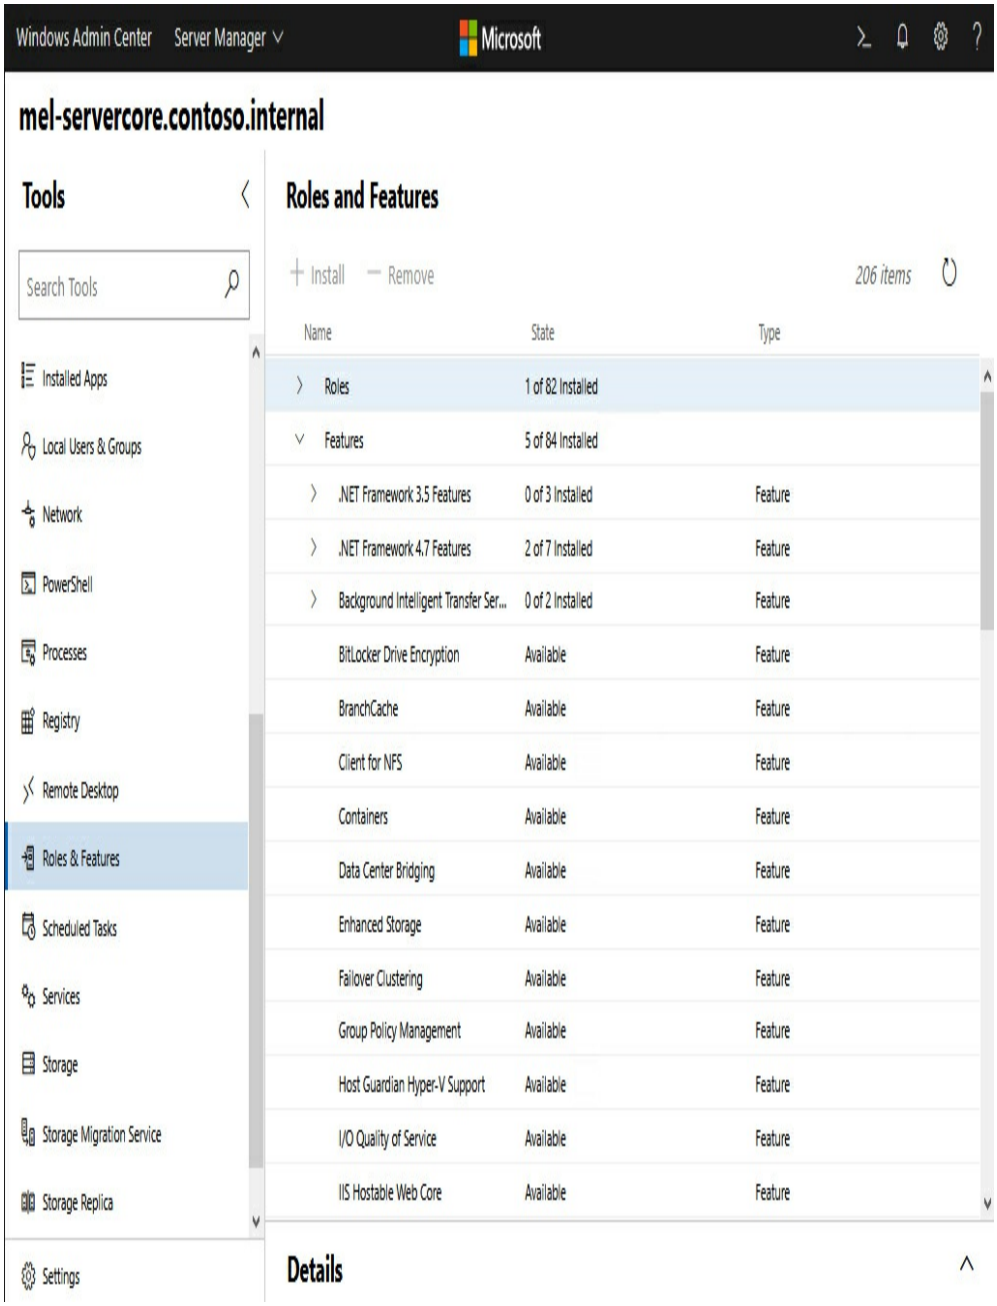


Figure 2-4 Server Core features

- .NET Framework 3.5 Features
- .NET Framework 4.7 Features

- Background Intelligent Transfer Service (BITS)
- BitLocker Drive Encryption
- BranchCache
- Client for NFS
- Containers
- Data Center Bridging
- Enhanced Storage
- Failover Clustering
- Group Policy Management
- Host Guardian Hyper-V Support
- I/O Quality of Service
- IIS Hostable Web Core
- IP Address Management (IPAM) Server
- iSNS Server service
- Management OData IIS Extension
- Media Foundation
- Message Queuing
- Multipath I/O
- Network Load Balancing
- Network Virtualization
- Peer Name Resolution Protocol
- Quality Windows Audio Video Experience
- Remote Differential Compression
- RPC over HTTP Proxy
- Setup and Boot Event Collection

- Simple TCP/IP Services
- SMB 1.0/CIFS File Sharing Support
- SMB Bandwidth Limit
- SNMP Service
- Software Load Balancer
- Storage Replica
- Telnet Client
- VM Shielding Tools for Fabric Management
- Windows Defender Antivirus
- Windows Internal Database
- Windows PowerShell
- Windows Process Activation Service
- Windows Server Backup
- Windows Server Migration Tools
- Windows Standards-Based Storage Management
- Windows subsystem for Linux
- WinRM IIS Extension
- WINS Server
- WoW64 Support

While it is possible to install and manage these roles using PowerShell, either directly or through a remote PowerShell session, most administrators are likely to deploy and manage these roles and features using Windows Admin Center.

As you learned in Chapter 1, Microsoft encourages

Technet24

organizations to manage all servers remotely. It doesn't really matter that Server Core doesn't have a GUI because you'll almost never be logging on directly; instead, you will be connecting using remote administration tools.

Inside OUT

Locked In

Windows Server 2012 and Windows Server 2012 R2 allowed you to switch between the Server Core and Server with Desktop Experience installation options. This gave administrators the ability to deploy the Server with Desktop Experience and then pare it back to Server Core if they found that they didn't need all the extra components included with Server with Desktop Experience. This gave server administrators who deployed the Server Core option an escape hatch. If they found that they couldn't run a specific application on Server Core, they could always bump it up to Server with Desktop Experience. Windows Server 2019 doesn't give you that option. The option you choose at deployment is the one you'll be stuck with. This makes sense in terms of Microsoft's philosophy around deployment, which is that if you want to make a substantive change, you should redeploy from scratch rather than reengineer the existing deployment. However, it probably means that fewer administrators will deploy Server Core because they know that applications will always work on Server with Desktop Experience, but they can't be sure that they might work with Server Core.

Server Core App Compatibility Features on Demand

One of the big challenges many organizations had with deploying Server Core is that many applications that seemingly didn't require GUI components ended up having some form of dependency on the GUI. This meant that many people who wanted to deploy Server Core were unable to do so because a dependency either blocked installation or rendered the application non-

functional.

App Compatibility Features on Demand (FOD) increases the compatibility of the Server Core installation option for a large number of applications. It also provides additional diagnostic tools for troubleshooting and debugging operations including:

- Microsoft Management Console (*mmc.exe*)
- Event Viewer (*Eventvwr.msc*)
- Performance Monitor (*PerfMon.exe*)
- Resource Monitor (*Resmon.exe*)
- Device Manager (*Devmgmt.msc*)
- File Explorer (*Explorer.exe*)
- PowerShell Integrated Scripting Environment (*PowerShell_ISE.exe*)
- Disk Management (*diskmgmt.msc*)
- Failover Cluster Manager (*CluAdmin.msc*)
- Microsoft Edge browser (Chromium version)

If the computer has connectivity to Windows Update, you can perform an online installation of the Server Core App Compatibility FOD by running the following command from an elevated PowerShell session and then restarting the computer:

Click here to view code image

```
Add-WindowsCapability -Online -Name
ServerCore.AppCompatibility~~~~0.0.1.0
```

If the Server Core instance doesn't have access to Windows Update, you can download the Features on Demand ISO from Microsoft's website and deploy these features by running the following commands; the first mounts the ISO, and the second installs FOD.

Click here to view code image

```
Mount-DiskImage -ImagePath X:\ISO_Folder\ISO_filename.iso

Add-WindowsCapability -Online -Name
ServerCore.AppCompatibility~~~~0.0.1.0 -Source
<Mounted_Server_FOD_Drive> -LimitAccess
```

More Info

Server Core Features on Demand

You can learn more about Server Core Features on Demand at https://docs.microsoft.com/en-us/windows-server/get-started-19/install-fod-19.

When to deploy Server Core

The key to deploying Server Core in your organization is determining whether it is suitable for a specific workload. Server Core is perfect for infrastructure-type roles such as Domain Controller, DNS server, DHCP server, and file server. Where Server Core may be less suitable is in hosting applications that have complex dependencies.

If you are thinking about deploying Server Core, you should perform tests to verify that each workload you intend to host on Server Core functions as expected.

Even though Server Core is Microsoft's current preferred deployment option, third-party server-hosted applications often have dependencies that require them to be installed on the server with the GUI installation option.

Inside OUT

Wget

Server Administrators often need to download files from a network or Internet location to a server computer to perform a task. This can get especially complicated if the server is on the perimeter network, meaning that you likely can't directly copy the file or files that you need from a network share. Server Core doesn't have a browser, which means that you can't use Internet Explorer to retrieve those files. What you can do is use the PowerShell wget command (which is just an alias to the Invoke-WebRequest cmdlet) with the UseBasicParsing parameter to download a file directly from a location on the Internet to the computer file system on which you are running the command. For example, to download the file https://www.contoso.com/folder/file.zip, use the command wget -UseBasicParsing https://www.contoso.com/folder/file.zip file.zip

Server Core can be installed from the Windows Server 2019 installation media, or it can be deployed using a variety of methods from the *install.wim* file that is located on the installation media. You'll learn more about deployment and configuration in Chapter 3, "Deployment and configuration."

Inside OUT

What happened to Nano?

Windows Server 2016 shipped with another installation option known as Nano Server. Nano Server was a very small version of the server operating system that only supported remote management and a limited number of roles. For a variety of reasons that Microsoft's Jeffrey Snover may (or may not) tell you about if you buy him some cake, Microsoft decided to discontinue Nano Server as an installation

option you could deploy on physical hardware or as a virtual machine. Like the infamous parrot in a Monty Python sketch, Nano Server isn't dead, but it lives on as a container image option for Windows Server containers.

SERVER WITH DESKTOP EXPERIENCE

Server with Desktop Experience, also known as Server with a GUI, is the version of Windows Server that IT pros who have been using and managing the product since the 1990s and 2000s are likely most familiar with. Server with Desktop Experience comes with a full desktop environment, including a web browser. You can also install software on Server with Desktop Experience that you would normally install on a computer running the Windows 10 operating system.

You are most likely to deploy Server with Desktop Experience if you need to host Remote Desktop Session Host servers. Additionally, you'll deploy Server with Desktop Experience if you need to host an application that has dependencies that are unavailable on Server Core.

Microsoft's recommendation is that you deploy Windows Server 2019 with the Desktop Experience only in situations where it is impractical to deploy Server Core. While many administrators who have worked with Windows Server for many years will choose the Server with Desktop Experience installation option out of

habit, deploying this version of Windows Server should be the exception rather than the rule.

ROLES AND FEATURES

If you're new to Windows Server 2019, you're likely to use the Add Roles And Features Wizard, shown in Figure 2-5, to add and remove roles and features from a computer running the Server with Desktop Experience option. While you can use this feature of the Server Manager console locally or remotely, you can only use it to manage roles and features on one server at a time. Windows Admin Center has a similar one-server-at-a-time limitation.

Technet24

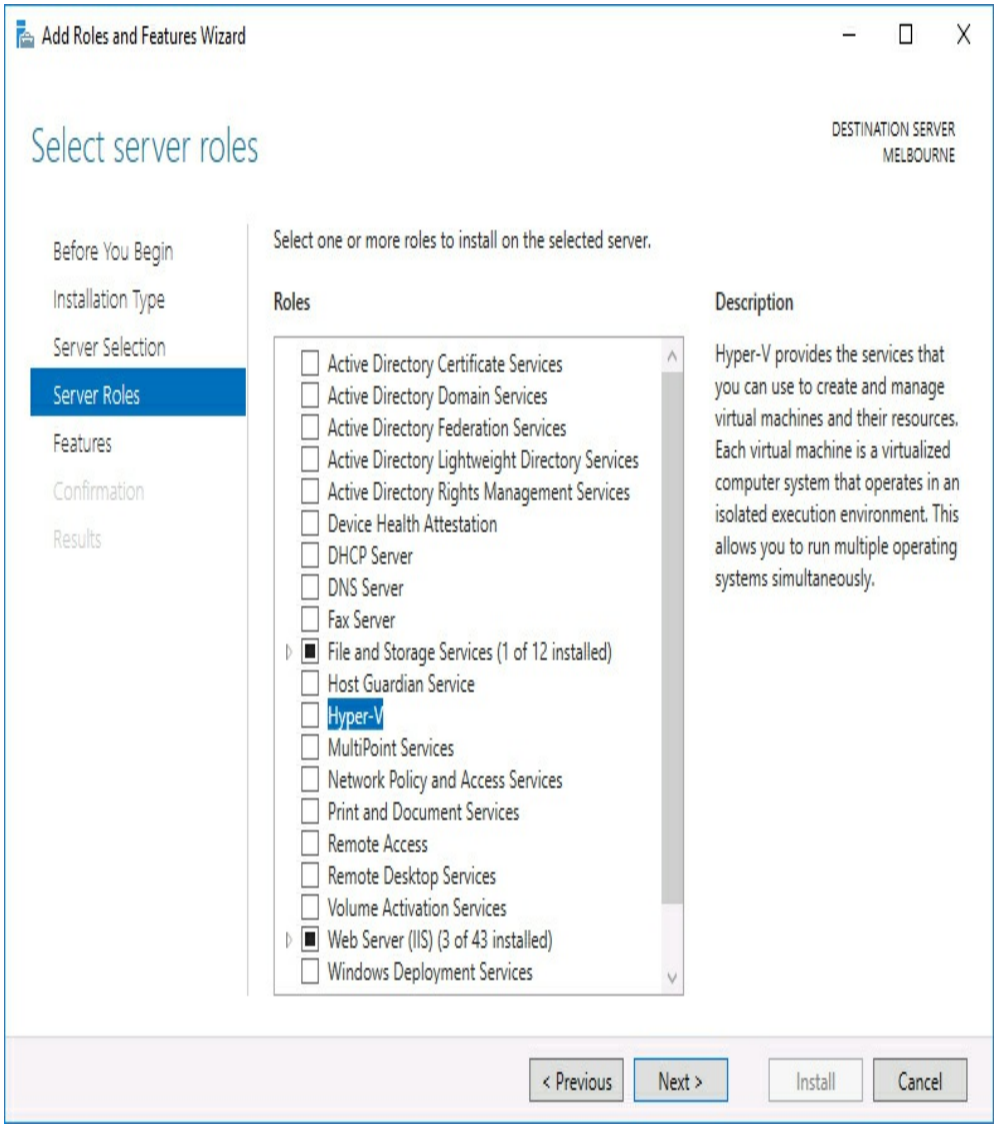


Figure 2-5 Add Roles And Features Wizard

PowerShell provides a better set of tools for adding, removing, and querying the status of features with the *Add-WindowsFeature*, *Get-WindowsFeature*, and *Remove WindowsFeature* cmdlets. For example, Figure 2-6 shows a list of all possible features on a computer, with several features, such as Active Directory

Certificate Services, Active Directory Domain Services, DHCP Server, DNS Server, and File and Storage Services installed. Figure 2-6 shows the output of the *Get-WindowsFeature | Out-Host -paging* command. This provides a list of features and roles that are installed or can be installed on the computer.

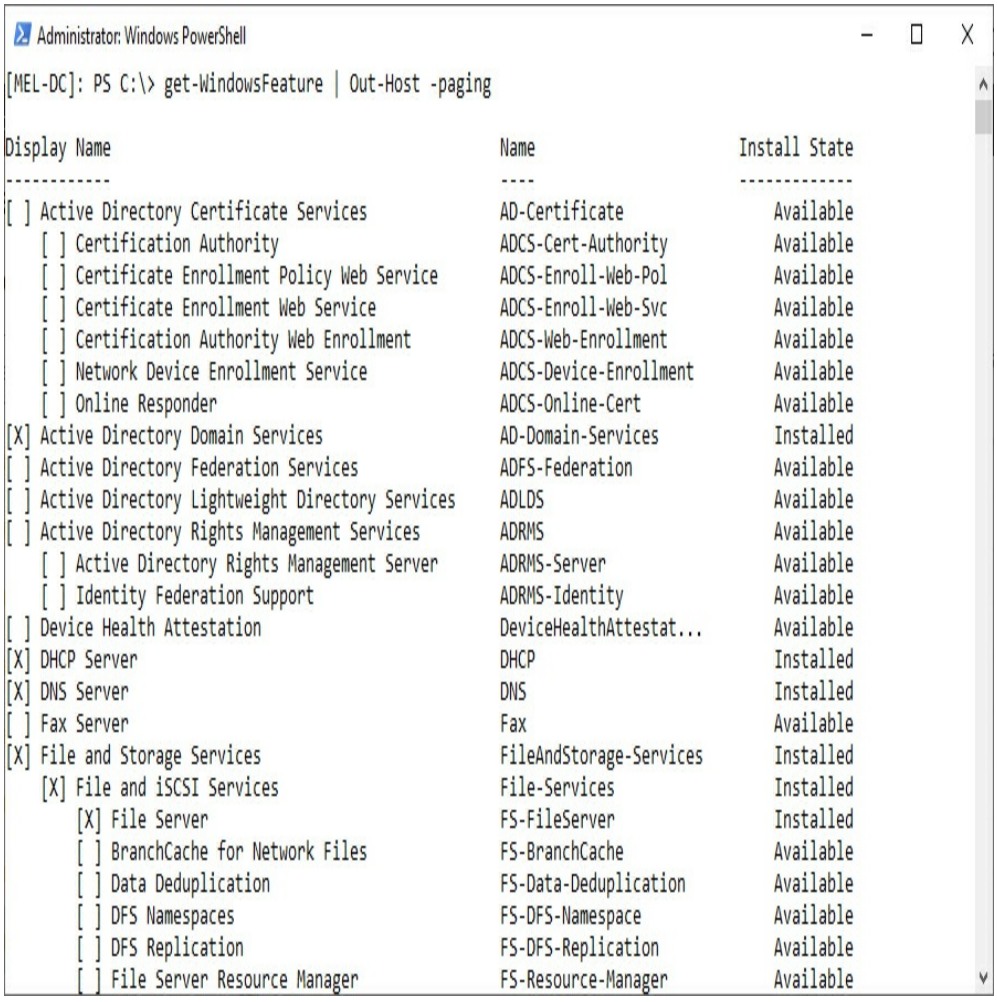


Figure 2-6 View installed features with PowerShell

For example, the following commands install the Hyper-

V role on all computers listed in the text file *virt-servers.txt* and then restarts those servers once the role installation is complete:

Click here to view code image

```
$Computers = Get-Content c:\virt-servers.txt
$Invoke-Command -ScriptBlock { Add-WindowsFeature Hyper-V -
Restart } -Computername $Computers
```

Inside OUT

Desired State

Adding and removing roles and features using Windows Admin Center, Server Manager, or PowerShell is a reasonable way to approach small numbers of server modifications. If you need to deploy many servers that have a consistent configuration of roles and features, you're better off using configuration management tools like PowerShell Desired State Configuration, Puppet, or Chef. With these tools, you deploy the server, and the configuration management system configures it with roles, features, and settings according to your specifications. You'll learn about using tools such as Desired State Configuration and Chef in Chapter 3, "Deployment and configuration."

Chapter 3

Deployment and configuration

Bare metal versus virtualized

Windows images

Answer files

Windows Deployment Services

Virtual Machine Manager

Infrastructure configuration as code

Desired State Configuration

Chef Infra Server

Puppet

Package-management utilities

Depending on the size of your organization, you might deploy a couple of servers and then leave them in

Technet24

production for years, or, as we saw when it came to the (ongoing) retirement of Windows Server 2008 R2, well over a decade. Other organizations deploy servers on a more frequent basis. Some even use a process where rather than applying software updates to a production server and incurring downtime as the update is applied and the server rebooted, they find it faster to deploy a newly patched computer, to migrate the existing workload to the new host, and then to decommission the original server.

This type of iterative deployment is possible because the tools around managing the configuration of servers have evolved. Today, deploying and configuring a new server is no more bothersome than deploying an application would have been a decade ago. In this chapter, you'll learn about how you can use Windows Server images, Windows Deployment Services, Virtual Machine Manager, Desired State Configuration, Puppet, and Chef to deploy and manage the configuration of computers and virtual machines running the Windows Server 2019 operating system.

BARE METAL VERSUS VIRTUALIZED

Today, almost all new workloads are virtualized. For most organizations, virtualization hosts are the primary

remaining physically deployed server, with almost all other workloads running as virtual machines. Unless you have specific reasons not to virtualize a workload, you should run Windows Server 2019 as a virtual machine (VM) rather than as a physically deployed server.

The security available with shielded VMs addresses one of the final objections that many organizations have had around deploying servers virtually rather than physically. With shielded VMs, you can provide the same level of security to a workload that you can to a physically deployed server sitting in a locked cage in a datacenter.

At present, your best bet when deploying virtualization hosts is to choose the Server Core installation option because this has a smaller installation footprint and a reduced attack surface compared to the Server with Desktop Experience option. When you deploy a Server Core virtualization host, you manage Hyper-V remotely from a privileged access workstation or a tool such as Windows Admin Center or Virtual Machine Manager.

WINDOWS IMAGES

With Windows *images*, the entire operating system, as well as associated drivers, updates, and applications, is stored within a single file known as an image file.

During installation, this image is applied to the target volume. Windows images use the Windows Imaging (WIM) file format, which has the following benefits:

- **Multiple deployment methods.** You can use a variety of ways to deploy Windows images. While unusual these days, you can deploy *.wim* files using a traditional DVD-ROM, from a bootable USB drive, from a network share, or through specialized deployment technologies such as Windows Deployment Services (WDS) or System Center Virtual Machine Manager. While it is possible to use System Center Configuration Manager to deploy Windows Server, Configuration Manager is primarily used with client operating systems rather than server operating systems.
- **Editable.** You can mount an image and edit it, enabling, disabling, or removing operating system roles and features as necessary.
- **Updatable.** You can update an image without having to perform an operating system image capture.

The Windows Server 2019 installation media contain two *.wim* files in the Sources folder: *Boot.wim* and *Install.wim*. The installation media uses *Boot.wim* to load the preinstallation environment that you use to deploy Windows Server 2019. *Install.wim* stores one or more operating system images. For example, as Figure 3-1 shows, the *Install.wim* file available with the evaluation version of Windows Server 2019 contains four different editions of Windows Server 2019. Depending on the specifics of the hardware on which you are attempting to install Windows Server, you might need to add extra drivers to the *boot.wim* file. For example, you will need to add extra drivers if the

Windows Server installation routine cannot access the storage device on which you want to install Windows Server because that device's driver is included in the default boot image.

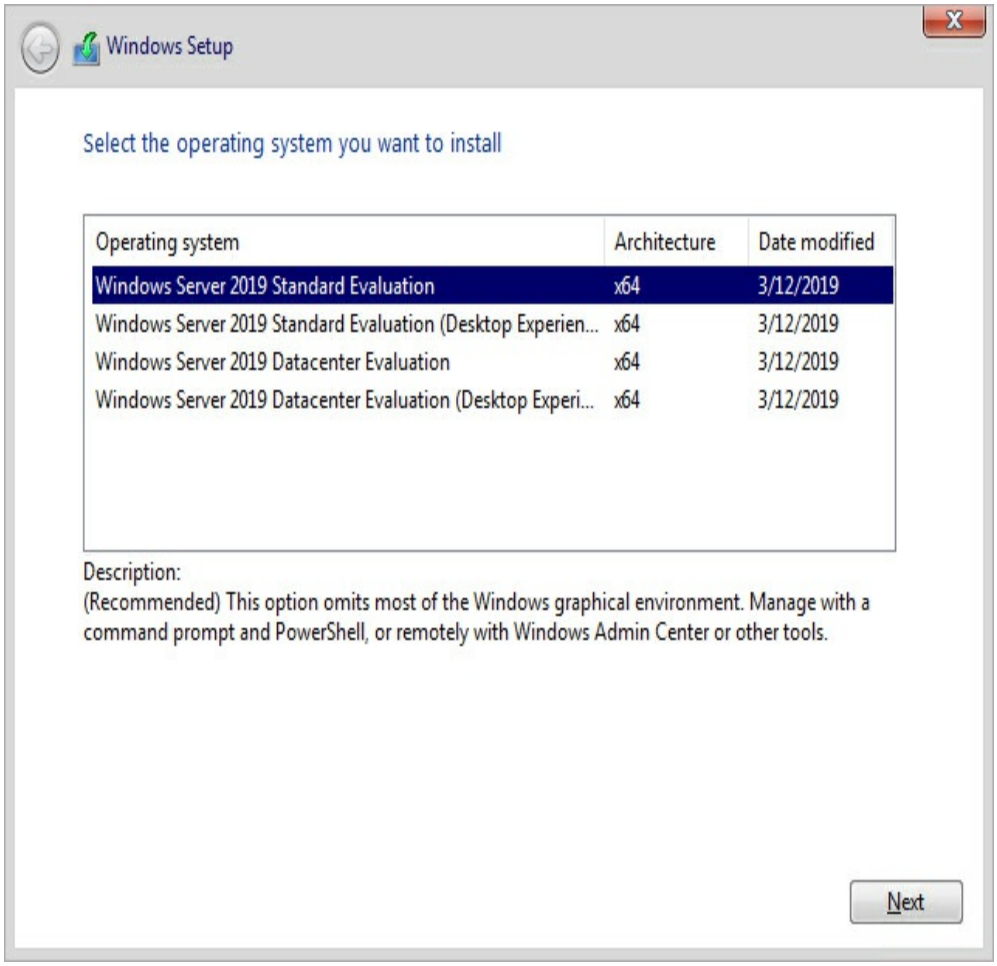


Figure 3-1 Windows Server 2019 editions

Modifying Windows images

The Deployment Image Servicing and Management (DISM) tool is a command-line tool that you can use to manage images in an offline state. You can use *Dism.exe*

to perform the following tasks:

- Enable or disable roles and features
- List roles and features
- Add, remove, and list software updates
- Add, remove, and list software drivers
- Add, remove, and list software packages in .appx format to a Windows image

For example, you can take the *Install.wim* file from the Windows Server installation media and use Dism.exe to mount that image, add new drivers and recent software updates to that image, and save those changes to an image all without having to perform an actual deployment. The advantage is that when you do use this updated image for deployment, the drivers and updates that you added are already applied to the image and you won't have to install them as part of your post-installation configuration routine.

You can use the Microsoft Update Catalog (*https://catalog.update.microsoft.com*) to search for drivers for images that you use with physically deployed servers. This site stores all the certified hardware drivers, software updates, and hotfixes published by Microsoft. Once you download drivers and software updates, you can add them to your existing installation images by using *Dism.exe* or the appropriate PowerShell cmdlets in the *DISM* PowerShell module.

Servicing Windows images

As an IT professional responsible for deploying Windows Server, you need to ensure that your deployment images are kept up to date. The latest software updates should be applied to the image, and any new device drivers for commonly used server hardware should be included.

The main goals of an image servicing strategy are the following:

- Ensure that the latest software updates and hotfixes are applied to the image before the image is deployed to new servers.
- Ensure that the latest drivers are applied to the image before the image is deployed to new servers.

If you don't take these steps, you'll have to wait until after you've deployed the operating system before you can apply updates and drivers. While Windows Server updates are cumulative now, and you won't need to spend hours updating an image from RTM, it's quicker to have the most recent update already applied to the image than it is to wait for the server to deploy, retrieve the latest update, and then wait for it to download and then install. Having updates apply after deployment has occurred consumes a significant amount of time and also substantively increases network traffic. One of your aims when performing a deployment should be to get the server operational and hosting workloads as quickly

as possible.

You can use the DISM (Deployment Image Servicing and Management) utility or the associated PowerShell cmdlets in the *DISM* PowerShell module to service the current operating system in an online state or perform offline servicing of a Windows image.

Servicing images involves performing the following general steps:

1. Mount the image so that it can be modified.
2. Service the image.
3. Commit or discard the changes made to the image.

Mounting images

By mounting an image, you can make changes to that image. When you mount an image, you link it to a folder. You can use File Explorer, Windows PowerShell, or Cmd.exe to navigate the structure of this folder and interact with it as you would any other folder located on the file system. Once the image is mounted, you can also use Dism.exe or PowerShell to perform servicing tasks, such as adding and removing drivers and updates.

A single WIM file can contain multiple operating system images. Each operating system image is assigned an index number, which you need to know before you can mount the image. You can determine the index number using *Dism.exe* with the */Get-wiminfo* switch. For

example, if you have an image named *Install.wim* located in the *C:\Images* folder, you can use the following command to get a list of the operating system images it contains.

click here to view code image

```
Dism.exe /get-wiminfo /wimfile:c:\images\install.wim
```

Figure 3-2 shows the result of this command and lists the images contained in Windows Server 2019. The Standard Evaluation Edition of Windows Server 2019 with the Server Core installation option is assigned index identity *1*, the Standard Evaluation with Desktop Experience is assigned index identity *2*, the Datacenter Evaluation with the Server Core installation option is assigned index identity *3*, and the Datacenter Evaluation with Desktop Experience is assigned index identity *4*.

Technet24

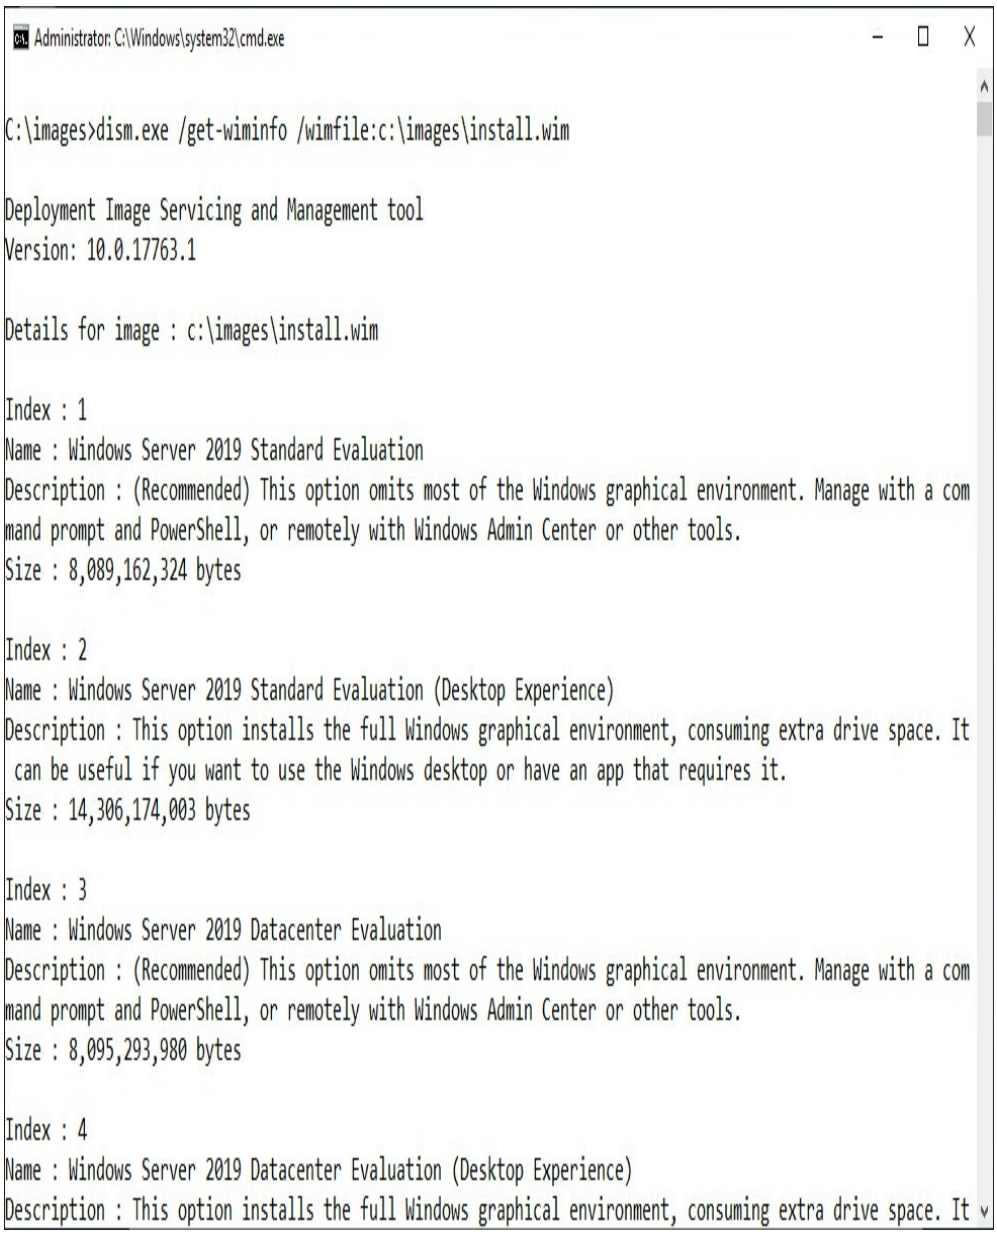


Figure 3-2 Details of an operating system image

You can accomplish the same task using the *Get-WindowsImage* PowerShell cmdlet. For example, to view the contents of the image *Install.wim* in the *C:\Images* folder, run the following command, as

shown in Figure 3-3.

click here to view code image

```
Get-WindowsImage -ImagePath c:\images\install.wim
```

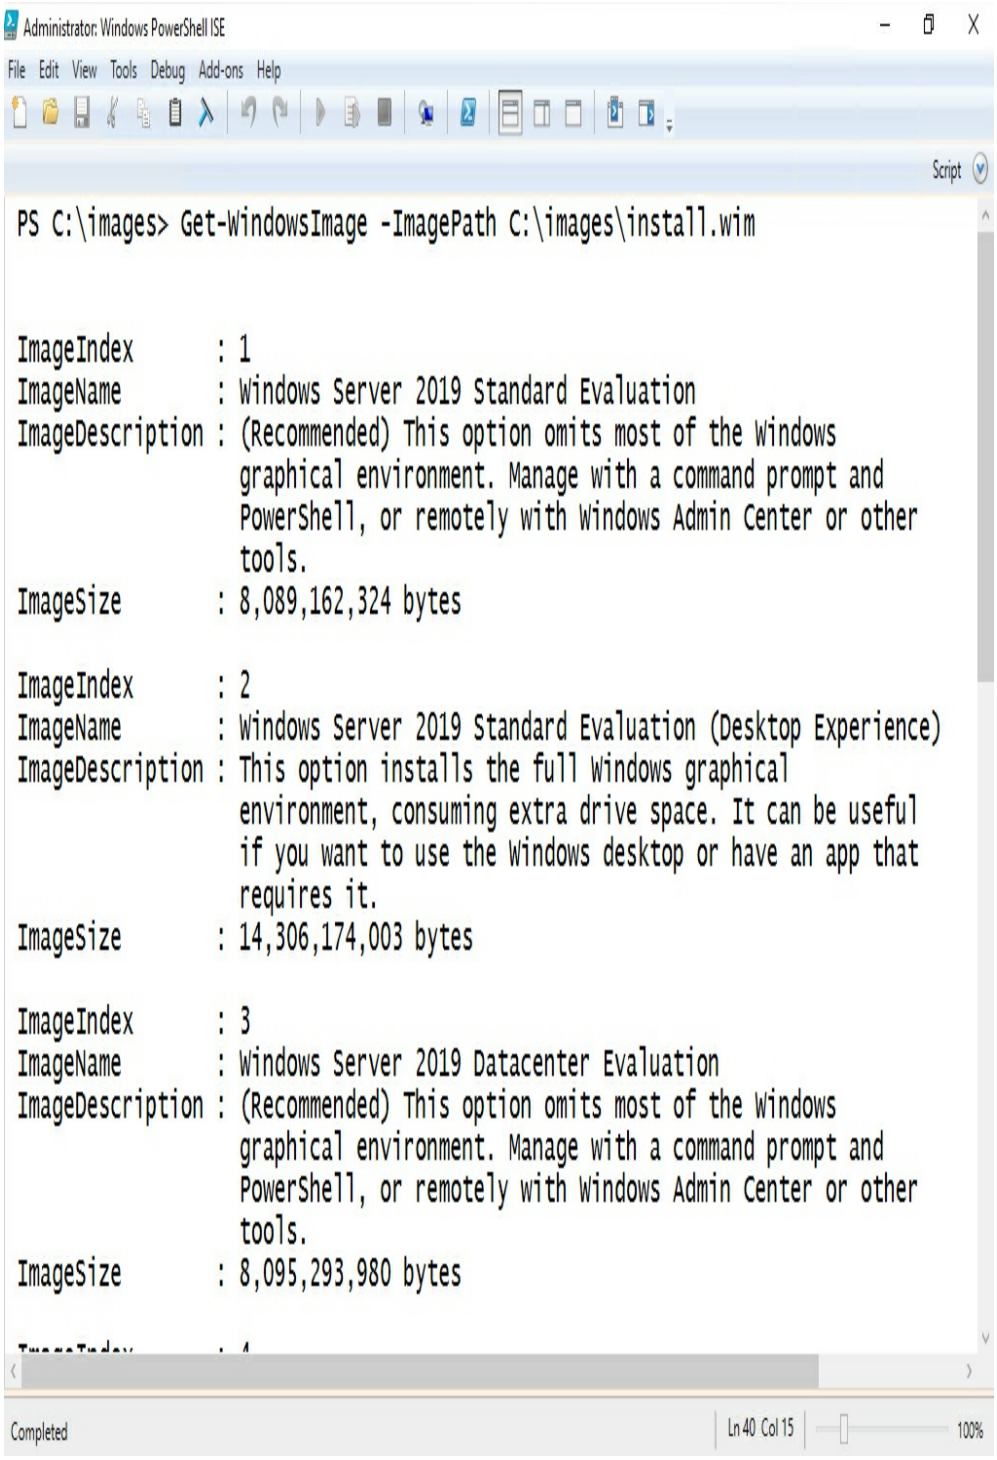


Figure 3-3 Windows image details

Inside OUT

DISM and PowerShell

Image management coverage is provided for both DISM and PowerShell because a large number of IT professionals are accustomed to the former and may still use DISM in scripts. In the long run, it's probably a good idea to use PowerShell, but it's also fair to say that DISM will be around for some time yet.

Once you have determined which operating system image you want to service, you can use the */Mount-image* switch with the *Dism.exe* command to mount that image. For example, to mount the Standard Edition of Windows Server 2019 from the *Install.wim* file that is available with the Evaluation Edition in the *C:\Mount* folder, issue this command:

click here to view code image

```
Dism.exe /mount-image /imagefile:c:\images\install.wim /index:2
/mountdir:c:\mount
```

Alternatively, you can accomplish the same goal using the following *Mount-WindowsImage* command:

click here to view code image

```
Mount-WindowsImage -ImagePath c:\images\install.wim -index 2 -
path c:\mount
```

Adding drivers and updates to images

Once you have mounted an image, you can start to service that image. When servicing images used to deploy Windows Server, the most common tasks are adding device drivers and software updates to the image.

You can use the */Add-Driver* switch with the *Dism.exe* command to add a driver to a mounted image. When using the switch by itself, you need to specify the location the driver's *.inf* file. Rather than adding a driver at a time, you can use the */Recurse* option to have all drivers located in a folder and its subfolders added to an image. For example, to add all of the drivers located in and under the *C: Drivers* folder to the image mounted in the *C:\Mount* folder, use the following command.

click here to view code image

```
Dism.exe /image:c:\mount /Add-Driver /driver:c:\drivers\ /recurse
```

Similarly, you could use the *Add-WindowsDriver* cmdlet to accomplish the same objective by issuing the command:

click here to view code image

```
Add-WindowsDriver -Path c:\mount -Driver c:\drivers -Recurse
```

You can use the */Get-Driver* option to list all drivers that have been added to the image and the */Remove-Driver* option to remove a driver from an image. In PowerShell, you use the *Get-WindowsDriver* cmdlet and the *Remove-WindowsDriver* cmdlets. You can remove only the drivers that you or someone else has added to an image. You can't remove any of the drivers that were present in the image when Microsoft published it. You might choose to remove an existing

driver if the driver you added in the past has since been updated.

You can use Dism.exe with the */Add-Package* switch to add packages that contain updates or packages in *.cab* or *.msu* format. Software updates are available from the Microsoft Update Catalog website in *.msu* format. For example, you can download an update from the Microsoft Update Catalog website named *Cumulative Update For Windows Server 2019 for x64-based Systems (KB4505056)* to the *C:\Updates* folder on a computer, as shown in Figure 3-4.

Technet24

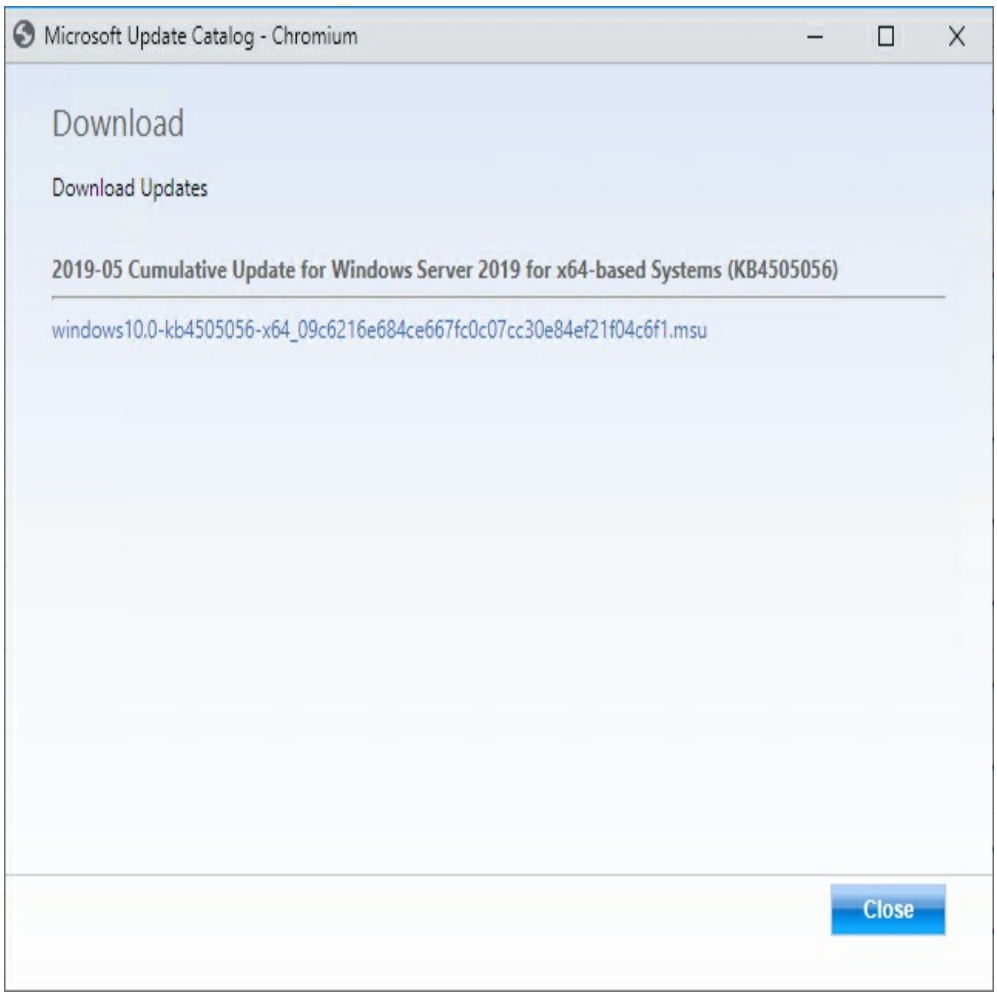


Figure 3-4 Downloading a cumulative update

If you mounted a WIM image of the Windows Server 2019 operating system in the *C:\Mount* folder, you could apply the update to the image by using the command:

click here to view code image

```
Dism.exe /image:c:\mount /Add-Package
/PackagePath:"c:\updates\windows10.0-kb4505056-
x64_09c6216e684ce667fc0c07cc30e84ef21f04c6f1.msu"
```

You can accomplish the same thing with the following *Add-WindowsPackage* command:

click here to view code image

```
Add-WindowsPackage -path c:\mount -packagepath "c:\updates\AMD64-
all-windows10.0-kb4505056-
x64_09c6216e684ce667fc0c07cc30e84ef21f04c6f1.msu"
```

Adding roles and features

You can determine which features are available in a mounted operating system image by using the */Get-Features* switch. Features include both Windows Server 2019 roles and features. For example, to learn which features are available in the image mounted in the *C:\Mount* folder, use this command, as shown in Figure 3-5:

click here to view code image

```
Dism.exe /image:c:\mount /Get-Features
```

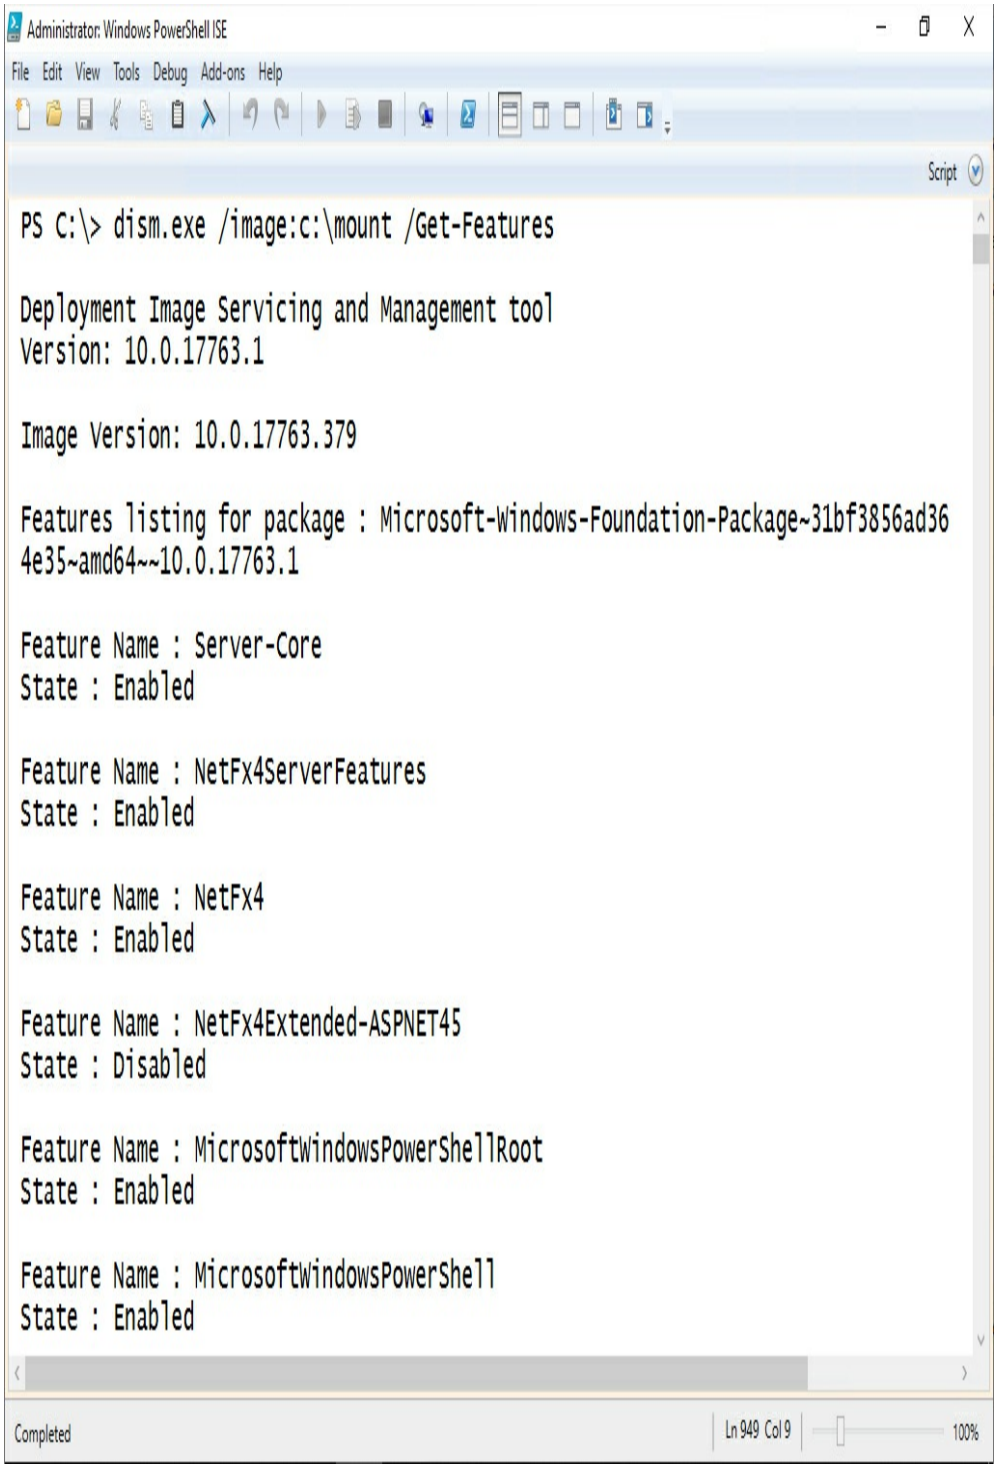


Figure 3-5 View features

You can enable or disable a specific feature using the */Enable-Feature* switch. For example, to enable the *NetFx3ServerFeatures* feature, which enables the .NET Framework 3.5 server features in an image, use this command.

click here to view code image

```
Dism.exe /image:c:\mount /Enable-Feature /all
/FeatureName:NetFx3ServerFeatures
```

Some features in the Windows Server image are in a state in which they are listed as having their payload removed, which means that the installation files for that feature are not included in the image. If you install a feature that had its payload removed when the operating system was deployed, the operating system can download the files from the Microsoft update servers on the Internet. You can also specify the location of the installation files. The installation files for the features that have had their payload removed in Windows Server are located in the *\Sources\sxs* folder of the volume in which the installation media is located.

You can add these payload-removed features to an image by using *Dism.exe* and specifying the source directory. For example, to modify an image mounted in the *C:\Mount* folder so that the Microsoft .NET Framework 3.5 features are installed and available, issue this command when the installation media is

located on volume D, as shown in Figure 3-6:

click here to view code image

```
Dism.exe /image:c:\mount /Enable-Feature /all /FeatureName:NetFx3
/Source:d:\sources\sxs
```

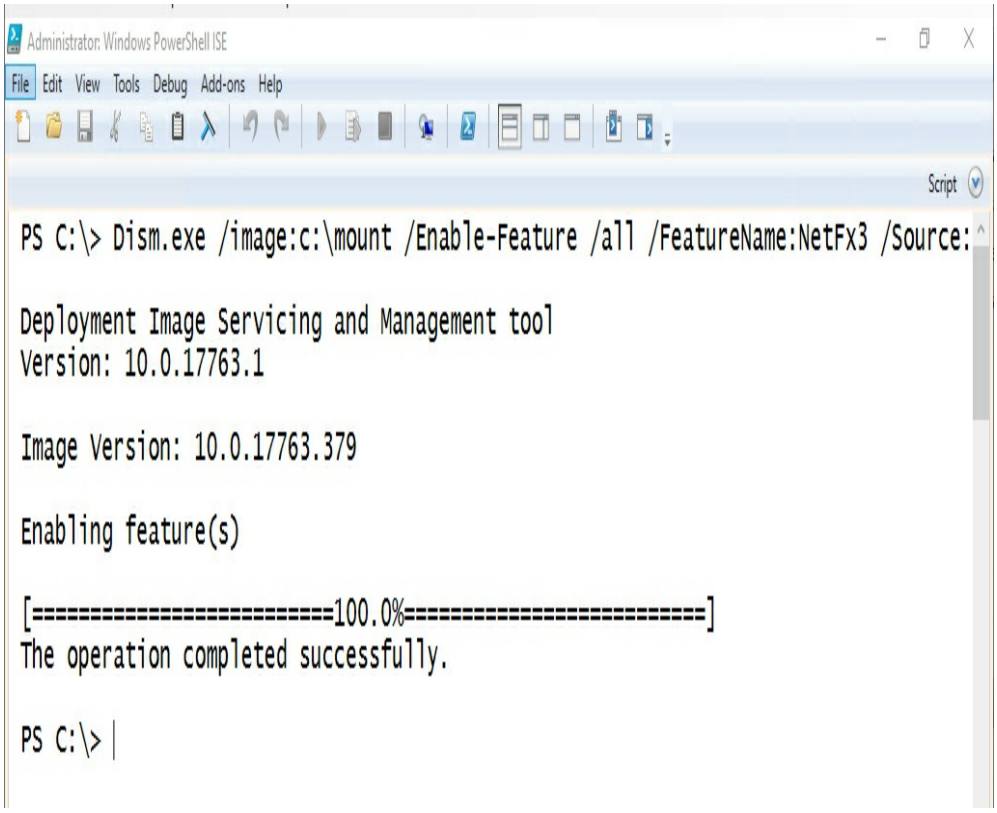


Figure 3-6 Enabling a feature

Inside OUT

.NET Framework 3.5

While features like WINS remain available within Windows Server 2019, the .NET Framework 3.5 is not only not installed by default, but it has been quarantined into its own separate installation folder. One might suspect that Microsoft wants it to be difficult to install. The problem is that such a huge variety of software depends on .NET Framework 3.5, which makes installing it more of a hassle than it needs to be.

Committing an image

When you finish servicing an image, you can save your changes using the */Unmount-Wim* switch with the */Commit* option. You can discard changes using the */Discard* option. For example, to make changes and then commit the image mounted in the *C:\Mount* folder, use this command.

click here to view code image

```
Dism.exe /Unmount-Wim /MountDir:c:\mount /commit
```

You can use the *Save-WindowsImage* PowerShell cmdlet to save changes to an image without dismounting the image. You use the *Dismount-WimImage* cmdlet with the *Save* parameter to save the modifications that you've made to an image and then dismount it. For example, to dismount and save the image mounted in the *C:\Mount* folder, run the command:

click here to view code image

```
Dismount-WimImage -Path c:\mount -Save
```

Once you have committed the changes, the *.wim* file that you originally mounted is updated with these modifications. You can then import this *.wim* file into WDS, use it to build a virtual hard disk, or use it with bootable USB installation media to deploy Windows Server 2019 with these updates and modifications

already applied.

Build and capture

When you perform a build and capture, you deploy an operating system; provision that operating system with updates, applications, and drivers; and then capture that operating system for deployment. Build and capture is used less often with server operating systems because they rarely require the same sort of application deployment that is required for client operating systems. If you can just pull a container with an updated application down onto an operating system after deployment, there is little reason to include it in the image. Additionally, with tools such as Desired State Configuration, Chef, and Puppet, many post-installation and configuration tasks can be completely automated, reducing the hassle of post-installation configuration.

If your deployment strategy does involve the deployment and capture of Windows Server 2019, you need to remember that you'll need to generalize the image prior to capture, removing any configuration information that is specific to the installation. You can perform this task using the Sysprep.exe utility. Sysprep.exe is included with Windows Server 2019 and is located in the *C:\Windows\System32\Sysprep* folder. When you use Sysprep.exe to prepare the image, you can configure the image to return to the system Out-of-

Box Experience (OOBE), as shown in Figure 3-7. This is the same experience you get when Windows Server boots for the first time, though in this case all the updates, applications, and drivers included in the captured image are included in the newly deployed image.

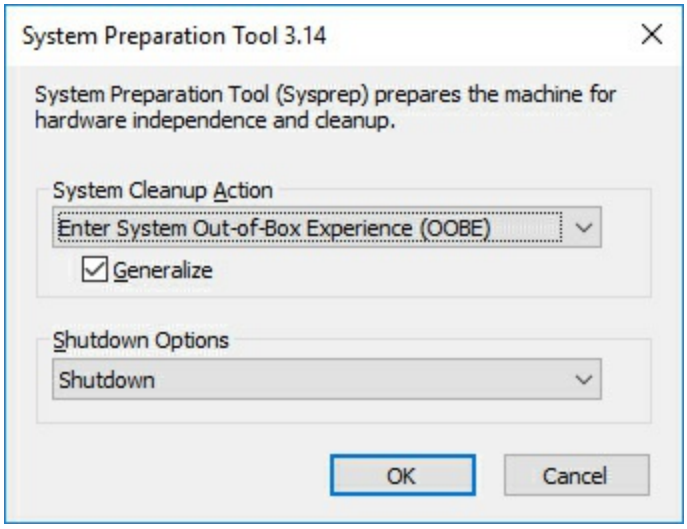


Figure 3-7 Sysprep

You can use Dism.exe with the */Capture-Image* switch or the *New-WindowsImage* PowerShell cmdlet to capture an image.

ANSWER FILES

Answer files present a traditional method of performing operating system deployment and configuration. While not as comprehensive as newer technologies such as Desired State Configuration (DSC) that assist not only

Technet24

in deployment but in ongoing configuration, answer files allow you to automate the process of deploying Windows Server.

Instead of having to manually select specific installation options and perform post–installation configuration actions such as joining a newly deployed server to an AD DS domain, you can automate the process with answer files. During setup, the Windows Server looks for a file on local and attached media named *Autounattend.xml*. If this file is present, Windows Server automatically uses the settings contained in the file to configure the new server deployment.

As its name suggests, *Autounattend.xml* uses the XML file format. Although it is certainly possible for you to manually edit this XML file using a text editor such as Notepad, this process is complicated, and you are likely to make errors that cause the file to not work. The Windows System Image Manager (known as Windows SIM) is a GUI-based tool that you can use to create an answer file. When using the tool, you must specify the image for which you want to create an answer file. Windows SIM then creates a catalog file for all the options that you can configure. After you configure all the settings that you want automated during installation and post-installation configuration, you can have the tool output an answer file using correct XML syntax. Windows SIM is included with the Windows

Assessment and Deployment Kit (Windows ADK), which you can download from the Microsoft website.

To create an answer file using Windows SIM, perform the following steps:

1. Download and install Windows ADK from the Microsoft website using the installation defaults. You can do this with Chocolatey, covered later in this chapter, by running the following command:

   click here to view code image

   ```
   Choco install -y windows-adk-all
   ```

2. Copy the *\Sources\install.wim* file from the Windows Server installation media to a temporary directory on the computer on which you have installed Windows ADK.
3. Open Windows SIM from the Start screen.
4. In the Windows SIM interface, click File, and then choose Select Windows Image. Open the *Install.wim* file.
5. Select an operating system image in the install image for which you wish to create an answer file. For example, Figure 3-8 shows the selection of the Standard edition with Desktop Experience operating system.

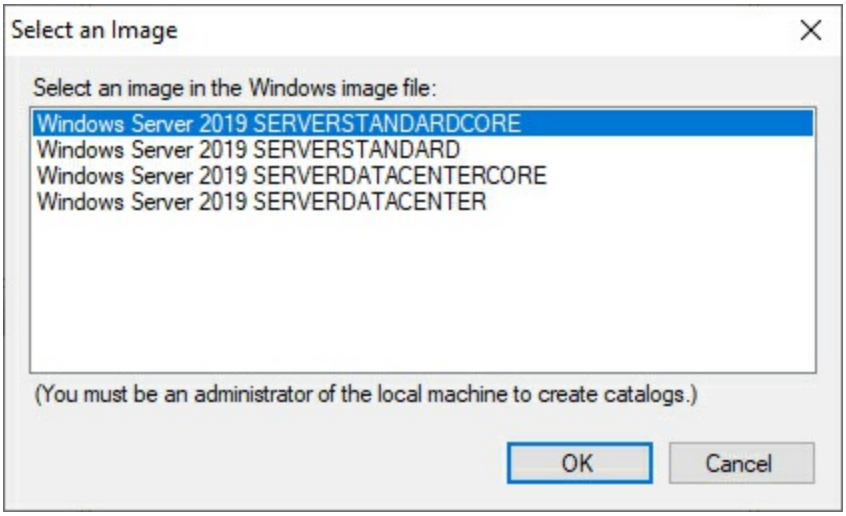


Figure 3-8 Select an image

6. When prompted to create a catalog file, click Yes.
7. Click File, New Answer File.
8. Use Windows SIM to select each component that you want to configure. Figure 3-9 shows how you can configure installation to join the *adatum.com* domain.

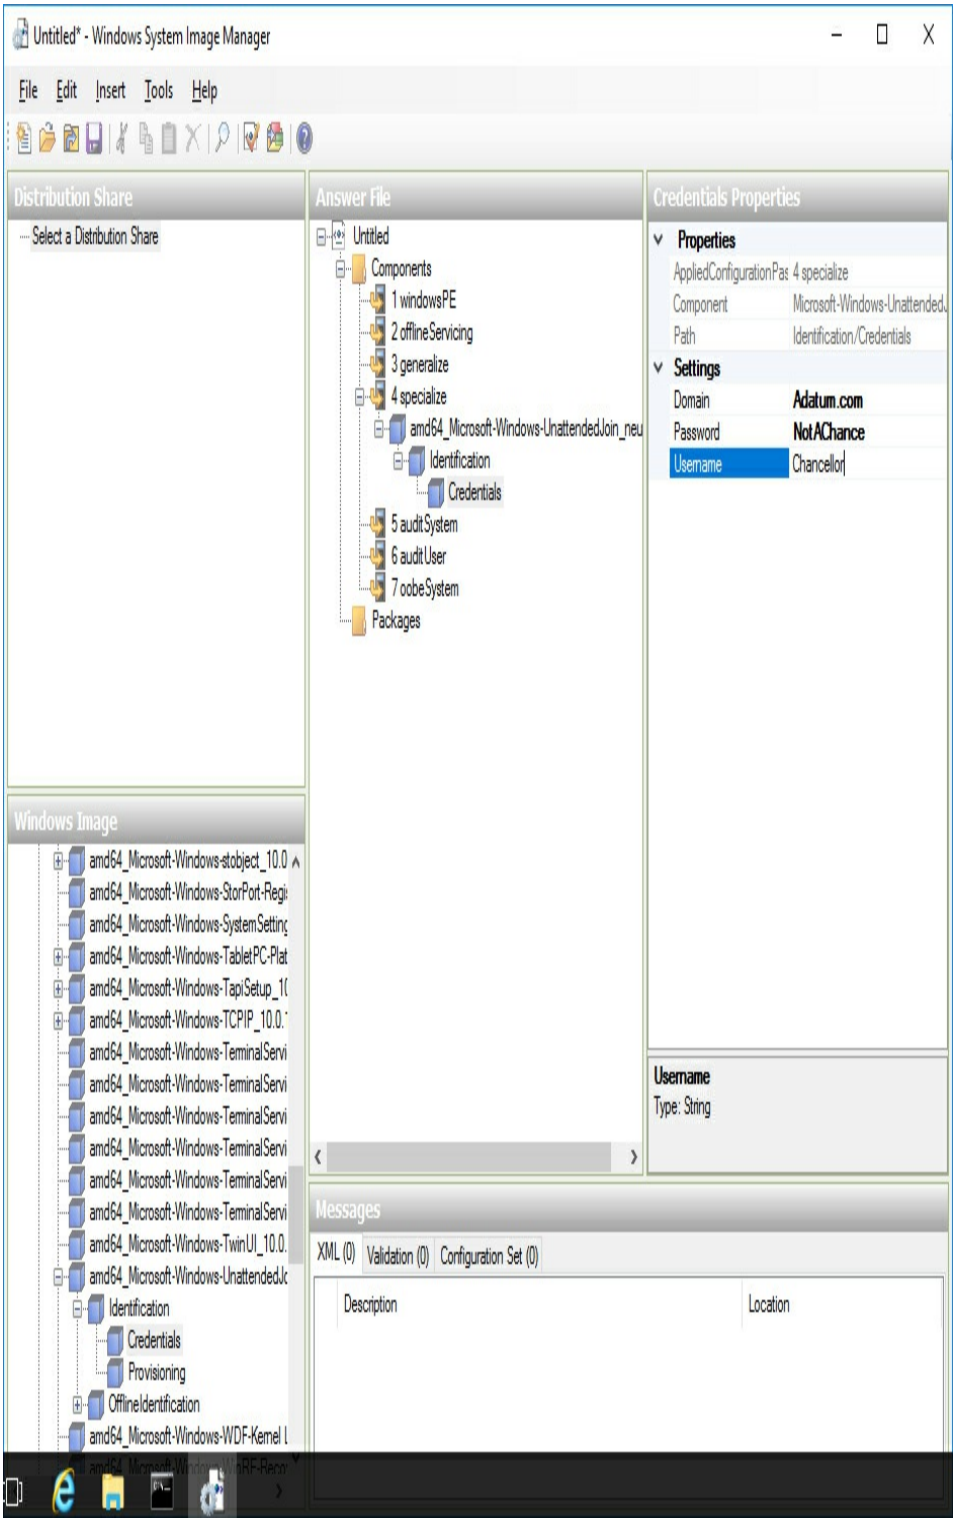


Figure 3-9 Windows System Image Manager

WINDOWS DEPLOYMENT SERVICES

Windows Deployment Services (WDS) is a server role that you can deploy on computers running Windows Server. WDS enables you to deploy operating systems, including Windows 8.1, Windows 10, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019, to computers over the network. WDS can send these operating systems across the network using multicast transmissions, which means that multiple computers receive the same operating system image while minimizing the use of network bandwidth. When you use multicast transmissions, the same amount of traffic crosses the network independently of whether you are deploying an operating system to 1 computer or to 50. WDS also can use unicast transmissions.

Deploying Windows Server through WDS involves performing the following steps:

1. An operating system deployment transmission is prepared on the WDS server.
2. The media access control (MAC) addresses of Pre-boot Execution Environment (PXE)–compliant network adapters are made available to the WDS server.
3. The computers that are targets of the transmission boot using their PXE-compliant network adapters.
4. These computers locate the WDS server and begin the operating system setup process. If the WDS server has been provisioned with an answer

file, the setup completes automatically. If the WDS server has not been provisioned with an answer file, an administrator must enter setup configuration information.

Each WDS server can have only one unattended installation file for each processor architecture. Because unattended installation files differ between server and client, you either need to swap unattended files when you are switching between client and server or have multiple WDS servers. WDS can be used in conjunction with other technologies such as Desired State Configuration where an answer file only performs basic configuration tasks, with the substantial tasks completed by an advanced configuration technology.

WDS requirements

WDS clients need a PXE-compliant network adapter, which is rarely a problem because almost all modern network adapters are PXE-compliant. You can also use WDS to deploy Windows Server 2012 and later to virtual machines running under Hyper-V. The trick to doing this is to use a legacy rather than a synthetic network adapter when creating the virtual machine as a Generation 1 virtual machine. This isn't necessary when using Generation 2 virtual machines because the Generation 2 virtual machine network adapters support PXE booting.

If you have a computer that does not have a PXE-

compliant network adapter, you can configure a special type of boot image known as a discover image. A discover image boots an environment, loading special drivers to enable the network adapter to interact with the WDS server. You create the boot image by adding the appropriate network adapter drivers associated with the computer that can't PXE boot to the *Boot.wim* file from the Windows Server installation media.

WDS has the following requirements:

- A Windows Server DNS server must be present on the local area network (LAN).
- Prior to Windows Server 1810, an authorized Dynamic Host Configuration Protocol (DHCP) server must be present on the network. You can host WDS and DHCP on the same computer as long as you configure the options shown in Figure 3-10. Versions of Windows Server after 1810—including Windows Server 2019—can be used with third-party DHCP servers. Because this limits your ability to use IP address tracking through IPAM, using a third-party DHCP server if you are using WDS is not a recommended strategy.

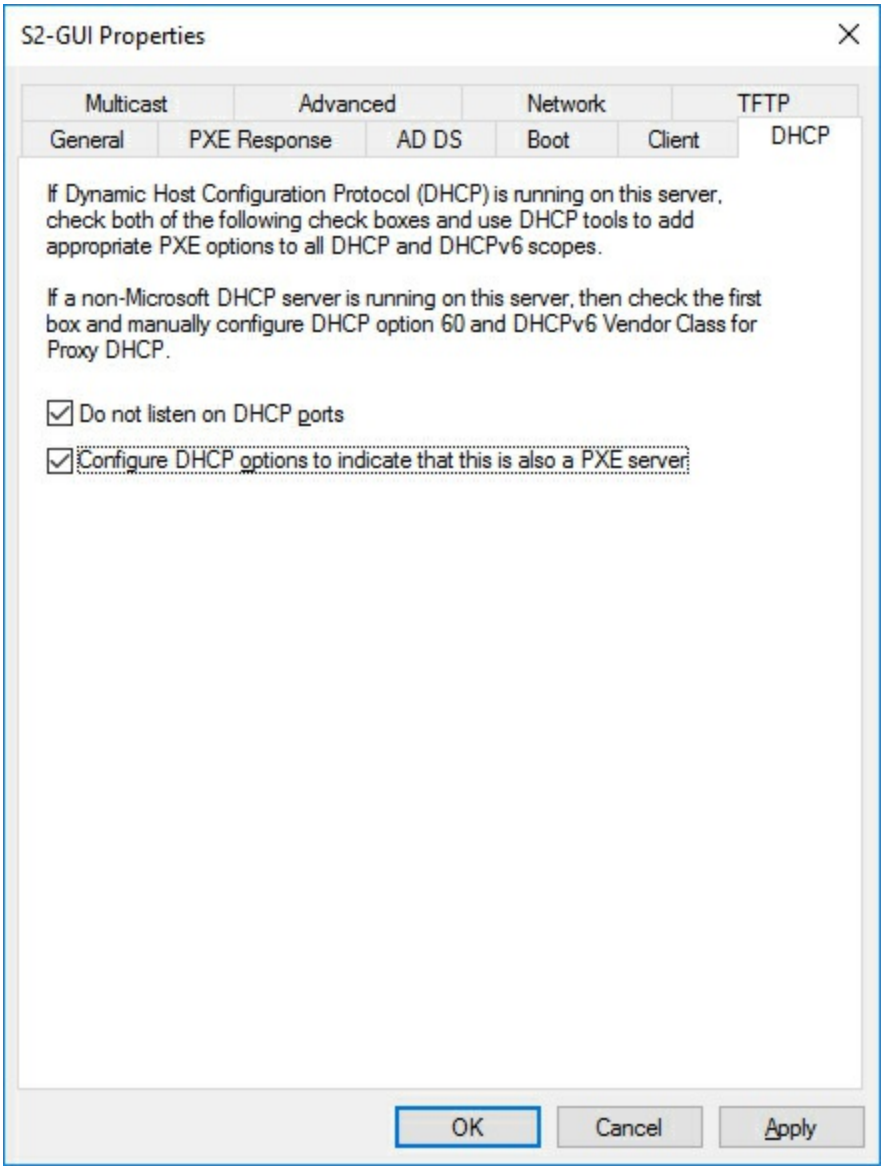


Figure 3-10 WDS DHCP settings

If you install WDS from the Add Roles And Features Wizard, you can configure these settings automatically. Although the WDS server does not require a static IP address, it is good practice to ensure that infrastructure roles such as WDS always use a consistent network address. You can install WDS on computers running the

Technet24

Server Core version of Windows Server.

When installing WDS on Server Core, you have to specify the location of the source files or ensure that the server has a connection to the Internet, which enables them to be downloaded automatically. Although it is possible to manage WDS from Windows PowerShell, most administrators use the graphicsal WDS Remote Server Administration Tools (RSAT) from a computer running Windows 10, Windows Server 2016, or Windows Server 2019 with Desktop Experience.

Managing images

Images contain either entire operating systems or a version of a special stripped-down operating system known as Windows PE. Windows PE functions as a type of boot disk, enabling a basic environment to be loaded from which more complex maintenance and installation tasks can be performed. WDS uses four image types: boot image, install image, discover image, and capture image.

- **Boot Image.** A special image that enables the computer to boot and begin installing the operating system using the install image. A default boot image, named *Boot.wim*, is located in the *sources* folder of the Windows Server installation media.
- **Install Image.** The main type of image discussed in this chapter. Contains the operating system as well as any other included components, such as software updates and additional applications. A default install image, named *Install.wim*, is present in the sources

folder of the Windows Server installation media. Install images can be in *.vhd* or *.vhdx* format, though you can only manage install images using the WDS console in Windows Server 2012 R2, Windows Server 2016, or Windows Server 2019.

- **Discover Image.** This special image is for computers that cannot PXE boot to load appropriate network drivers to begin a session with a WDS server.
- **Capture Image.** A special image type that enables a prepared computer to be booted so that its operating system state can be captured as an install image. You add capture images as boot images in WDS.

To import an image into WDS, perform the following steps:

1. Open the Windows Deployment Services console.
2. Click Install Images. From the Action menu, click Add Install Image.
3. Choose whether to create a new image group or to use an existing image group.
4. Specify the location of the image file.
5. In the Available Images page of the Add Image Wizard, select the operating system images that you want to add. When the image or images are added, click Next, Finish.

Configuring WDS

The installation defaults for WDS are suitable when you deploy the role in small environments. If you are deploying WDS in larger environments and do not choose to implement System Center Virtual Machine Manager for server operating system deployments, you might want to configure the options discussed in the following sections, which are available by editing the

properties of the WDS server in the Windows Deployment Services console.

PXE response settings

With PXE response settings, you can configure how the WDS server responds to computers. As Figure 3-11 shows, you can configure WDS not to respond to any client computers (this effectively disables WDS), to respond to known client computers, or to respond to all computers but require an administrator to manually approve an unknown computer. Known computers are those that have prestaged accounts in Active Directory. You can prestage computers if you know the MAC address of the network interface card (NIC) that the computer uses. Vendors often supply a list of MAC addresses associated with computers when you purchase those computers, and you can use this list to prestage computer accounts.

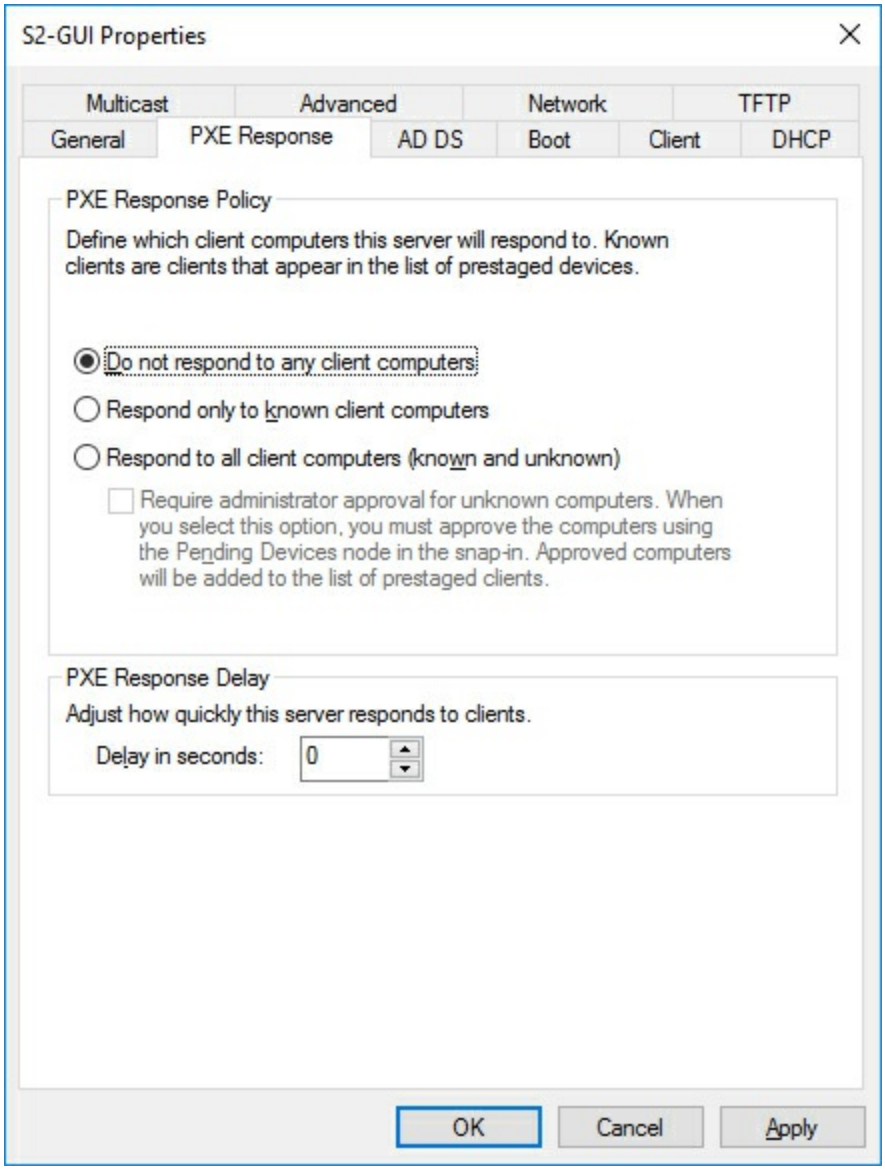


Figure 3-11 PXE Response settings

You use the PXE Response Delay setting when you have more than one WDS server in an environment. You can use this setting to ensure that clients receive transmissions from one WDS server over another, with the server configured with the lowest PXE response

delay having priority over other WDS servers with higher delay settings.

Client Naming Policy

The Client Naming Policy enables you to configure how computers installed from WDS are named if you aren't using deployment options that perform the action. You can also use the settings on the AD DS tab, shown in Figure 3-12, to configure domain membership and organizational unit (OU) options for the computer account.

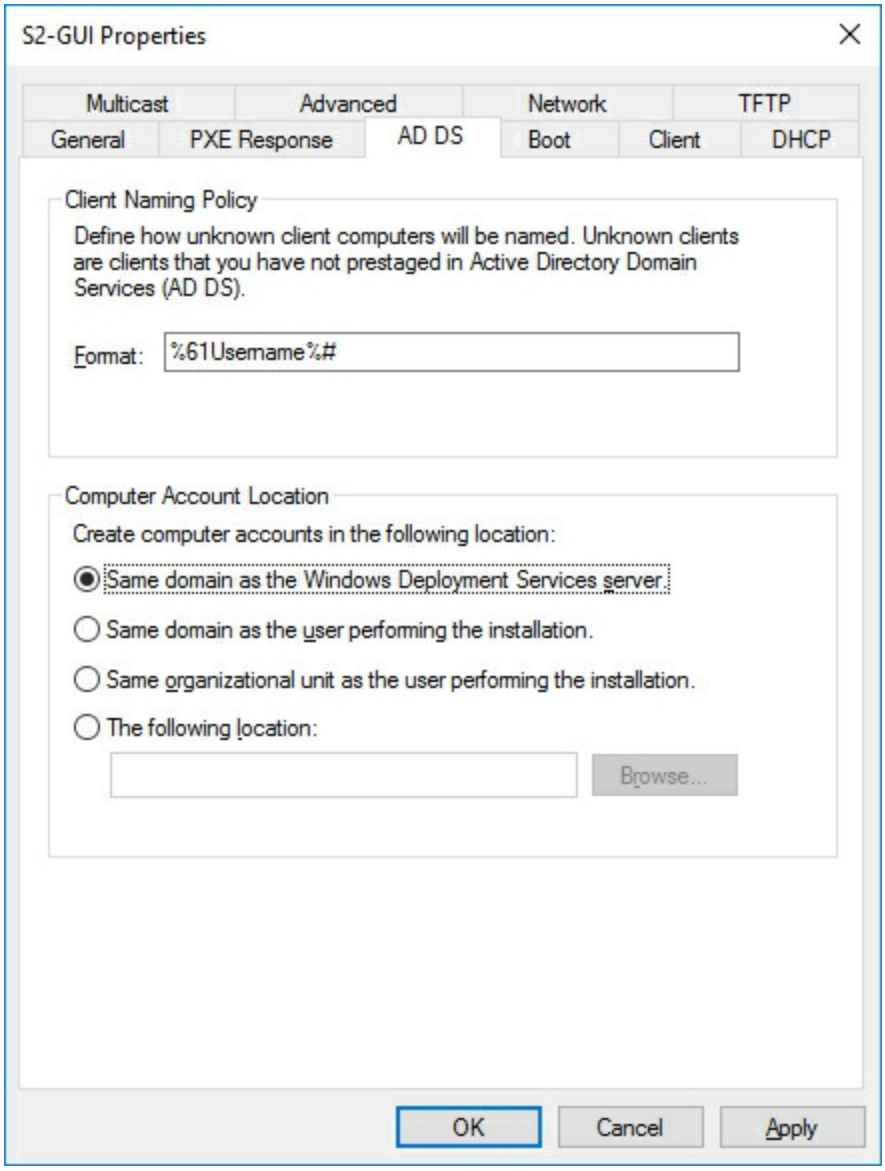


Figure 3-12 Client Naming Policy

WDS boot options

In the Boot options tab of the WDS server's Properties dialog box, shown in Figure 3-13, you can configure how clients that PXE boot interact with the WDS server. You can also configure a default boot image for each

Technet24

architecture supported by WDS. By default, once a client has connected to a WDS server, someone must press the F12 key to continue deploying the operating system. In environments in which you are performing a large number of simultaneous deployments, requiring this level of manual intervention might substantially delay the deployment.

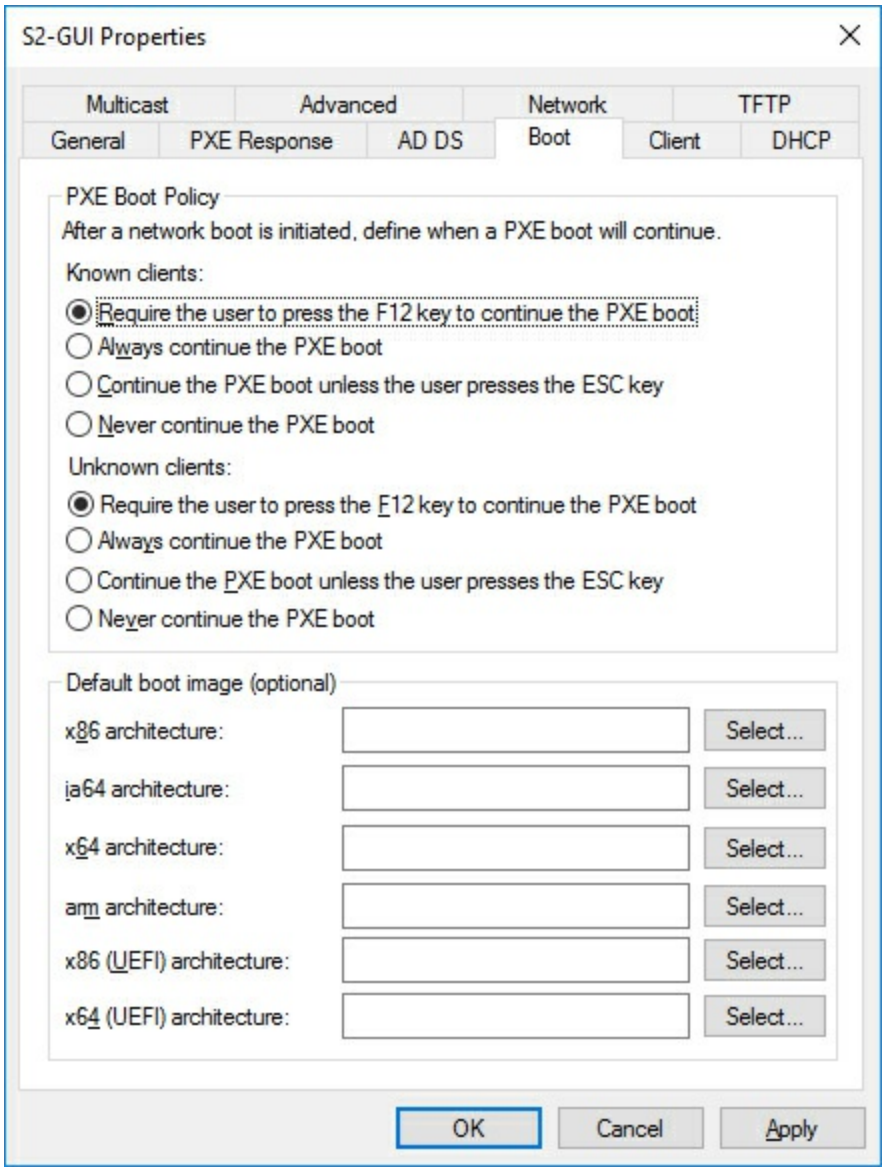


Figure 3-13 Boot options

Multicast options

The default settings of WDS have all computers that join the multicast transmission receiving the installation image at the same speed. If you frequently deploy operating systems, you are aware that sometimes there are one or two computers that have network adapters that slow transmission; transmissions that should take only 15 minutes now take half a day. You can configure the transfer settings on the Multicast tab, shown in Figure 3-14, so that clients are partitioned into separate sessions depending on how fast they can consume the multicast transmission. You still have those slow computers taking a long time to receive the image, but the other computers connected to the transmission can complete the deployment quicker.

Technet24

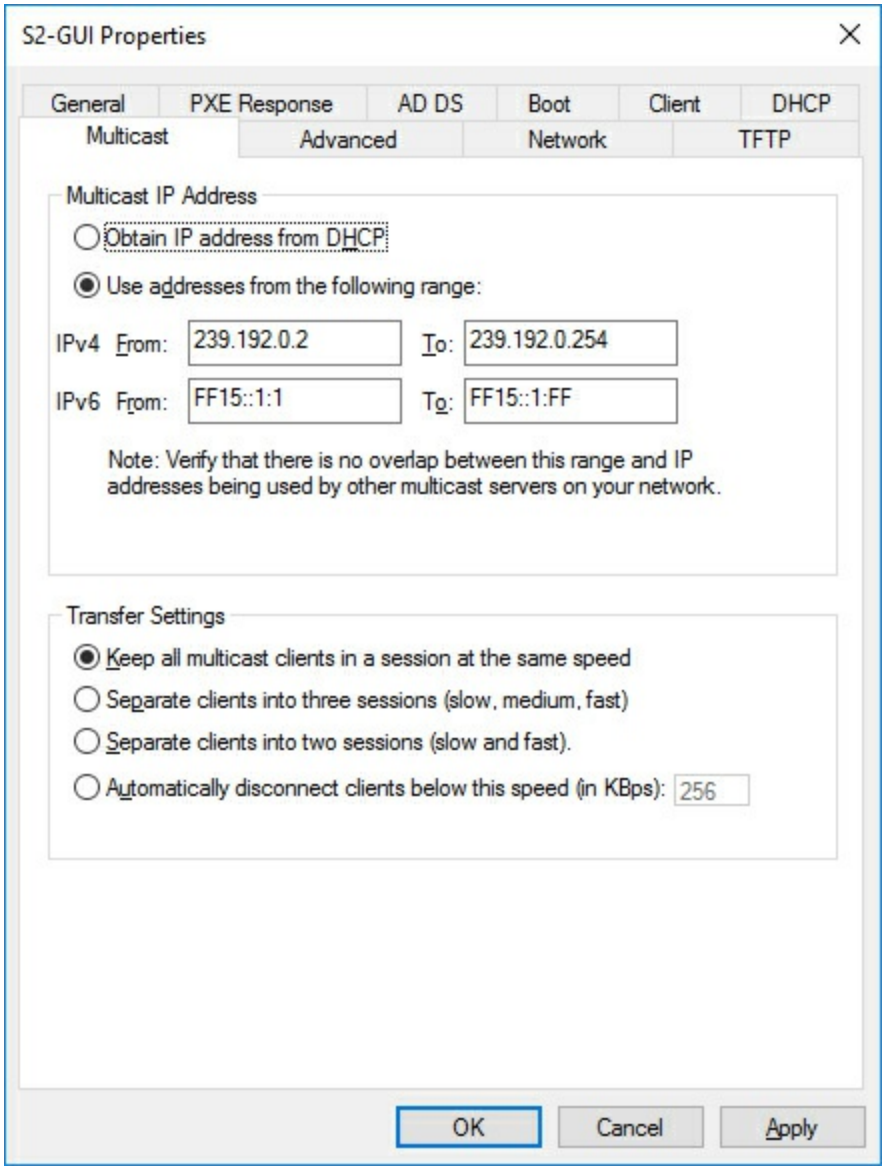


Figure 3-14 WDS Multicast tab

Other options

Although you are less likely to need them, you can configure other options on the following tabs:

- **Advanced tab.** You can configure WDS to use a specific domain controller and global catalog (GC) server. You can also configure

whether WDS is authorized in DHCP. DHCP authorization occurs automatically when you install the WDS role.

- **Network tab.** You can specify a User Datagram Protocol (UDP) port policy to limit when UDP ports are used with transmissions. You can also configure a network profile to specify the speed of the network, minimizing the chance that WDS transmissions slow the network down.
- **TFTP tab.** You can specify maximum block size and Trivial File Transfer Protocol (TFTP) window size.

Configuring transmissions

You use WDS transmissions to set WDS to transfer the operating system image to PXE clients. When configuring a WDS transmission, you need to decide what type of multicast transmission you are going to perform in the Multicast Type page of the Create Multicast Transmission Wizard, as shown in Figure 3-15.

Technet24

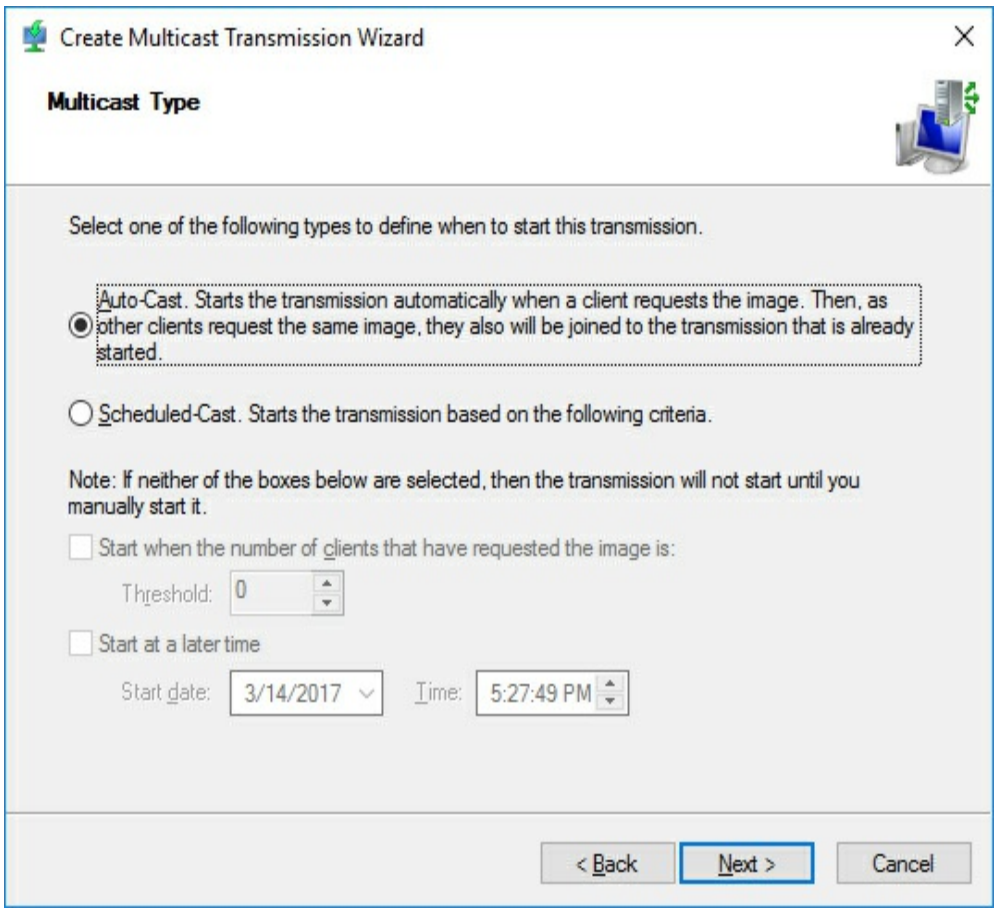


Figure 3-15 Multicast type

The difference between these options is as follows:

- **Auto-Cast.** A transmission starts when a client requests the image. If another client requests the same image, the client joins the existing transmission, caching data from the current transfer and then retrieving data that was transmitted before the client joined the transmission. This is the best option to use when you are performing one-off deployments.
- **Scheduled-Cast.** You choose either to start the transmission when a specified number of clients have joined, or you start the transmission at a particular date and time. Scheduled-Cast is the best option to use when you are deploying the same operating

system image to a large number of computers.

To configure a WDS transmission, perform the following steps:

1. Open the Windows Deployment Services console, expand the WDS server from which you want to perform the deployment, and click Multicast Transmissions. In the Action menu, click Create Multicast Transmission.
2. Provide a name for the multicast transmission.
3. On the Image Selection page, specify which operating system image you want to deploy using the transmission.
4. On the Multicast Type page, specify whether you use Auto-Cast or Scheduled-Cast. If you choose Scheduled-Cast, select the number of clients or the transmission start time.

Driver groups and packages

You can stage device drivers on a WDS server by importing the device driver as a package. A driver package contains the extracted driver files. You can import the driver package into WDS by locating the driver's *.inf* file. When using the WDS console, you can either import individual driver packages or all the drivers in a set of folders.

In the WDS console, you can organize drivers into driver groups. A driver package can be a member of more than one group, and deleting a driver group does not delete the associated driver packages. You can use driver groups with filters to limit which driver packages are available to WDS clients.

VIRTUAL MACHINE MANAGER

For most organizations, the majority of their server deployments involve virtual machines rather than deploying to bare metal physical hardware. While tools such as the Hyper-V console and WDS are adequate for smaller environments, if you need to deploy hundreds or thousands of virtual machines each year, you need a more comprehensive set of tools than those that are included with the Windows Server operating system.

System Center Virtual Machine Manager is one tool that you can use to manage your organization's entire virtualization infrastructure, from virtualization hosts, clusters, and VMs, to managing the entire networking and storage stack. In this section, you'll learn about VMM templates, storage, networking, and host groups.

Virtual machine templates

A Virtual Machine Manager VM template allows you to rapidly deploy virtual machines with a consistent set of settings. A VMM VM template is an XML object that is stored with a VMM library, and it includes one or more of the following components:

- **Guest Operating System Profile.** A guest operating system profile that includes operating system settings.
- **Hardware Profile.** A hardware profile that includes VM hardware settings.
- **Virtual Hard Disk.** This can be a blank hard disk or a virtual

hard disk that hosts a specially prepared—sysprepped, in the case of Windows-based operating systems— version of an operating system.

You can create VM templates based on existing virtual machines deployed on a virtualization host managed by VMM, based on virtual hard disks stored in a VMM library, or by using an existing VM template.

VM templates have the following limitations:

- A VM template allows you to customize IP address settings, but you can only configure a static IP address for a specific VM from a pool when deploying that VM from the template.
- Application and SQL Server deployment are only used when you deploy a VM as part of a service.
- When creating a template from an existing VM, ensure that the VM is a member of a workgroup and is not joined to a domain.
- You should create a separate local administrator account on a VM before using it as the basis of a template. Using the built-in administrator account causes the sysprep operation to fail.

More Info

Virtual machine templates

You can learn more about virtual machine templates at https://docs.microsoft.com/en-us/system-center/vmm/vm-template?view=sc-vmm-2019.

VMM storage

VMM can use local and remote storage, with local storage being storage devices that are directly attached to the server and remote storage being storage available

through a storage area network. VMM can use:

- **File storage.** VMM can use file shares that support the SMB 3.0 protocol. This protocol is supported by file shares on computers running Windows Server 2012 and later. SMB 3.0 is also supported by third-party vendors of network-attached storage (NAS) devices.
- **Block storage.** VMM can use block-level storage devices that host LUNs (Logical Unit Numbers) for storage using either the iSCSI, Serial Attached SCSI (SAS), or Fibre Channel protocols.

VMM supports automatically discovering local and remote storage. This includes automatic discovery of:

- Storage arrays
- Storage pools
- Storage volumes
- LUNs
- Disks
- Virtual disks

Using VMM, you can create new storage from storage capacity discovered by VMM and assign that storage to a Hyper-V virtualization host or host cluster. You can use VMM to provision storage to Hyper-V virtualization hosts or host clusters using the following methods:

- **From available capacity.** Allows you to create storage volumes or LUNs from an existing storage pool.
- **From writable snapshot of a virtual disk.** VMM supports creating storage from writable snapshots of existing virtual disks.
- **From a clone of a virtual disk.** You can provision storage by creating a copy of a virtual disk. This uses storage space less

efficiently than creating storage from snapshots.

- **From SMB 3.0 file shares.** You can provision storage from SMB 3.0 file shares.

VMM supports the creation of a thin provisioned logical unit on a storage pool. This allows you to allocate a greater amount of capacity than is currently available in the pool, and it is only possible when:

- The storage array supports thin provisioning
- The storage administrator has enabled thin provisioning for the storage pool

VMM 2019 supports balancing of virtual disks across cluster-shared volumes to ensure that no single cluster shared volume is over-committed.

> More Info
>
> **Storage in VMM**
>
> You can learn more about storage in VMM at https://docs.microsoft.com/en-us/system-center/vmm/manage-storage?view=sc-vmm-2019.

> Inside OUT
>
> **Storage classifications**
>
> Storage classifications allow you to assign a metadata label to a type of storage. For example, you might name a classification used with a storage pool that consists of high-speed solid-state disks as Alpha, a classification used with Fibre Channel RAID 5 SAS storage as Beta, and iSCSI SATA RAID 5 as Gamma. The labels that you use should be appropriate to your environment.

VMM networking

A VMM logical network is a collection of network sites,

VLAN information, and IP subnet information. A VMM deployment needs to have at least one logical network before you can use it to deploy VMs or services. When you add a Hyper-V based virtualization host to VMM, one of the following happens:

- If the physical adapter is associated with an existing logical network, it remains associated with that network once added to VMM.
- If the physical adapter is not already associated with a logical network, VMM creates a new logical network, associating it with the physical adapter's DNS suffix.

You can create logical networks with the following properties:

- **One Connected Network.** Choose this option when network sites that compose this network can route traffic to each other, and you can use this logical network as a single connected network. You have the additional option of allowing VM networks created on this logical network to use network virtualization.
- **VLAN-Based Independent Networks.** The sites in this logical network are independent networks. The network sites that compose this network can route traffic to each other, though this is not required.
- **Private VLAN (PVLAN) Networks.** Choose this option when you want network sites within the logical network to be isolated independent networks.

You create network sites after you have created a VMM logical network. You use network sites to associate IP subnets, VLANs, and PVLANs with a VMM logical network.

More Info

VMM logical network

You can learn more about VMM logical networks at https://docs.microsoft.com/en-us/system-center/vmm/plan-network?view=sc-vmm-2019.

Logical switches

VMM logical switches store network adapter configuration settings for use with VMM-managed virtualization hosts. You configure the properties of one or more virtualization host's network adapters by applying the logical switch configuration information.

You should perform the following tasks before creating a logical switch:

- Create logical networks and define network sites.
- Install the providers for any Hyper-V extensible virtual switch extensions.
- Create any required native port profiles for virtual adapters that you use to define port settings for the native Hyper-V virtual switch.

When you configure a VMM logical switch, you configure the following:

- Extensions
- Uplinks
- Virtual Ports

Extensions

You use logical switch extensions to configure how the logical switch interacts with network traffic. VMM

includes the following switch extensions:

- **Monitoring.** Allows the logical switch to monitor but not modify network traffic.
- **Capturing.** Allows the logical switch to inspect but not modify network traffic.
- **Filtering.** Allows the logical switch to modify, defragment, or block packets.
- **Forwarding.** Allows the logical switch to alter the destination of network traffic based on the properties of that traffic.

Uplink port profiles

Uplink port profiles specify which set of logical networks should be associated with physical network adapters. In the event that there are multiple network adapters on a virtualization host, an uplink port profile specifies whether and how those adapters should participate in teaming. Teaming allows network adapters to aggregate bandwidth and provide redundancy for network connections.

Virtual port profiles

You use port classifications to apply configurations based on functionality. The following port profiles are available:

- **SR-IOV.** Allows a virtual network adapter to use SR-IOV (Single Root Input Output Virtualization)
- **Host Management.** For network adapters used to manage the virtualization host using RDP, PowerShell, or another management technology

- **Network Load Balancing.** To be used with network adapters that participate in Microsoft Network Load Balancing
- **Guest Dynamic IP.** Used with network adapters that require guest dynamic IP addresses, such as those provided by DHCP
- **Live Migration Workload.** Used with network adapters that support VM live migration workloads between virtualization hosts
- **Medium Bandwidth.** Assign to network adapters that need to support medium-bandwidth workloads
- **Host Cluster Workload.** Assign to network adapters that are used to support host clusters
- **Low Bandwidth.** Assign to network adapters that need to support low-bandwidth workloads
- **High Bandwidth.** Assign to network adapters that are used to support high-bandwidth workloads
- **iSCSI Workload.** Assign to network adapters that are used to connect to SAN resources using the iSCSI protocol

More Info

Port profiles and logical switches

You can learn more about port profiles at https://docs.microsoft.com/en-us/system-center/vmm/network-port-profile?view=sc-vmm-2019.

Virtual machine networks

In VMM, virtual machines connect to a VMM logical network through a VMM virtual machine network. You connect a virtual machine's network adapter to the virtual machine network rather than the logical network. You can have VMM automatically create an associated virtual machine network when you create a

logical network. If you have configured a logical network to support network virtualization, you can connect multiple VM networks to the logical network, and they will be isolated from each other. Also, you can configure virtual networks to support encryption.

You can use network virtualization to configure logical networks in such a manner that different VM tenants can utilize the same IP address space on the same virtualization host without collisions occurring. For example, tenant alpha and tenant beta use the 172.16.10.x address space when their workloads are hosted on the same virtualization host cluster. Even though tenant alpha and tenant beta have virtual machines that use the same IPv4 address, network virtualization ensures that conflicts do not occur.

When you configure network virtualization, each network adapter is assigned two IP addresses:

- **Customer IP address.** This IP address is the one used by the customer. The customer IP address is the address visible within the VM when you run a command such as *ipconfig or Get-NetIPConfiguration*.
- **Provider IP address.** This IP address is used by and is visible to VMM. It is not visible within the VM operating system.

MAC address pools

A MAC address pool gives you a pool of MAC addresses that can be assigned to virtual machine network

adapters across a group of virtualization hosts. Without MAC address pools, virtual machines are assigned MAC addresses on a per-virtualization host basis. While unlikely, it is possible the same MAC address may be assigned by separate virtualization hosts in environments with a very large number of virtualization hosts. Using a central MAC address pool ensures that doesn't happen. When creating a MAC address pool, you specify a starting and an ending MAC address range.

More Info

MAC address pools

You can learn more about MAC address pools at https://docs.microsoft.com/en-us/system-center/vmm/network-mac?view=sc-vmm-2019.

Static IP address pools

An IP address pool is a collection of IP addresses that, through an IP subnet, is associated with a network site. VMM can assign IP addresses from the static IP address pool to virtual machines running Windows operating systems if those virtual machines use the logical network associated with the pool. Static IP address pools can contain default gateway, DNS server, and WINS server information. Static IP address pools aren't necessary because VMs can be assigned IP address information from DHCP servers running on the network.

More Info

IP address pools

You can learn more about IP address pools at https://docs.microsoft.com/en-us/system-center/vmm/network-pool?view=sc-vmm-2019.

Private VLANS

VLANs segment network traffic by adding tags to packets. A VLAN ID is a 12-bit number, allowing you to configure VLAN IDS between the numbers 1 and 4095. While this is more than adequate for the majority of on-premises deployments, large hosting providers often have more than 5,000 clients, so they must use an alternate method to segment network traffic. A PVLAN is an extension to VLANS that uses a secondary VLAN ID with the original VLAN ID to segment a VLAN into isolated sub networks.

You can implement VLANs and PVLANs in VMM by creating a Private VLAN logical network. Private VLAN logical networks allow you to specify the VLAN and/or PVLAN ID, as well as the IPv4 or IPv6 network.

Adding a WDS to VMM

In a scalable environment, you'll need to add additional Hyper-V host servers on a frequent basis as either a standalone server or as part of a failover cluster to increase your capacity. While it's possible to use another technology to deploy new Hyper-V host servers to bare

metal, the advantage of integrating virtualization host deployment with VMM is that you can fully automate the process. The process works in the following general manner:

1. Discovery of the chassis occurs. This may be through providing the chassis network adapter's MAC address to VMM.
2. The chassis performs a PXE boot and locates the Windows Deployment Services (WDS) server that you have integrated with VMM as a managed server role. When you integrate WDS with VMM, the WDS server hosts a VMM provider that handles PXE traffic from the bare metal chassis started using the VMM provisioning tool.
3. The VMM provider on the WDS server queries the VMM server to verify that the bare metal chassis is an authorized target for managed virtualization host deployment. In the event that the bare metal chassis isn't authorized, WDS attempts to deploy another operating system to the chassis. If that isn't possible, PXE deployment fails.
4. If the bare metal chassis is authorized, a special Windows PE (Preinstallation Environment) image is transmitted to the bare metal chassis. This special Windows PE image includes a VMM agent that manages the operating system deployment.
5. Depending on how you configure it, the VMM agent in the Windows PE image can run scripts to update firmware on the bare metal chassis, configure RAID volumes, and prepare local storage.
6. A specially prepared virtual hard disk (in either *.vhdx* or *.vhd* format) containing the virtualization host operating system is copied to the bare metal chassis from a VMM library server.
7. The VMM agent in the Windows PE image configures the bare metal chassis to boot from the newly placed virtual hard disk.
8. The bare metal chassis boots into the virtual hard disk. If necessary, the newly deployed operating system can obtain additional drivers not included in the virtual hard disk from a VMM library server.
9. Post-deployment customization of the newly deployed operating system

Technet24

occurs. This includes setting a name for the new host and joining an Active Directory Domain Services domain.

10. The Hyper-V server role is deployed, and the newly deployed virtualization host is connected to VMM and placed in a host group.

The PXE server needs to provide the PXE service through Windows Deployment Services. When you add the VMM agent to an existing Windows Deployment Services server, VMM only manages the deployment process if the computer making the request is designated as a new virtualization host by VMM.

To integrate the WDS server with VMM to function as the VMM PXE server, you need to use an account on the VMM server that is a member of the local Administrators group on the WDS server. PXE servers need to be on the same subnet as the bare metal chassis to which they deploy the virtualization host operating system.

VMM host groups

Host groups allow you to simplify the management of virtualization hosts by allowing you to apply the same settings across multiple hosts. VMM includes the All Hosts group by default. You can create additional host groups as required in a hierarchical structure. Child host groups inherit settings from the parent host group. However, if you move a child host group to a new host group, the child host group retains its original settings,

except for any PRO configuration. When you configure changes to a parent host group, VMM provides a dialog box asking if you would like to apply the changed settings to child host groups.

You can assign network and storage to host groups. Host group networks are the networks that are assigned to the host group. These resources include IP address pools, load balances, logical networks, and MAC address pools. Host group storage allows you to allocate logical units or storage pools that are accessible to the VMM server for a specific host group.

More Info

VMM host groups

You can learn more about VMM host groups at https://docs.microsoft.com/en-us/system-center/vmm/host-groups?view=sc-vmm-2019.

VMM virtualization host requirements

To be able to configure a bare metal hardware chassis so that it can function as a VMM managed Hyper-V virtualization host, the hardware chassis needs to meet the following requirements:

- **X64 processor.** This needs to support hardware-assisted virtualization and hardware-enforced Data Execution Prevention (DEP). In some cases, it may be necessary to enable this support in BIOS.
- **PXE boot support.** The hardware chassis must be able to PXE boot from a PXE-enabled network adapter. The PXE-enabled

network adapter must be configured as a boot device.

- **Out-of-band (OOB) management support.** System Center VMM can discover and manage the power states of hardware chassis that support BMC (Baseboard Management Controller). VMM supports the following protocols:
    - SMASH (Systems Management Architecture for Server Hardware) version 1 over WS-Management
    - DCMI (Datacenter Management Interface) version 1.0
    - IPMI (Microsoft Intelligent Platform Management Interface) version 1.5 or version 2.0

Placement rules

Placement rules allow you to configure how VMM identifies a suitable host for a VM deployment. Usually, this is based on the available resources on the virtualization host or the host cluster. By configuring host group placement rules, you can create rules that dictate the conditions under which a new VM can be placed on a virtualization host in the host group.

To add a placement rule, edit the properties of the host group, and on the Placement tab, click Add. You then specify a custom property and one of the following requirements:

- Virtual Machine Must Match Host
- Virtual Machine Should Match Host
- Virtual Machine Must Not Match Host
- Virtual Machine Should Not Match Host

Host reserves

Host reserves allow you to configure the resources that VMM should set aside for the host operating system. When VMM is going to place a new VM on a host, that host must be able to meet the VM's resource requirements without exceeding the configured host reserves. You can configure host reserves for:

- CPU
- Memory
- Disk I/O
- Disk Space
- Network I/O

Dynamic optimization

Dynamic optimization allows virtualization host clusters to balance workloads by transferring VMs between nodes according to the settings configured at the host group level. Whether the transfer occurs depends on whether the hardware resources on a node in the virtualization host cluster fall below the configured settings. Dynamic optimization applies only to clustered virtualization hosts and does not apply hosts that are not members of a cluster.

More Info

Virtual Machine Manager

You can learn more about Virtual Machine Manager at https://docs.microsoft.com/en-us/system-center/vmm/overview?view=sc-vmm-

2019.

INFRASTRUCTURE CONFIGURATION AS CODE

As infrastructure—including servers, storage, and networks—has increasingly become virtualized, configuration of that infrastructure has increasingly become defined by code. This means rather than having to build out a template image, it's increasingly possible to define the properties of a server and the workload it hosts through a template written in a format such as JavaScript Object Notation (JSON), Ruby, or XML.

For example, with Azure, it's possible to load a JSON template that defines the properties of a virtual machine workload as well as the network and storage settings associated with that workload. It also becomes increasingly possible to alter the properties of a running workload by modifying and reapplying the template rather than going through the workload's management interface and applying changes manually.

Other advantages of infrastructure configuration as code include:

- You can version that code so that you can roll back changes by simply applying an earlier version of a set of templates in the event that a set of changes cause unforeseen problems.
- You can test the infrastructure code in a development environment

before deploying it in a production environment. For example, you can ensure the infrastructure code to build your production VM workloads functions as expected on a test Hyper-V server before deploying the same code to your production Hyper-V failover cluster.

The WSLab project, hosted on GitHub at *https://github.com/microsoft/WSLab*, provides a set of PowerShell scripts that allow for the automatic creation of entire Hyper-V-based virtual machine testing labs. Running scripts allows you to do the following:

- Create updated base operating system virtual machine hard disk images by:
  - Automatically downloading required operating system evaluation software from the Internet.
  - Using the WIM images stored on those evaluation versions of Windows operating systems, create sysprepped, virtual machine parent disks.
  - Applying the most recent software updates to ensure that the base operating system virtual machine hard disks are up to date.
- Create Hyper-V internal and external networks.
- Create fully functional virtual machines, including:
  - Creating new virtual machines from the base operating system virtual hard disks.
  - Connecting these newly created virtual machines to appropriate Hyper-V internal and external networks.
  - Injecting Desired State Configuration resources into the VMs from the host machine.
  - Injecting *Unattend.xml* files into the newly created VMs.

- Invoking the DSC Local Configuration Manager to trigger the application of the DSC resources.

In Virtual Machine Manager, you can deploy a multitier virtual machine service by creating, configuring, and deploying the appropriate VM and service templates. Tools such as WSLab demonstrate that it's possible to deploy entire environments quickly from scratch just with scripts, tools such as DSC, and the base operating system's installation media.

Inside OUT

Infrastructure Developer

Infrastructure Developer provides an alternative way of describing some of the descriptions that are associated with DevOps and Site Reliability Engineering (SRE). Many of us who came up through IT operations are suspicious of any new approach, simply because most of us have been in the industry some time, and we've seen different fads come and go. Many IT professionals staff actually work in environments that don't have a single developer. That means the idea of adopting "a DevOps" or "SRE" mindset seems irrelevant to the way they actually do their jobs on a day-to-day basis because deployments of brand-new software happen infrequently rather than every day. However, because configuration is increasingly defined through code, it's not unreasonable for IT Operations staff who work with configuration code to be called "infrastructure developers." Rather than creating and modifying application code, we'll be creating and modifying the code that defines our organization's infrastructure. For the most part, we're doing the same sort of things we were doing when we got into the industry in the 1990s, but we're just using a new set of code-based tools now.

DESIRED STATE CONFIGURATION

Desired State Configuration (DSC) is a configuration-management technology that's built into Windows Server through PowerShell version 5 and higher. At a high level, a DSC configuration file describes how a

computer should be configured rather than just a script that configures a computer. A component of DSC, called the Local Configuration Manager, takes a configuration file that describes how the server should be configured and ensures the server is configured in that way.

Almost all the tools required to use DSC are built into Windows Server 2019. You can also use DSC in a homogeneous Windows environment without the need to deploy a third-party operating system. A current disadvantage of DSC is that it can be challenging and complicated to configure and that it isn't as developed as other third-party configuration tools where the tool forms the basis of an organization's business model.

Not only does DSC perform initial configuration of a server, but it can also enforce that configuration. You can configure DSC so that changes made to the server's configuration not made through DSC will automatically roll back the next time DSC runs. This makes DSC an excellent tool for ensuring that servers in your organization retain configurations that meet compliance requirements. If you manage configuration entirely through DSC, you perform configuration changes by modifying DSC configurations rather than signing onto a server and modifying the configuration manually.

DSC can also work in audit mode. When using DSC in audit mode, you can generate log data describing how a

computer's configuration deviates from that specified in a configuration file. This can be useful in determining how a computer's configuration may have drifted from a baseline, and it can also be used to determine what updates need to be made to a DSC configuration file when modifications are made to a server that you wish to retain.

Inside OUT

Idempotency

Idempotency is an important concept when using tools including Desired State Configuration, Chef, and Puppet. Idempotency allows the same technique to be applied multiple times, always producing the same result. This is important because these tools apply a specific configuration to a system on a periodic basis. In the case of configuration, applying the configuration tool should either bring the existing configuration into line with the desired configuration, or if the desired configuration is already set, should not alter that configuration further.

DSC configuration files

DSC configuration files describe how a computer should be configured. Configuration files require at least one node and require at least one configuration item per node. A configuration file can have multiple nodes, with different configuration items in each node. Generally, you create one configuration for each different computer type in your organization. For example, you create one for computers that host the DHCP server role, and you create another for computers that function as file servers.

The following hypothetical configuration would ensure that the WINS and Windows Server Backup features were present on a computer.

click here to view code image

```
Configuration TestDscConfiguration {
Node “TEST-SVR” {
WindowsFeature InstallWINS {
Ensure = "Present"
Name = "WINS"
}
WindowsFeature InstallBackup {
Ensure = "Present"
Name = "Windows-Server-Backup"
}
}
}
TestDscConfiguration
```

When you run a DSC configuration script, it generates an MOF file. The MOF file is what the LCM (Local Configuration Manager) uses when determining which configuration to apply to a computer. You compile a configuration script by *dot-sourcing* it. For example, if you saved the above script as *TestConfiguration.ps1* in the *c:\DSCTEST* folder and *.sourced* it, the script would create a new folder named *TestDscConfiguration*; a new file named *TEST-SVR.mof* would be created in the same folder.

MOF files contain the configuration data that you want to apply to the Windows Server node. The MOF files are created by running DSC configuration files. You can deploy MOF files to the LCM on a Windows Server node

by using the Push method to inject a file into a running computer, or you can configure a pull server from which nodes can retrieve MOF files and DSC resources.

> **More Info**
>
> **DSC Configuration**
>
> You can learn more about DSC configuration files at https://docs.microsoft.com/en-us/powershell/scripting/dsc/configurations/configurations.

Local Configuration Manager

The Local Configuration Manager (LCM) is present on each computer that uses DSC. The LCM is the component that takes a MOF file and DSC resources and uses them to enact a DSC configuration. You can view the current settings of the LCM using the *Get-DscLocalConfigurationManager* cmdlet, as shown in Figure 3-16.



Figure 3-16 LCM settings

DSC resources

DSC resources are used by the LCM. They can be PowerShell scripts or an item compiled into a DLL. Resources are stored in resource modules, which are distributed to the LCM to allow it to carry out the configuration instructions stored in the MOF. DSC resource modules must be present on any computer where they are used to create or apply a configuration script.

The PowerShell Gallery hosts a wealth of DSC resources that can form the basis of your DSC configuration scripts and configurations. DSC clients can obtain resources directly from DSC pull servers. If you are using another method of deploying configurations—such as the push method or copying method—you'll need to ensure that the appropriate DSC resources are stored on the computer on which you want to configure using DSC.

DSC push model

You can push a MOF file to a target computer using the *Start-DscConfiguration* cmdlet and a PowerShell remoting connection. This is effective in smaller and test environments and can be used to push MOF files to target computers. For example, you might use this to push a MOF file to a test virtual machine running under client Hyper-V on your Windows 10 computer. The DSC push model requires that PowerShell remoting be

enabled on each computer that is to be the target of a DSC push action. This is less of an issue with Windows Server 2019 because PowerShell remoting is enabled by default on this operating system.

Each time you use *Start-DscConfiguration* to push a new configuration to a node, the push action overwrites any existing configuration on the target node. Once you push the new configuration, the LCM on the target node begins immediately processing the MOF file. This makes the DSC push model a great way of testing configurations because you don't have to wait for the DSC cycle to restart to process the update, something that is necessary when using DSC pull servers.

A substantial drawback to using DSC push is that any resources that are to be used by the MOF file must already be present on the node you are configuring through push. Should you wish to use new resources for a configuration, you'll need to transfer them to the target node prior to pushing the new configuration.

DSC pull server

DSC pull servers are computers running Windows Server 2012 or later that function as central DSC repositories. Rather than push DSC configurations to each node, each node can be configured to poll a DSC pull server, downloading new MOF files and resources as they become available. DSC pull servers are almost

always configured as web servers. It's possible to use a standard file server as a pull server, but file servers lack some of the functionality of IIS-based pull servers, such as the ability to function as compliance servers. Compliance servers allow you to view a central repository of information about each node's compliance with DSC configuration settings.

When you deploy a pull server, you specify a location to store MOF files and a location to store DSC resource files in zipped format. MOF files stored on a pull server use a Globally Unique Identifier (GUID) for identification. Each MOF file and DSC resource module zip file must be stored with an associated checksum file. You can generate checksum files using the *New-DSCChecksum* cmdlet.

More Info

DSC

You can learn more about DSC at https://docs.microsoft.com/en-us/powershell/scripting/dsc/overview/dscforengineers?view=powershell-7.

CHEF INFRA SERVER

Chef Infra Server (referred to as Chef throughout this section) is a third-party configuration management solution. Unlike DSC, where you must build most of the infrastructure out yourself and you're generally reliant upon a community to build resources for you unless you

can build them yourself, Chef provides you with software and resources that you can use to manage the configuration of computers running a variety of operating systems. Chef allows you to manage servers running both Windows and Linux based operating systems.

> **Inside OUT**
>
> **Chef in Azure**
>
> While you can spin up your own Linux instance in Azure and install Chef on that IaaS VM, there is also a fully preconfigured Linux IaaS VM in the Azure Gallery.

Chef servers

Chef uses a client/server architecture. You manage the configurations that you want to deploy on a server as cookbooks, and the clients of those servers apply those cookbooks and recipes to Chef nodes to meet your configuration objectives. When deploying Chef, you have two primary options that you are responsible for managing: Hosted Chef and a Chef Server.

Hosted Chef

Hosted Chef is an instance of Chef Server hosted in the cloud by the Chef organization. The big benefit of Hosted Chef is that like any software as a service (SaaS) solution, the Chef organization takes care of the entire infrastructure. Hosted Chef requires that clients have a connection to the Internet, so they can communicate

with the server. This makes Chef great for workloads running on Internet-connected networks but problematic for organizations who need configuration management on isolated networks.

Installing Chef Server

You can deploy Chef Server on-premises on a server running a Linux operating system. At present, you can't deploy Chef Server on a computer running Windows Server 2019, though you can manage a Chef Server from any computer with a browser.

Chef Server is supported on the following operating systems:

- Centos 6 and 7
- Red Hat Enterprise Linux 6 and 7
- SUSE Linux Enterprise Server 11 and 12
- Ubuntu 16.04 and 18.04
- Oracle Enterprise Linux 6 and 7

To install Chef Server on Linux, perform the following tasks:

1. Obtain the appropriate Chef Server package for the Linux server that you will use to host Chef Server from *https://downloads.chef.io/chef-server and* store it in a temporary directory on the Linux server that you want to function as the Chef server. In the following example, the package for Ubuntu LTS 16.04 is used.
2. Install the package. On a Linux server running Ubuntu, you would do this with a command similar to the following:

click here to view code image

```
Sudo dpkg -i ./chef-server-core_12.19.31-
1_amd64.deb
```

3. Once the package has installed, run the following command to begin configuration, which takes a moderate amount of time:

   click here to view code image

   ```
   Sudo chef-server-ctl reconfigure
   ```

4. Once configuration has completed, create a new user. For example, run the following command to create a new user named *sydneymelborne* and store the certificate associated with that user in the *~/certs* folder of the account used to deploy Chef server. (Note that the username is in lower case, and the password does need to be passed to the console.)

   click here to view code image

   ```
   Sudo chef-server-ctl user-create
   sydneymelbourne Sydney Melbourne
   sydney@company.pri P@ssw0rd --filename
   ~/certs/sydneymelbourne.pem
   ```

5. You can create an organization and associate the user that you created with that organization using the following command. In this case, the organization name is *contoso*. (Be sure that you enter the organization name in lowercase.)

   click here to view code image

   ```
   Sudo chef-server-ctl org-create contoso
   "Contoso-Chef" --association_user
   sydneymelbourne --filename ~/certs/contoso.pem
   ```

6. To install the web-based management console, run the following command:

click here to view code image

```
Sudo chef-server-ctl install chef-manage
```

7. Then use the following commands to configure the console:

   click here to view code image

   ```
   Sudo check-server-ctl reconfigure
   Sudo chef-manage-ctl reconfigure --accept-
   license
   ```

8. You can then verify that the server is functioning by using a standards-compliant web browser to navigate to the FQDN of the Chef server, as shown in Figure 3-17. You should sign in with the account created in step 4.



Figure 3-17 Chef console

9. If you navigate to the Administration node, you'll be able to view the organization that you created and the users who are members of that organization, as shown in Figure 3-18.

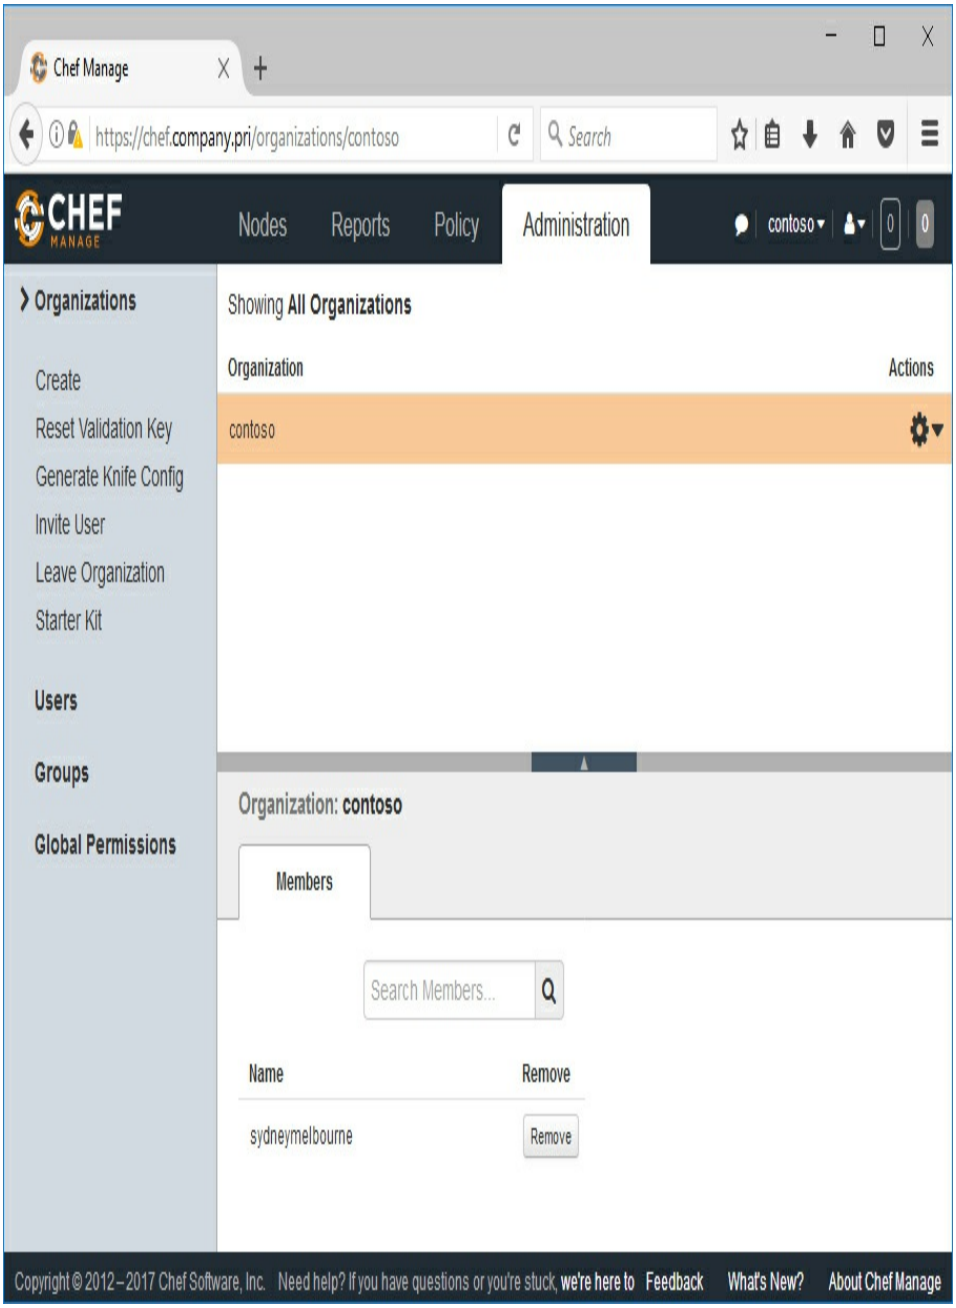


Figure 3-18 Chef organization

10. Once you have verified the configuration of the console, you can install and configure Chef's reporting capability by running the following commands:

    click here to view code image

```
Sudo chef-server-ctl install opscode-reporting
Sudo chef-server-ctl reconfigure
Sudo opscode-reporting-ctl reconfigure --
accept-license
```

More Info

Installing Chef Server

You can learn more about installing Chef Server at https://docs.chef.io/install_server.html.

Chef Development Kit

You use the Chef Development Kit (ChefDK) to author recipes, cookbooks, and policies. You can install the Chef Development Kit by downloading the relevant files from Chef's website; also, if you have Chocolatey configured, you can install the Chef Development Kit, as described later in this chapter, by running the following command from an elevated command prompt:

```
Choco install -y chefdk
```

Performing this installation gives you access to the *ChefDK* PowerShell prompt, which allows you to use PowerShell to interact with ChefDK.

Starter Kit

The Starter Kit is a package that you can download from your Chef server that contains all the information and files needed to connect to an organization hosted on

your Chef server. You can find this package by clicking Starter Kit in the sidebar menu when you have the organization selected in the Administration tab, as shown in Figure 3-19.

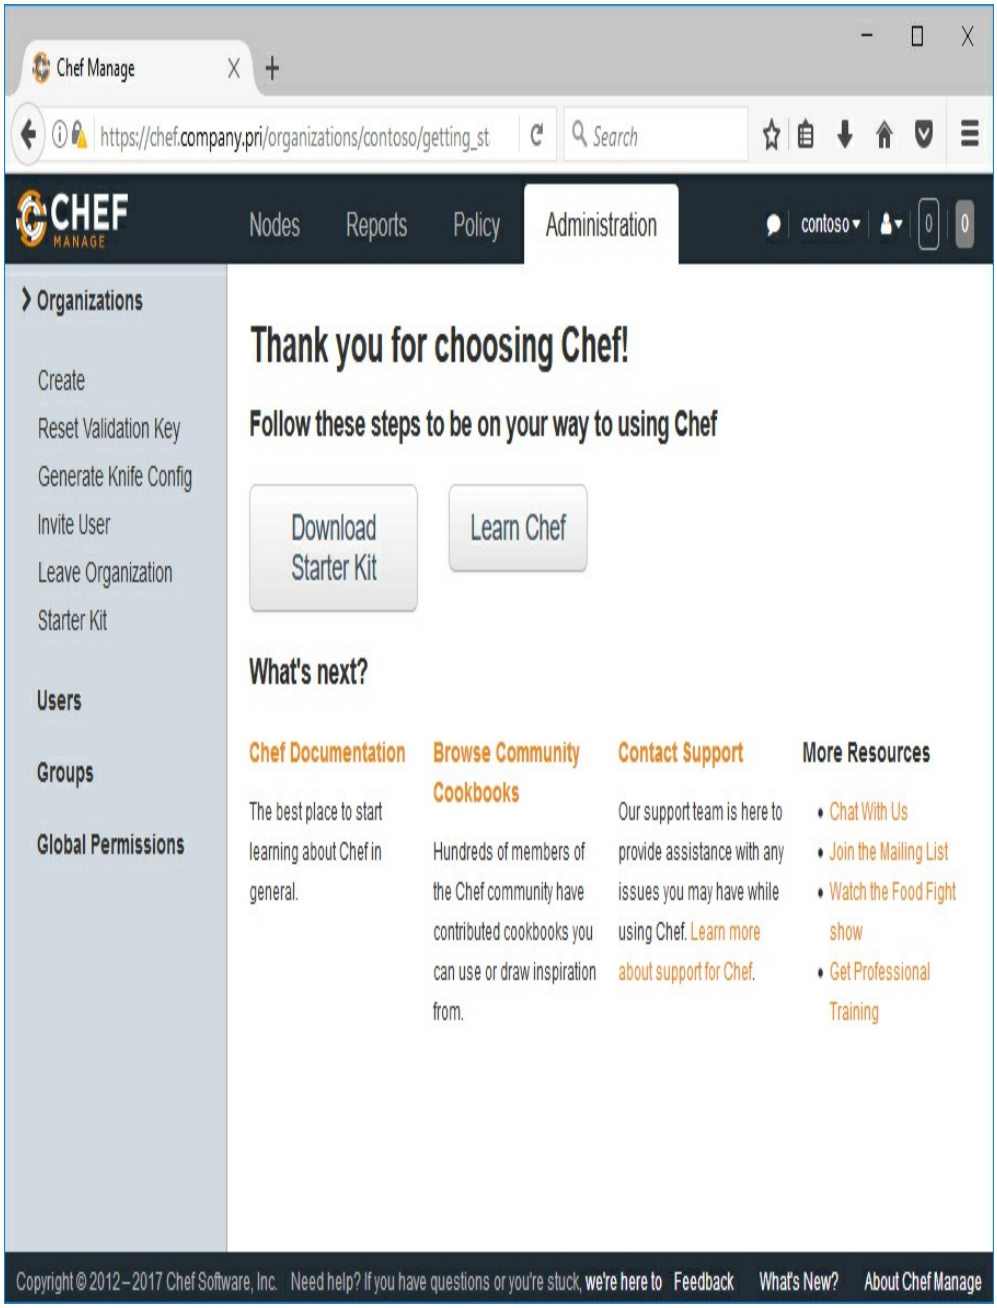


Figure 3-19 Starter Kit

The Starter Kit is a zipped file that has the basic structure of a Git repository. You can add this file structure to a Git repository if you want to implement source control. The file *knife.rb*, shown opened in Visual Studio Code in Figure 3-20, contains the settings of the organization on the Chef Server.

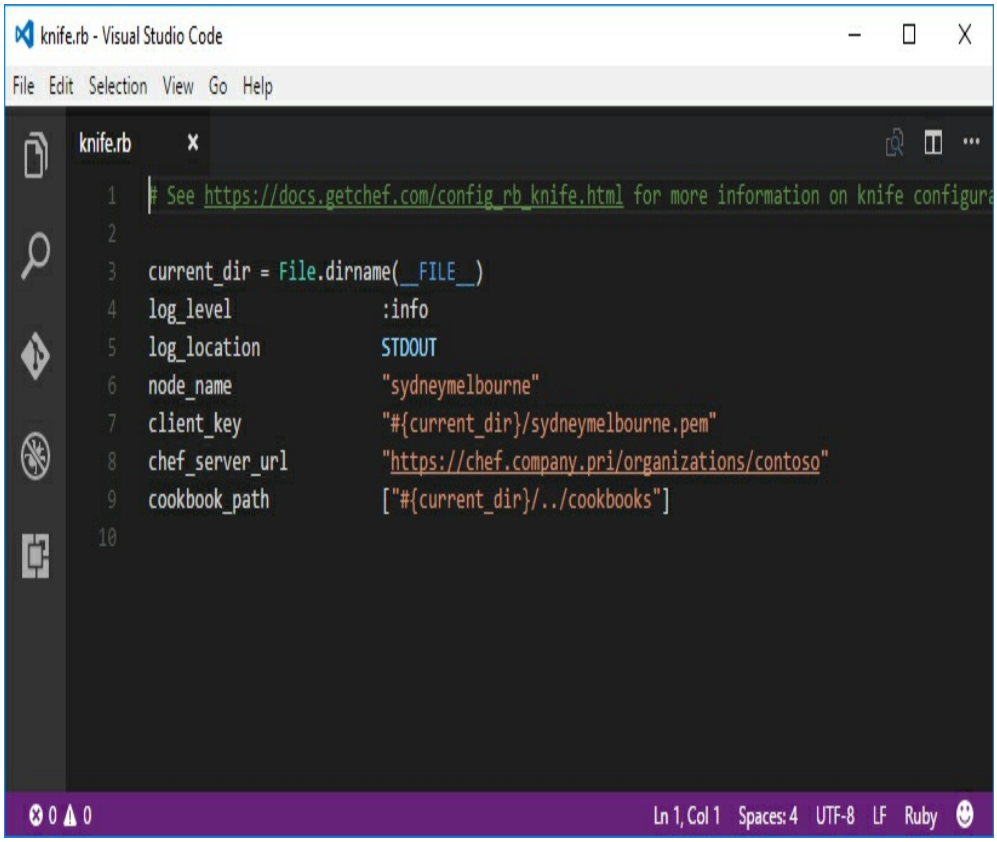


Figure 3-20 Chef organization settings

To use ChefDK with the Chef Server, you need to have the TLS certificates assigned to the server present on the

Technet24

workstation on which you have ChefDK installed. Even if your computer trusts the TLS certificate, you'll still need a local copy of the public certificate to use the Knife utility in the ChefDK to interact with your organization's Chef server. You can retrieve the TLS certificate that is assigned to the Chef server by running the following command in the *chef-repo* folder from the Chef Development Kit shell that becomes available when you install the ChefDK. The certificate is added to the trusted certificates area of the Chef Git repository folder structure that was created when you unzipped the Starter Kit, as shown in Figure 3-21.

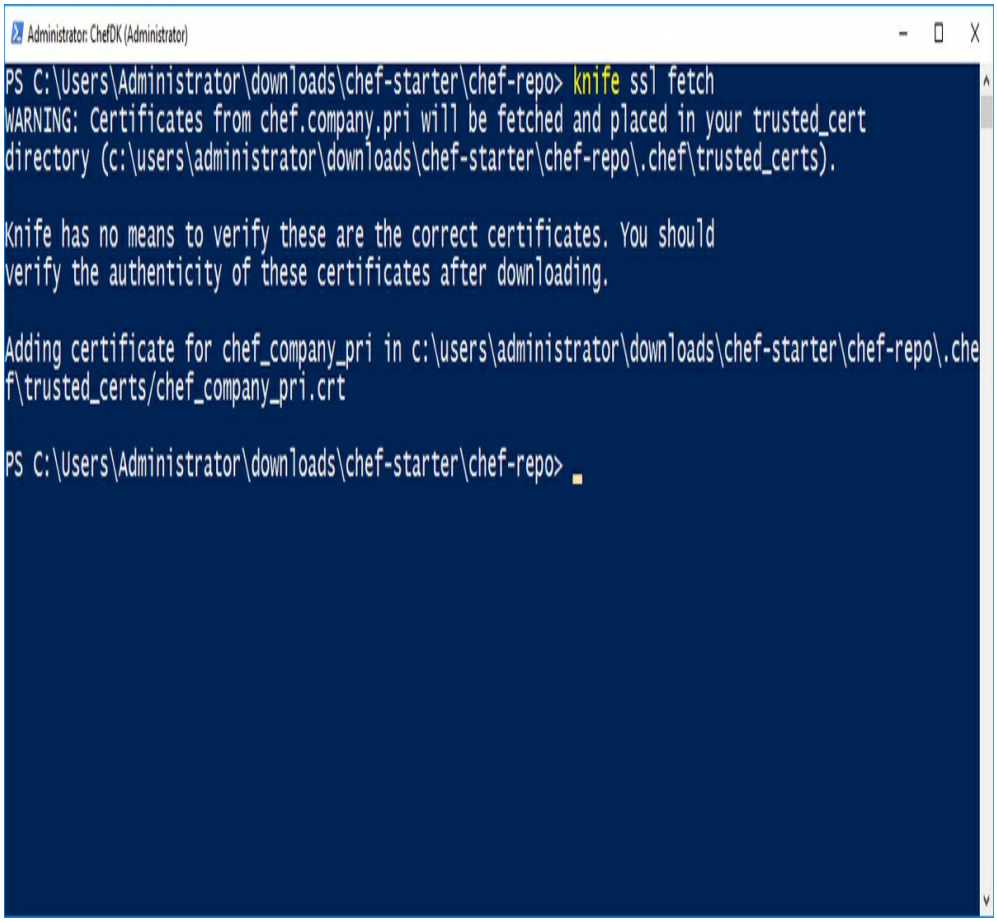


Figure 3-21 Acquire Chef certificate

You can verify that the certificate has installed correctly by running the following command from the *chef-repo* folder using the *ChefDK* PowerShell prompt:

```
Knife SSL check
```

Recipes and cookbooks

You use a Chef recipe to specify the configuration a Chef node should have. You bundle recipes into cookbooks

and deploy cookbooks against Chef clients. A default Windows cookbook is available from the Chef website. You can download the current example Windows cookbook from Chef and the ohai cookbook (which is a dependency of the Windows cookbook) by running the following command from the *chef-repo\cookbooks* folder using the ChefDK PowerShell prompt:

click here to view code image

```
Knife cookbook site download windows
Knife cookbook site download ohai
```

Doing so downloads both the Windows cookbook and the ohai cookbook in *.tar* and *.zip* format. You can use a utility such as 7-Zip to extract the cookbooks. Once extracted, place the cookbooks in the cookbooks folder.

You can use the Windows cookbook as a base cookbook on which you can create additional specific cookbooks. For example, to create a cookbook named *wins-windows* that installs the WINS server feature on a Windows Server computer, perform the following steps:

1. In the *Chef-repo\Cookbooks* folder, run the following ChefDK PowerShell command:

   click here to view code image

   ```
   Chef generate cookbook wins-windows
   ```

2. Open the *Chef-repo\cookbooks\wins-windows\recipes\default.rb* file in Visual Studio Code and add the following lines, as shown in Figure 3-22.

click here to view code image

```
windows_feature “WINSRuntime” do
action :install
end
```

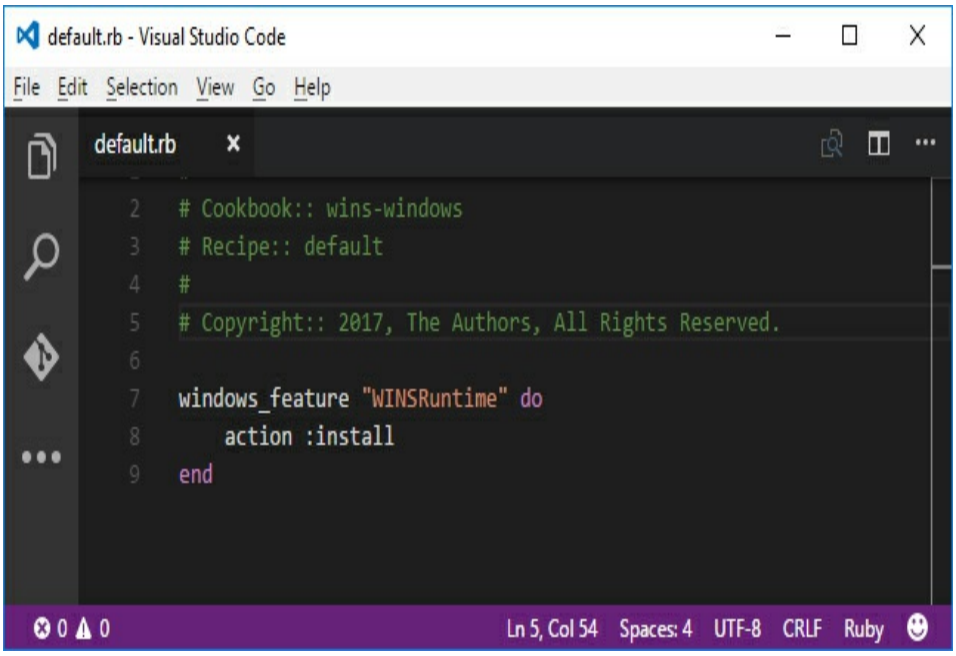


Figure 3-22 Default.rb

3. To configure the new cookbook’s dependency on the existing example Windows cookbook, open the *Chef-repo\cookbooks\wins-windows\metadata.rb* file and add the following line, as shown in Figure 3-23:

```
depends ‘windows’
```

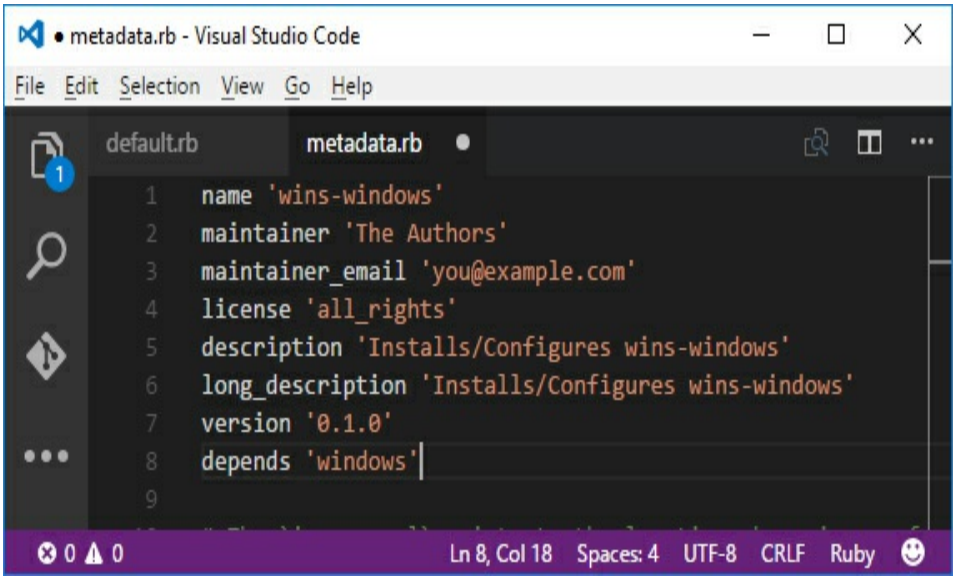


Figure 3-23 Configure dependencies

4. In the *Chef-repo\Roles* folder, create a new file named *wins-windows.rb* and enter the following text, as shown in Figure 3-24.

click here to view code image

```
name "wins-windows"
description "test WINS install"
run_list "recipe[wins-windows]"
```

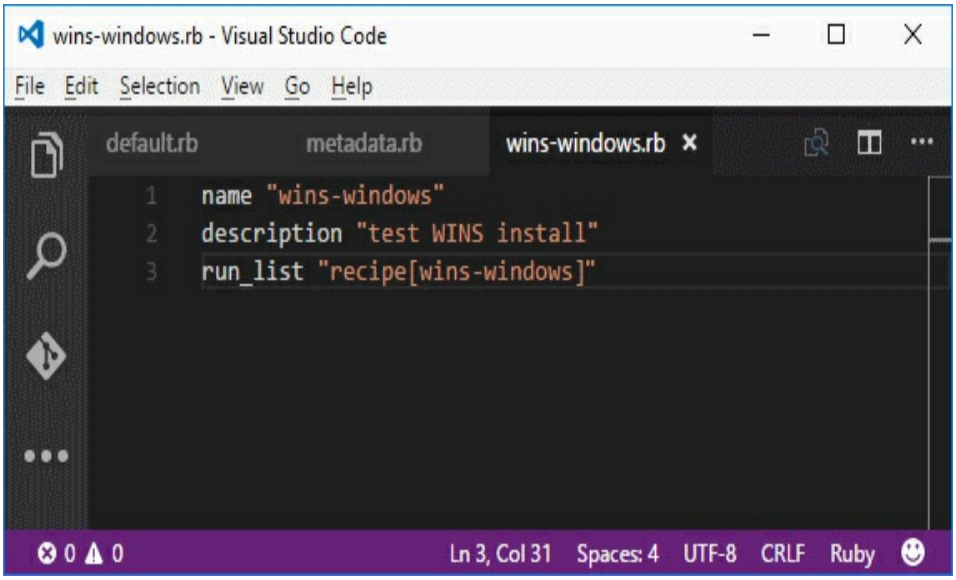


Figure 3-24 wins-windows.rb

5. From the *Chef-repo\Cookbooks* folder, run the following command to push all current cookbooks from the computer that hosts ChefDK to the Chef server:

```
Knife cookbook upload --all
```

6. Next, change into the *Chef-repo\Roles* folder and run the following command to create the relevant role:

   click here to view code image

```
Knife role from file .\wins-windows.rb
```

7. You can verify that cookbooks are present on the Chef Server on the Policy tab of the Chef console, as shown in Figure 3-25.

Technet24

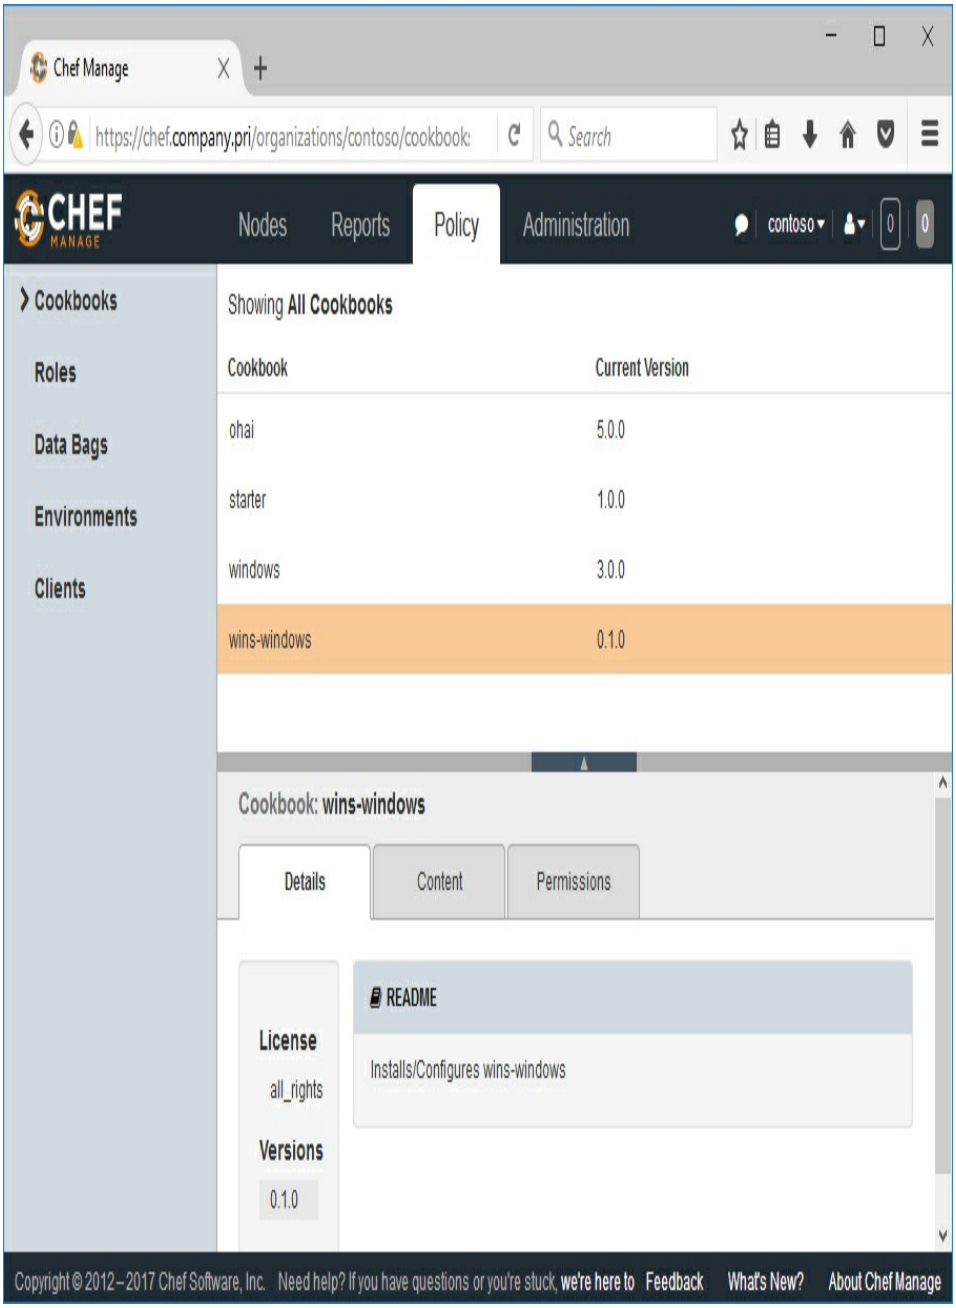


Figure 3-25 All Cookbooks

8. You can verify that the roles have uploaded correctly by checking the Roles section when the Policy tab is selected, as shown in Figure 3-26.

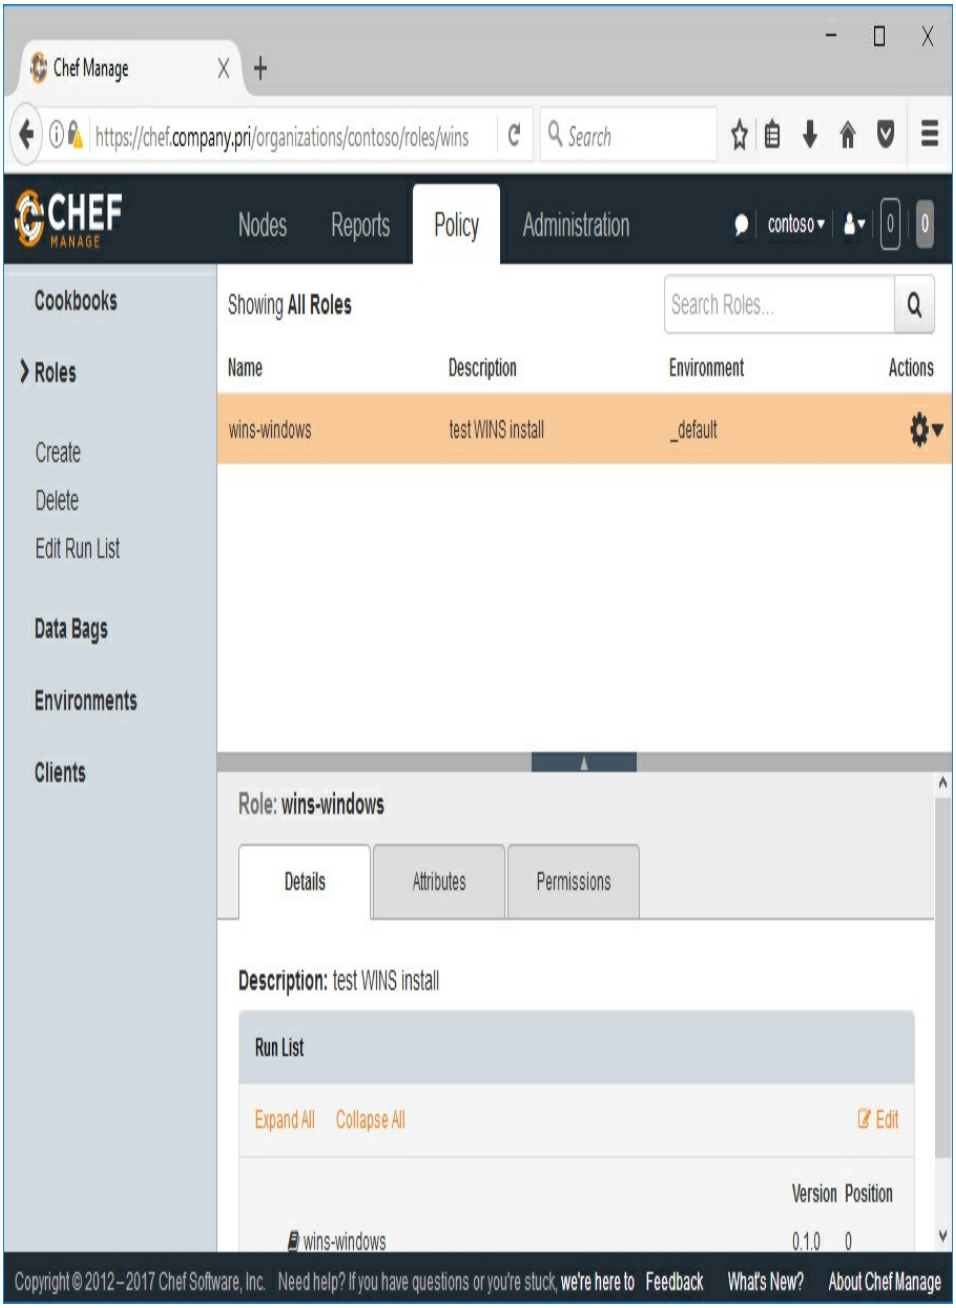


Figure 3-26 Chef roles

Deploying Chef agents

Deploying the Chef agent to a computer makes that computer a node in the Chef organization. The most

reliable method of ensuring that a Chef node has all the relevant certificates and configuration settings required to communicate with a Chef Server is to perform a *knife* bootstrap deployment from a computer with ChefDK installed.

Prior to performing this deployment, you should update your tools to ensure that you have the latest versions of the deployment tools available. To do this, run the following command from the ChefDK prompt in the *Chef-repo* folder:

```
Gem update knife-windows
```

To deploy an agent to a Windows Server, you need to ensure that WinRM is configured on both the source and the target computer, and run the following command (where *prodserver* is the name of the target computer):

click here to view code image

```
Knife bootstrap windows winrm prodserver.company.pri --winrm-user
Administrator --winrm-password P@ssw0rd --node-name prodserver
```

You can verify that the client is installed on a node and registered with the Chef server by checking the Nodes tab in the management console.

Deploying Chef cookbooks and recipes

By applying a role to a node, you deploy the contents of

that role, including cookbooks and their associated recipes. You can apply a role or recipe directly to a Chef node using the console by selecting the node in the console, editing the node's Run List, and then assigning a Recipe or a Run List, as shown in Figure 3-27.

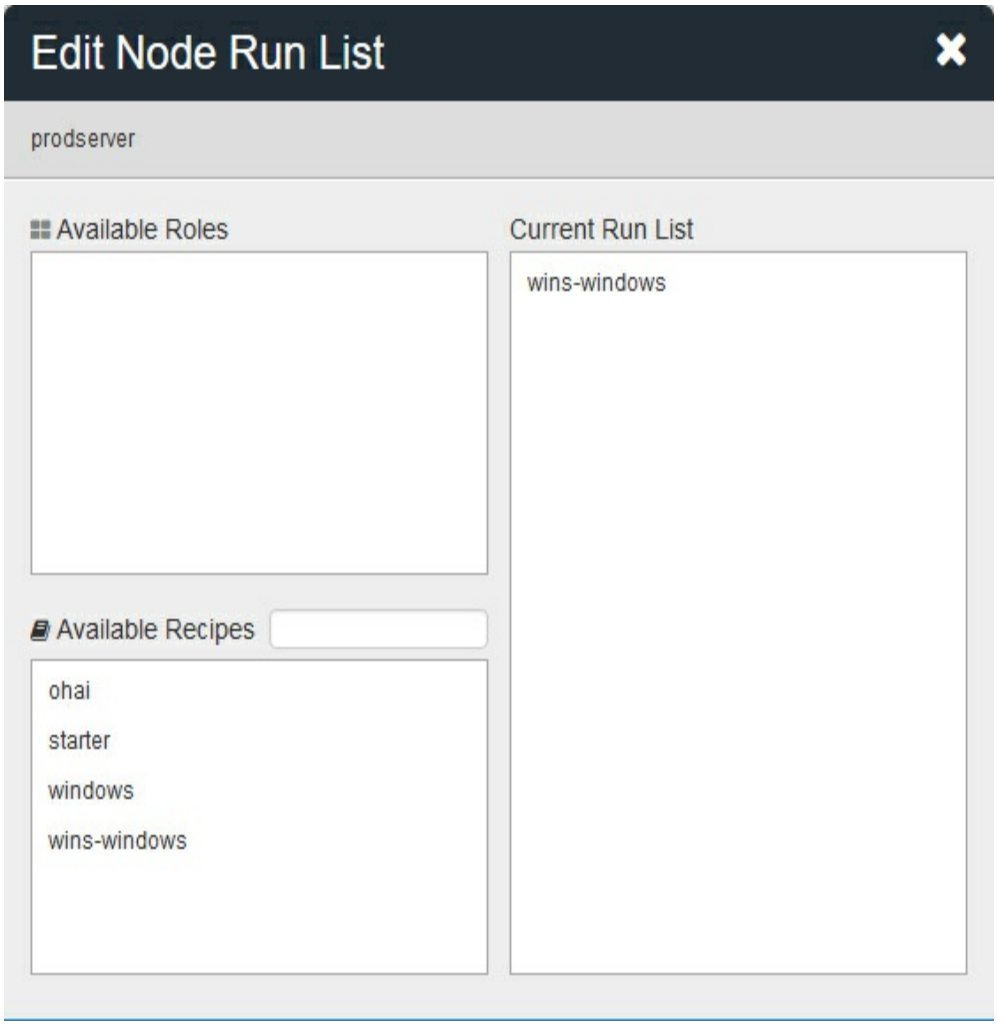


Figure 3-27 Edit Node Run List

The Run List will apply the next time the run applies. You can also manually trigger a run on the client by

Technet24

executing *chef-client.bat* from the C:\Windows\System32 folder of an administrative command prompt. You can verify that the run has completed correctly by checking the Reports tab on the Chef console.

More Info

Chef

You can learn more about Chef at https://www.chef.io.

PUPPET

Puppet is suitable for heterogeneous environments where you need to manage the configuration of both Linux and Windows servers. Puppet uses an architecture where a server named *Puppet Master* provides a central management point to computers that have the *Puppet* agent installed.

More Info

Puppet

You can learn more about Puppet at https://puppet.com/.

Puppet Master Server

The Puppet Master Server is only currently supported on hosts running a Linux- or UNIX-based operating system. If you are responsible for an environment that only has Windows servers, Puppet may not be a good fit

because the Puppet Master Server can only be deployed on computers running the following operating systems:

- RedHat Enterprise Linux, CentOS, Scientific Linux, and Oracle Linux Version 7 and 6 (x86_64)
- Ubuntu version 16.04 (amd64)
- SLES 12 SP1 (x86_64)

> **Inside OUT**
>
> **Linux distro selection**
>
> Chef and Puppet can be run on a variety of Linux distributions. Most heterogeneous environments have standardized on a single distribution of Linux, so your first choice for a Chef or Puppet host is likely to be the distribution you are already using. If you don't have Linux servers deployed in your environment and do want to use Chef or Puppet, you might consider one of the Ubuntu LTS options as a host. This is because a version of Ubuntu is now available in Windows 10, and there are likely to be more synergies involved in running Ubuntu as part of Windows Subsystem for Linux and running Ubuntu as the host server for Puppet or Chef.

The Puppet Master Server stores configuration for all Puppet clients on the network. The Puppet Master Server requires at least 6 GB of RAM, four cores, a hostname that is accessible to other computers on the network, and accessibility on port 3000. You'll likely need to manually configure the firewall to be open on the appropriate port. The administrator of the Puppet Master Server must have root user or have Sudo privileges on the host server.

To install the Linux computer as a Puppet Server, first connect to the server with an account that has the ability to elevate privileges and verify the computer's

hostname. Then perform the following steps:

1. You'll need to download a tarball that is specific to your Linux distribution from the Puppet website to the Linux computer. Extract the tarball downloaded from the Puppet website using the following command:

   ```
   Tar -xvf puppet-enterprise*
   ```

2. This extracts the Puppet installation files. Change to the directory that was created when you extracted the tarball, and start the installation script by typing the following command:

   click here to view code image

   ```
   Sudo ./puppet-enterprise-installer
   ```

3. When prompted, select the guided installation. Once the installation is complete, you'll need to open a browser and navigate to port 3000 on the address of the newly deployed Puppet Server. You'll need to ignore any certificate errors and click Let's Get Started, as shown in Figure 3-28.

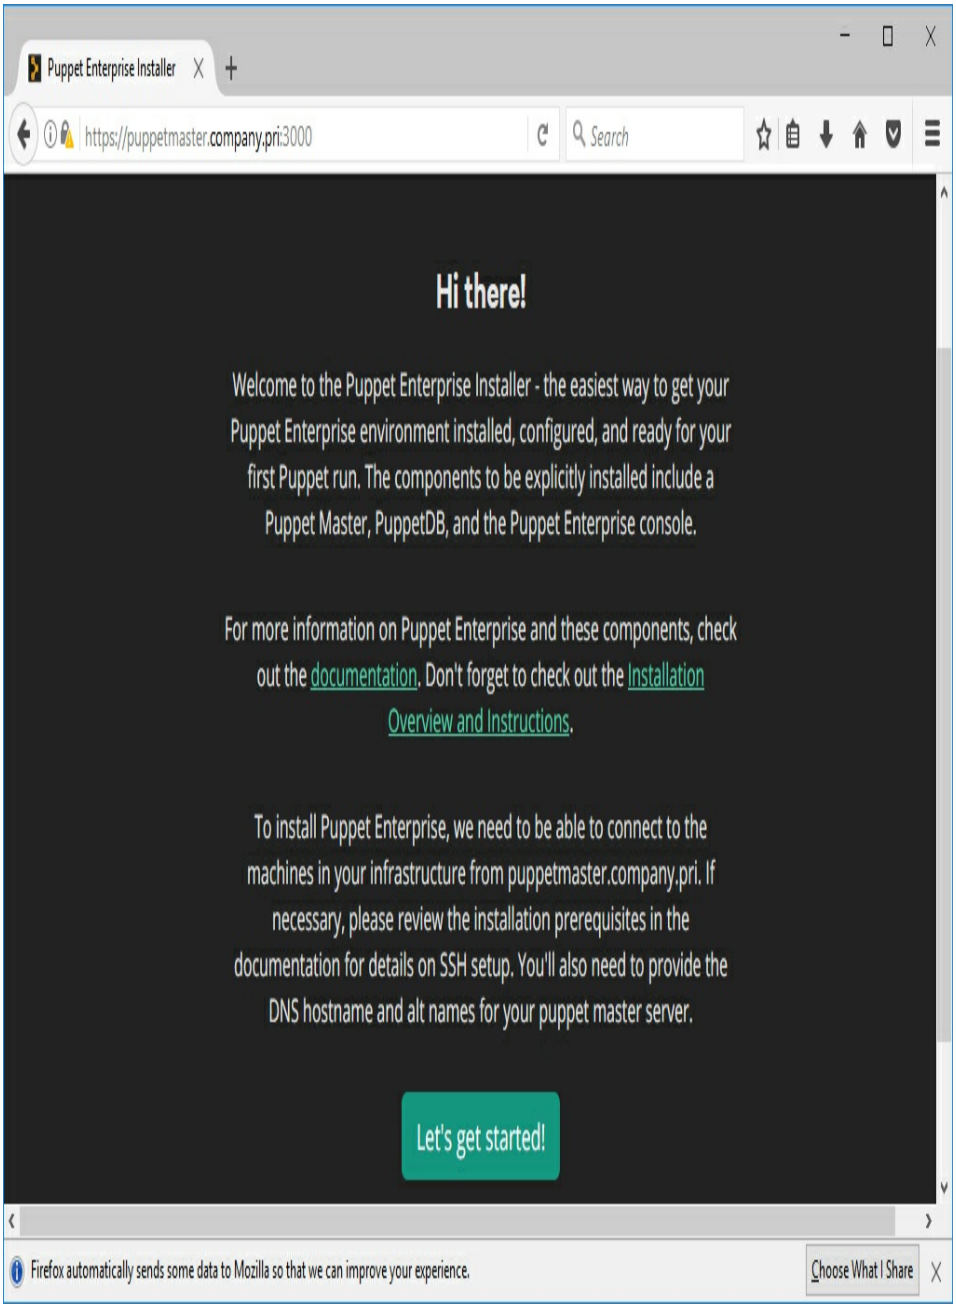


Figure 3-28 Puppet Server configuration

4. On the Puppet Master Component page, provide:
    1. The FQDN of the Puppet Master Server
    2. Any aliases that will be used by Puppet clients

3. Choose whether to install a PostgreSQL instance for the PuppetDB, and provide a console admin user password, as shown in Figure 3-29.

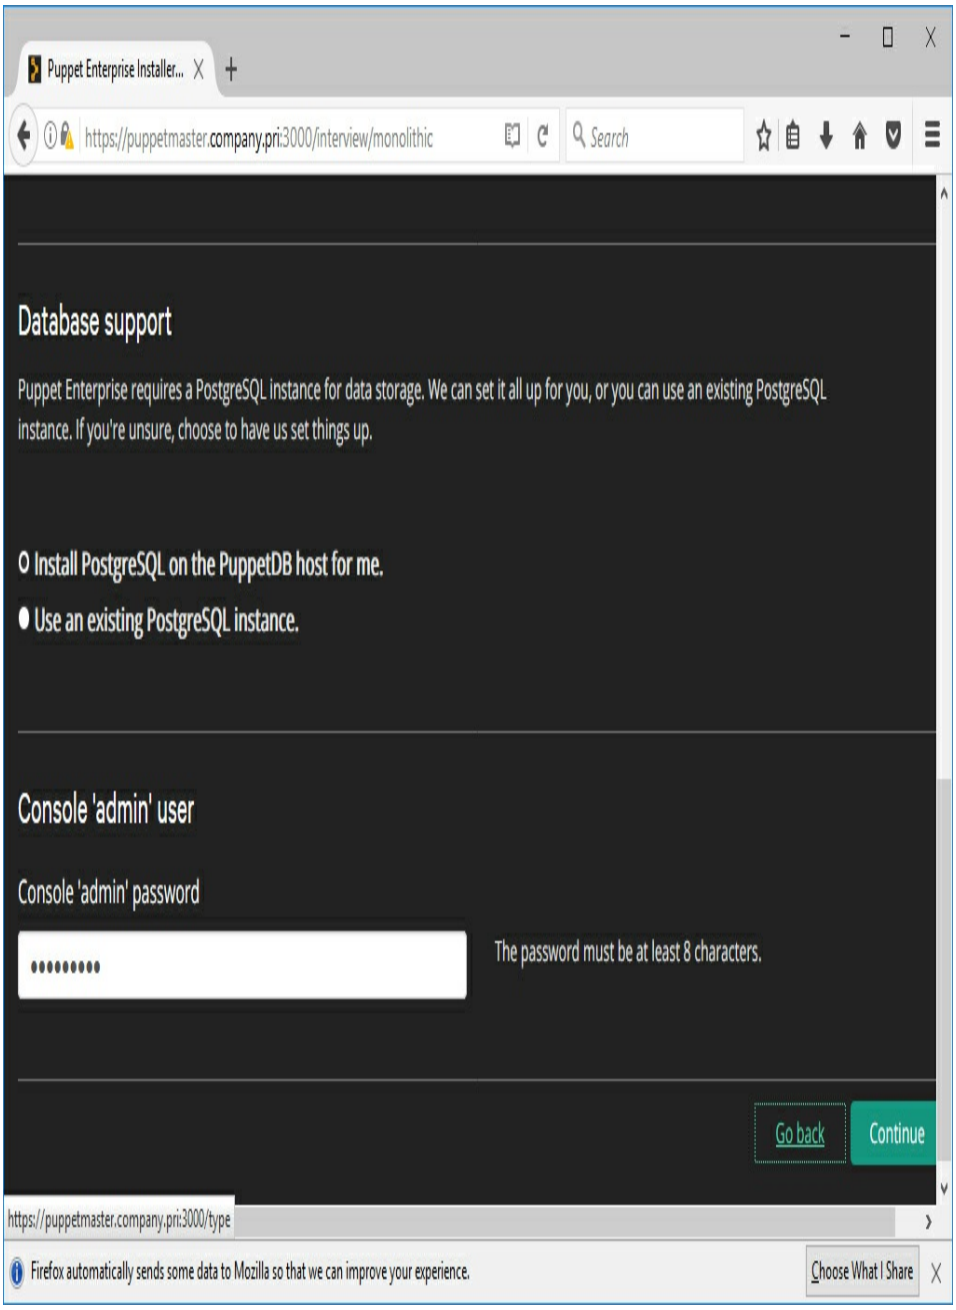


Figure 3-29 Install database and configure admin password

5. After the validation occurs, click Deploy Now. Puppet Server is deployed

on the Linux server.

You can now connect to the web-based administration interface using the username *admin* and the password provided in step 4. The Inventory node lists all computers that are currently being managed by the Puppet server.

> More Info
>
> **Deploying a Puppet Server**
>
> You can learn more about deploying Puppet Servers at https://docs.puppet.com/pe/latest/windows_installing.html.

Deploying Puppet agent to Windows Server

The first step in deploying a Puppet agent to a computer running Windows Server 2019 is to configure the Puppet Master Server so that it has the correct agent installation software for an x64 version of Windows. You can do this by taking the following steps:

1. In the Puppet administration console, click Classification under the Nodes menu, and in the PE Infrastructure section shown, click PE Master.
2. On the Classes tab, in the Add New Class text box, type **Windows**, and then click *pe_repo::platform::windows_x86_x64*, as shown in Figure 3-30. Click Add Class, and then click Commit Changes. Unless you run the Puppet agent *-t* command on the Puppet Master Server, you'll need to wait up to 30 minutes for the Puppet agent to wake up.

Technet24

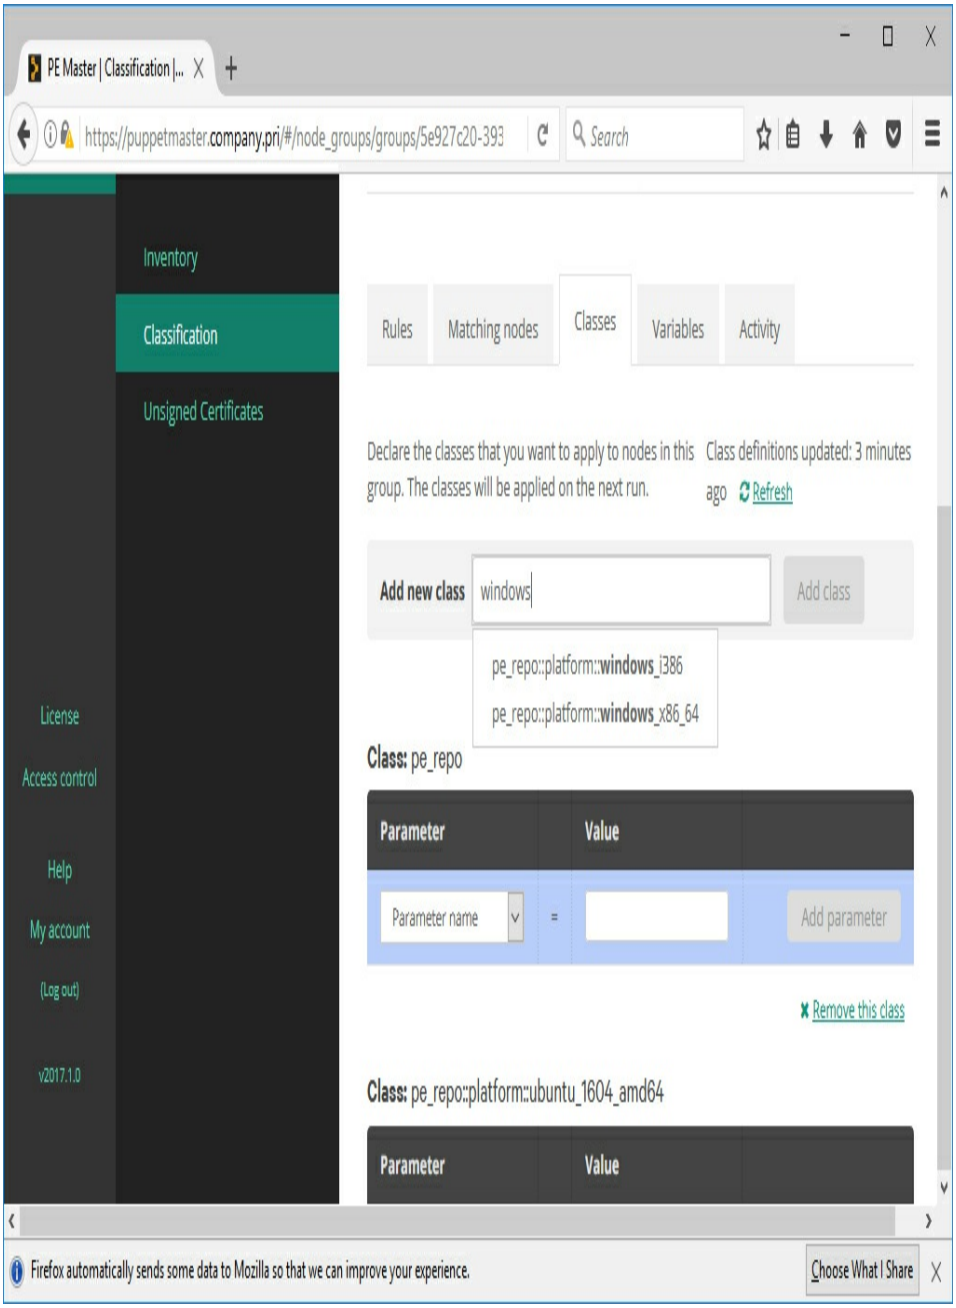


Figure 3-30 Configure Windows as a class

3. On your Windows Server, open a browser and navigate to *https://<YOUR PUPPET MASTER HOSTNAME>:8140/packages/current/windows-x86_64/puppet-agent-x64.msi*.

4. Run the downloaded installer. On the Welcome page, click Next, and then accept the terms of the License Agreement.
5. On the Configuration page, ensure that the FQDN of the Puppet Server matches the address of your Puppet Server, as shown in Figure 3-31, and click Install. Click Finish when the installation completes.

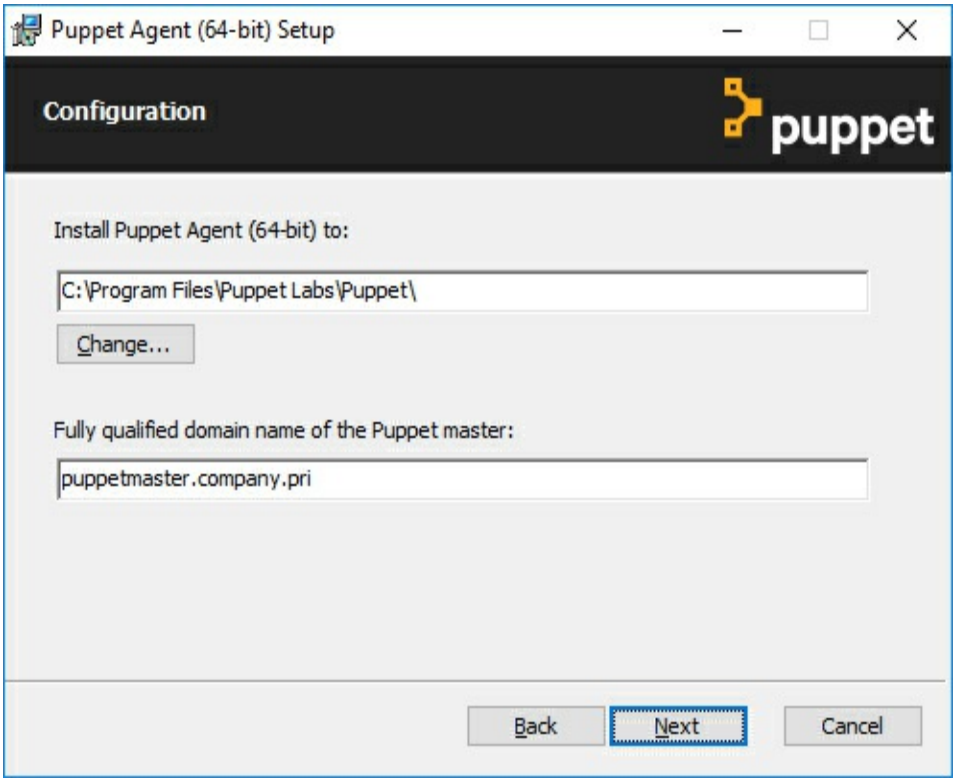


Figure 3-31 Puppet Agent Setup dialog box

6. Restart the computer, start an administrative PowerShell session, and run the following command, which initiates communication with the Puppet Master Server, allowing you to then trigger:

```
puppet agent -t
```

7. You'll then need to return to the Puppet console, navigate to the Unsigned Certificates tab under Nodes, and then sign the unsigned node of the newly detected certificate of your new Windows Server Puppet client, as shown in Figure 3-32.

Technet24

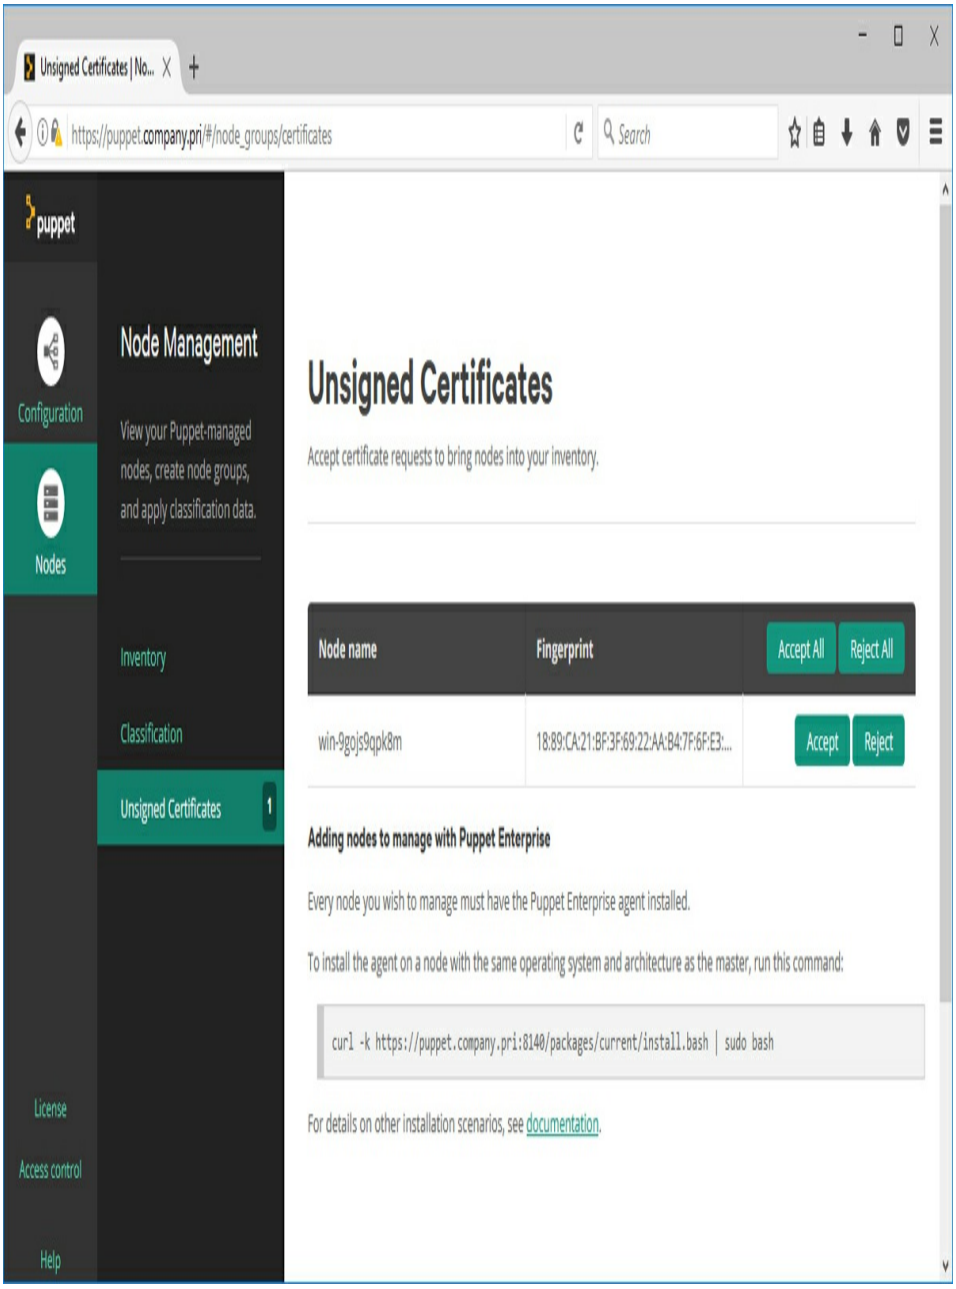


Figure 3-32 Sign unsigned certificates

Managing Windows Server configuration

Puppet manifest files describe a desired state of a managed server in terms of a list of resources and a set

of attributes. You can use any text editor to write a manifest, though you might want to use a tool like Visual Studio Code because it supports syntax highlighting for Puppet syntax. Puppet manifests use the *.pp* extension.

You can use the Puppet Master web-based console to specify which manifests apply to particular Puppet clients, also termed nodes. Once Puppet is installed on a client, you can run Puppet commands from an elevated PowerShell session. Generally, you'll build and test manifests on a development node using the command line tools. Then you can configure the manifests to be distributed and enforced by the Puppet Master Server.

You can perform the following tasks from an elevated PowerShell prompt on a computer that has the Puppet agent installed:

- You can validate the syntax on a manifest using the *Puppet parser validate* command. For example, run the following command to check the syntax of a manifest file named *services.pp*:

  click here to view code image

  ```
  Puppet parser validate
  c:\test_manifests\services.pp
  ```

- Manifests are applied using the *Puppet apply* command. You can simulate a Puppet cycle using the *--noop* option, which runs Puppet in non-enforcement mode. For example, to simulate a Puppet cycle against the *services.pp* manifest, run the command:

click here to view code image

```
Puppet apply c:\test_manifests\services.pp
--noop
```

- Once you've checked a manifest in non-enforcement mode, you can enforce the configuration using Puppet Apply.
- Puppet agent allows you to retrieve compiled catalogs from a Puppet Master. By default, this cycle runs every 30 minutes.

More Info

Basic Windows configuration

You can learn more about basic Windows Server configuration using Puppet at https://docs.puppet.com/pe/latest/windows_basic_tasks.html.

Puppet Windows Module Pack

Puppet Forge is an online community for modules that are submitted by people in the Puppet community, as well as people who work for the organization. Because an extensive community exists around Puppet, it's likely that someone has already written a module to perform a configuration task that you want to perform. You can take their hard work and use it to help configure the systems in your organization.

The Windows module pack is a collection of Puppet modules that you can use to perform common administrative tasks. To install the Windows Module Pack, open an administrative PowerShell prompt on a computer with the Puppet Agent installed, and run the

following command:

click here to view code image

```
Puppet module install puppetlabs/windows
```

The Puppet Windows Module pack includes tools that allow you to use Puppet to manage the following:

- Read, create, and write registry keys with the *registry* module.
- Remotely manage the execution of PowerShell scripts with the *powershell* module.
- Manage PowerShell DNS resources using the *dsc* module.
- Configure fine-grained access control permissions using the *acl* module.
- Manage WSUS client configuration using the *wsus_client* module.
- Configure which roles and features are installed using the *windowsfeature* module.
- Download files from various locations using the *download_file* module.
- Build and manage IIS sites and virtual applications using the *iis* module.
- Manage the deployment of software packages using the *chocolatey* module.
- Configure environment variables using the *windows_env* module.

While not comprehensive, you should now have a foundation-level understanding of the functionality of what you can do with Puppet, and whether it is worthy of further investigation in terms of inclusion as a part of your Windows Server administration toolkit.

More Info

Windows Module Pack

You can learn more about the Windows Module Pack for Puppet at https://docs.puppet.com/pe/latest/windows_modules.html.

Inside OUT

DSC, Puppet, or Chef?

There is no right answer to the question: "DSC, Puppet, or Chef?" If your organization's infrastructure is almost exclusively based on Microsoft products, you might want to consider DSC as a first option to see if it does what you need it to do. The intention of Microsoft is for DSC to be a fully supported way of configuring their products in the future, though it is in its early days. Puppet and Chef have longer pedigrees in the Open Source world, and if you already use one of these products, it makes sense to use it to manage Windows Server 2019. Microsoft certainly intends for DSC to work with Linux in the future, but it's also definitely playing catch up in this area compared to Puppet and Chef. What's important to understand, though, is that Microsoft is working with both of these companies, rather than against them; Microsoft's attitude is that if all of these technologies provide you with an excellent toolset to manage the configuration of Windows Server 2019, you'll be much more likely to choose Windows Server 2019 than if you could only do configuration management through Microsoft's in-house solution.

PACKAGE-MANAGEMENT UTILITIES

Configuration of servers and applications through tools like DSC, Puppet, and Chef provide an example of one way that the process of Windows server administration has changed in the last decade. Another way that things have changed is through the use of package-management utilities. Rather than the traditional process of downloading the current version of software from a vendor's website, vendors are increasingly

publishing their software in centralized repositories that you can access through package-management utilities. The benefit of package-management utilities is that they provide a central location for obtaining software, updating, and removing software applications.

If you've used a Debian-derived flavor of Linux, you'll be familiar with the *Apt-get* utility. *Apt-get* allows you to install any package prepared for Debian. It also allows you to update packages and the operating system. You can even use *Apt-get* to perform an entire operating system upgrade.

While not yet as advanced as *Apt-get*, package-management on Windows follows the paradigm that's been available in the Open Source world for several decades. Windows Server both directly and indirectly now supports package management, either directly through cmdlets in the *PackageManagement* Windows PowerShell module or indirectly through package providers such as Chocolatey.

Windows Server 2019 includes native support for the installation of a variety of different package-management providers through PowerShell. Some providers make PowerShell resources available, and others make pre-packaged software available that you would otherwise download from the Internet and manually install. You can view which package providers

are accessible by default to a Windows Server 2019 installation by using the *Find-PackageProvider* cmdlet, shown in Figure 3-33. In this chapter, you'll learn about the PowerShell Gallery and Chocolatey. In later chapters, you'll learn about specific providers, such as the *DockerMsftProvider*, which contains the tools used to manage Windows Server and Hyper-V containers; in Chapter 10, you'll learn about the ability to download containers that run full containerized isolated versions of Microsoft server applications including SQL Server 2016 community edition.

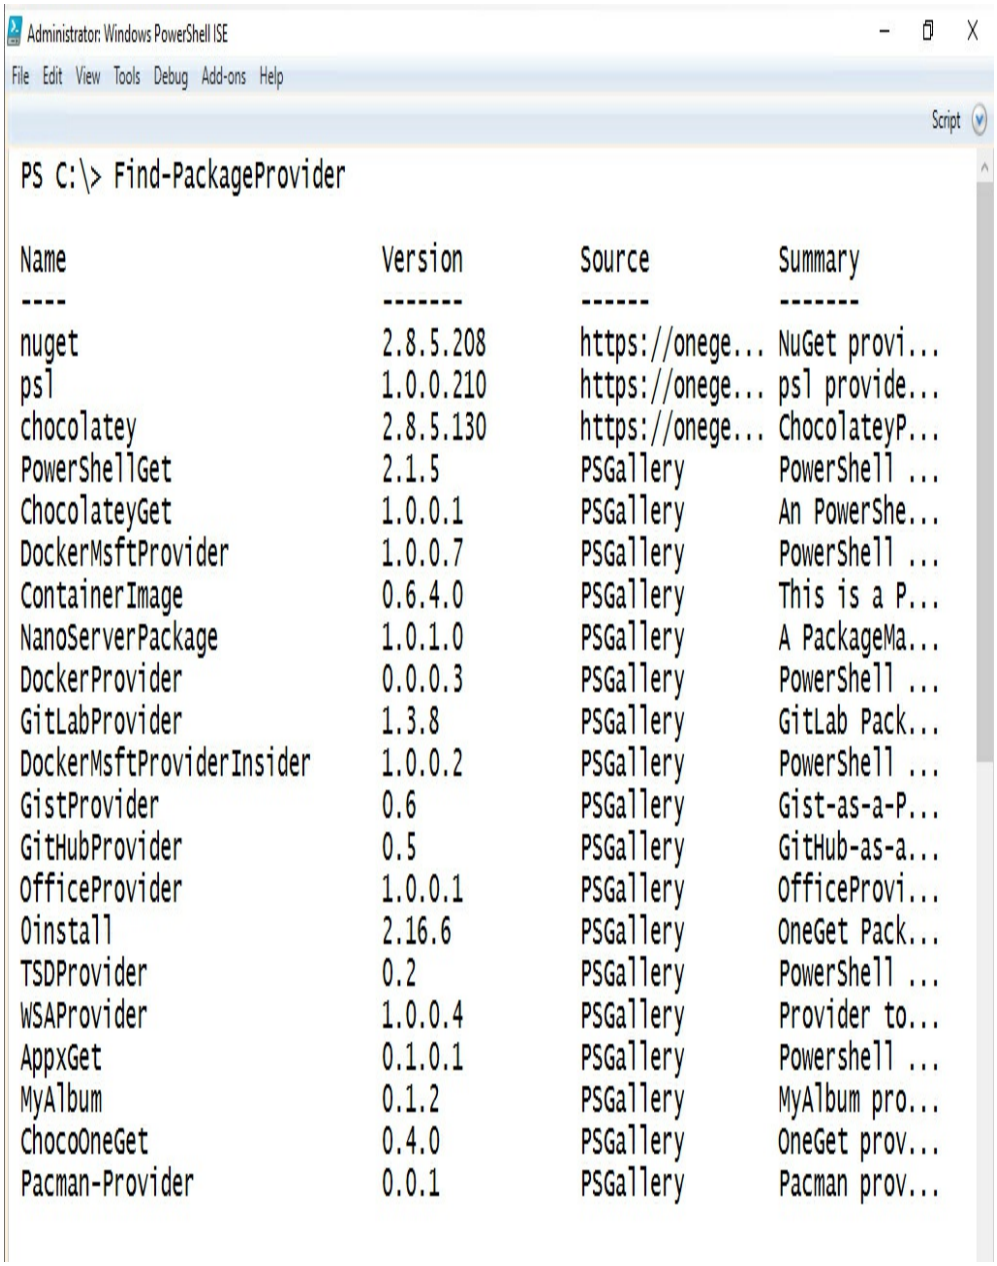


Figure 3-33 Package providers

PowerShell Gallery

PowerShell Gallery is a provider that functions as a clearinghouse for PowerShell-related tools and content.

The PowerShell team publishes content to the gallery, as do other PowerShell enthusiasts and MVPs. This means that it's fairly likely that if there is a script, cmdlet, or a module that does something you think quite a few people would need to do, someone has probably already created it and published it to the PowerShell gallery. You can interact with PowerShell Gallery with the *PowerShellGet* module. You can use this module to access cmdlets that do the following:

- Use *Find-Module* and *Find-Script* to find items in the gallery.
- Use *Save-Module* and *Save-Script* to save items from the gallery.
- Use *Install-Module* and *Install-Script* to install items from the gallery.
- Use *Publish-Module* and *Publish-Script* to publish items to the gallery.
- Add a custom repository to the gallery using *Register-PSRepository*.

Inside OUT

PowerShell community

Sites such as PowerShell.org are run and managed by the PowerShell community and contain a wealth of knowledge and advice on how you can best leverage PowerShell to accomplish your goals.

More Info

PowerShell Gallery

You can learn more about PowerShell Gallery at https://www.powershellgallery.com/.

Chocolatey

Chocolatey is perhaps the most well-known package provider available for Windows operating systems. It has a very extensive collection of software packages, with anything from commonly used applications such as 7-Zip and Visual Studio Code, to administration platforms such as Windows ADK. Chocolatey can also be integrated with Puppet, Chef, and DSC.

Chocolatey also runs very well on the Server Core installation option of Windows Server 2019, as well as within Windows Server and Hyper-V containers. It can be substantially quicker to use Chocolatey to retrieve and install a package; for example, it is quicker to use Chocolatey to install the Windows ADK with all options in the default location on a computer running the Server with Desktop Experience option. To do so, run the following command:

click here to view code image

```
Choco install -y windows-adk-all
```

This requires substantially less time than searching the Microsoft website, downloading all the necessary prerequisites, and running the installation routine.

To install Chocolatey, you'll need to set the PowerShell execution policy to *unrestricted* using the following command:

click here to view code image

```
Set-ExecutionPolicy Unrestricted
```

Then you'll need to run the following PowerShell command:

click here to view code image

```
iwr https://chocolatey.org/install.ps1 -UseBasicParsing | iex
```

This installs Chocolatey and all its dependencies on the local computer.

> **Inside OUT**
>
> **Install packages at your own risk**
>
> In a secure environment, you would not set your execution policy to unrestricted, and you certainly wouldn't download and install packages from a public repository no matter how reputable it may seem. In the long run, you don't know anything about the people building the packages, and while they are most likely simply helpful people trying to make your life easier by packaging software, they could also be a part of the great galactic space lizard malware conspiracy and simply finding a new vector to get malware deployed on your server. Before using a package in production, make sure you test it. When it comes to building secure servers, you'll need to validate the integrity of all software that you deploy, which means that the shortcuts provided by package providers are probably not an option.

Chocolatey has the following advantages when it comes to software deployment on servers:

- Makes the deployment of software extremely simple. Once Chocolatey is installed, further packages can be installed with one line of code.
- Packages are downloaded and installed with one command. You don't need to download a file, unzip it, run the installer, and answer the questions of an inquisitive wizard.

- Package dependencies are automatically installed when you attempt to install a package.
- Maintainers often regularly update packages. This means that you can update all software installed via Chocolatey on a system by running the following command:

```
Chocolatey update all
```

Chocolatey has the following disadvantages:

- Usually requires that the computer you are deploying software on be connected to the Internet. You can configure internal repositories, but you'll have to manage and maintain them yourself, including validating the quality of the software found in those repositories.
- Packages are usually installed in default locations with default settings. You don't get the installation customization options that you get when you install software manually.
- The provenance of the software isn't clear. While it's reasonable to assume that the software is safe, even large vendor app stores have apps that are infected with malware.

Using a simple script, you can install Chocolatey and then use it to download and install a common set of packages. You can also set up your own Chocolatey repository of packages that hosts only packages that you trust. You can then point a local Chocolatey instance at that repository and use it to quickly install software on newly deployed computers running a Windows operating system.

Inside OUT

Great for clients, too

Chocolatey is really great for clients and has far more packages designed for computers running the Windows client operating system than it does for servers. It's far simpler to run two lines of code to install Chocolatey, and then download a utility such as 7-Zip than it is to open a browser, navigate to the appropriate website, and then locate, download, and install the appropriate file.

More Info

Chocolatey

You can learn more about Chocolatey at https://chocolatey.org/. You can learn more about integrating Chocolatey with Puppet, Chef, and DSC at https://chocolatey.org/docs/features-infrastructure-automation.

Chapter 4

Active Directory

Managing Active Directory

Domain controllers

AD DS structure

Accounts

Group policy

Restoring deleted items

Managing AD DS with PowerShell

Active Directory, the identity glue that binds on-premises Microsoft networks, is at the center of almost all on-premises networks. Although each computer can have its own unique individual user and service accounts, Active Directory provides a central user, a computer, and a service account store.

But Active Directory is more than an identity store; it can also be used to store data for Active Directory-aware

Technet24

applications. One example of this is Microsoft Exchange Server, which stores server configuration information in Active Directory. Other applications, such as System Center Configuration Manager, are also highly dependent on Active Directory.

Active Directory is a topic so vast that completely covering it would require a multi-volume set. In this chapter, we look at managing Active Directory. We also cover deploying domain controllers, forests, domains, service accounts, and Group Policy, as well as examining how to restore deleted items.

MANAGING ACTIVE DIRECTORY

Domain Controllers are one of, if not the, highest value target for attackers on your network. An attacker that is able to take control of a domain controller has control over every domain joined computer in the organization. Also known as "domain dominance," having control of a domain controller means having control of all the authentication and authorization processes in the domain.

While we will discuss server hardening in more detail in Chapter 19, "Hardening Windows Server and Active Directory," one strategy for increasing server security is to reduce the attack surface. As mentioned in Chapter 2, where possible, you should deploy roles on computers

installed in the Server Core configuration because this configuration has a smaller attack surface than a server that includes desktop components. Server Core is a great option for domain controller deployment because it provides a reduced attack surface, which strengthens the security of the installation.

Remote rather than local administration

You should perform Active Directory–management tasks remotely using Windows Admin Center, management consoles, or PowerShell rather than signing on to the domain controller directly using RDP. If you're doing all your administrative tasks remotely using Windows Admin Center, management consoles, or PowerShell, it won't make any difference to you that you've deployed the domain controller in the Server Core configuration. Using remote consoles also reduces the chance of malware being introduced to the domain controller. There are countless stories of organizations having their security compromised because an administrator signed in to a server using Remote Desktop, went to download a utility from the Internet using the built-in web browser, and ended up with more than they bargained for in terms of malware because they weren't careful about their browsing destinations.

A general rule about consoles is that if you need to perform a task a couple of times, you should use

Windows Admin Center or a console. If you are new to a task, you're less likely to mess it up if you use a GUI tool. This is because a GUI tool holds your hand and assists you through the task. If you must repeatedly perform the same task, and it's a task that you are familiar with, you should automate it using tools such as PowerShell. You'll learn about the Active Directory related PowerShell cmdlets later in this chapter.

There are a number of consoles that you can use to perform Active Directory administrative tasks. These include the following:

- Active Directory Administrative Center
- Active Directory Users and Computers
- Active Directory Sites and Services
- Active Directory Domains and Trusts

At the time of writing, Windows Admin Center provides some Active Directory administrative functionality but not nearly enough that AD administrators could use it as their primary tool for managing Active Directory.

Active Directory Administrative Center

Active Directory Administrative Center (ADAC) was introduced with Windows Server 2012, but it never really caught on as the primary method of managing Active Directory for most administrators. Active Directory Administrative Center allows you to manage

users, computers, and service accounts to perform tasks with the Active Directory, such as Recycle Bin, and to manage functionality, such as Dynamic Access Control.

Active Directory Administrative Center is a newer console that has better search functionality than the other consoles listed in this chapter, which haven't substantively changed since the release of Windows 2000. Hopefully, Windows Admin Center will improve to the point that by the next edition of this book, it can be used as the primary GUI admin tool for Active Directory administration.

You can use Active Directory Administrative Center to manage the following:

- User, computer, and service accounts
- Domain and Forest functional level
- Fine-grained password policies
- Active Directory Recycle Bin GUI
- Authentication Policies
- Dynamic Access Control

ADAC is built on PowerShell, meaning that it provides a graphical interface to build and enact PowerShell cmdlets. You can use the PowerShell History viewer, shown in Figure 4-1, to see which cmdlets were used to carry out a task that you configured in the GUI. This simplifies the process of automating tasks because you

Technet24

can copy code straight out of the PowerShell history and then paste it into tools such as PowerShell ISE or Visual Studio Code.

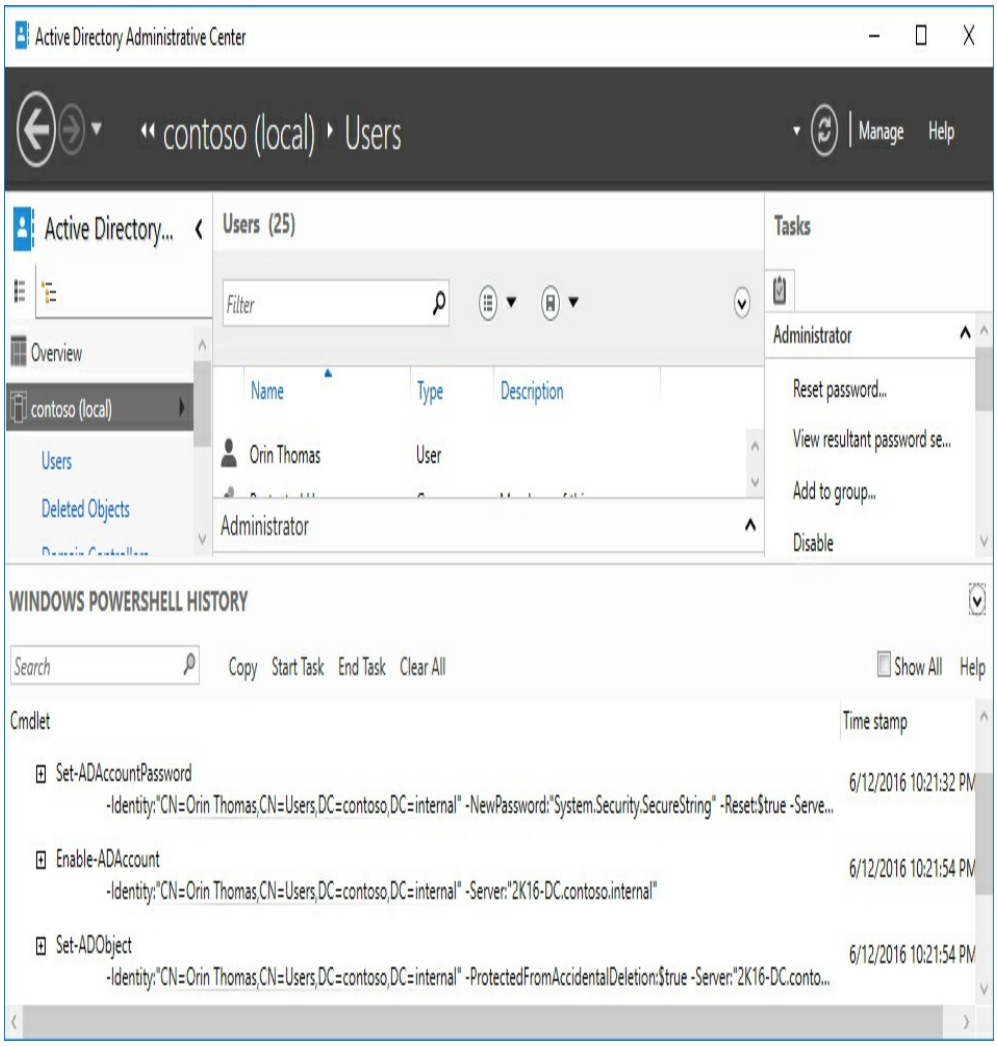


Figure 4-1 Active Directory Administrative Center

One of the most useful elements of Active Directory Administrative Center is the search functionality. You can use this functionality to locate accounts that might

require further attention, such as users who haven't signed in for a certain period of time, users configured with passwords that never expire, or users with locked accounts. Using the Add Criteria option in the Global Search node of the Active Directory Administrative Center, you can search based on the following criteria:

- Users with disabled/enabled accounts
- Users with expired passwords
- Users whose passwords have an expiration date or no expiration date
- Users with enabled but locked accounts
- Users with enabled accounts who have not signed in for more than a given number of days
- Users with a password expiring in a given number of days
- Computers running as a given domain controller type
- Last modified between given dates
- Object type is *user/inetOrgPerson/computer/group/organizational unit*
- Directly applied password settings for a specific user
- Directly applied password settings for a specific global security group
- Resultant password settings for a specific user
- Objects with a given last known parent
- Resource property lists containing a given resource property
- Name
- Description

- City
- Department
- Employee ID
- First name
- Job Title
- Last Name
- *SamAccountName*
- State/Province
- Telephone number
- UPN
- Zip/Postal Code
- Phonetic company name
- Phonetic department
- Phonetic display name
- Phonetic first name
- Phonetic last name

There are several tasks that you can't do with Active Directory Administrative Center or with PowerShell, such as running the delegation of Control Wizard. Generally speaking, when managing Active Directory using a console on Windows Server 2019, you should use Active Directory Administrative Center as your first option and fall back to Active Directory Users and Computers if you can't accomplish what you need to in ADAC.

Active Directory Users and Computers console

The Active Directory Users and Computers console is the one that many system administrators use to perform basic AD-related tasks. They use this console primarily out of habit because almost all functionality present in this console is also present in Active Directory Administrative Center. Active Directory Users and Computers has been around since the days of Windows 2000 Server.

Active Directory Users and Computers allows you to perform tasks including:

- Running the Delegation Of Control Wizard
- Administering different domains within the forest
- Selecting which domain controller or LDAP port the tool connects to
- Finding objects within the domain
- Raising the domain functional level
- Managing the RID Master, PDC Emulator, and Infrastructure Master FSMO role locations
- Creating and editing the properties of
    - Computer accounts
    - User accounts
    - Contacts
    - Groups
    - *InetOrgPerson*

    - *msDS-ShadowPrincipalContainers*
    - *msImaging-PSPs*
    - MSMQ Queue Alias
    - Organizational Unit
    - Printers
    - Shared Folder
- Resultant Set of Policy planning

The View Advanced Features function allows you to see more details of the Active Directory environment. You enable this from the View menu of Active Directory Users And Computers. Enabling this view allows you to see containers that aren't visible in the standard view. If you've ever read a set of instructions that tell you to locate a specific object using Active Directory Users And Computers, and you haven't been able to find that object, chances are that you haven't enabled the View Advanced Features option.

The Delegation Of Control Wizard is only available in Active Directory Users And Computers. This wizard allows you to delegate control over the domain and organizational units (OU). For example, you use this wizard to delegate the ability for a specific group to reset user passwords in an OU. This wizard is useful when you want to delegate some privileges to a group of IT ops staff, but you don't want to grant them all the privileges that they'd inherit if you made them a

member of the Domain Admins group. You'll learn more about delegating permissions in Chapter 19, "Hardening Windows Server and Active Directory."

Active Directory Sites and Services console

You use the Active Directory Sites And Services console to manage Active Directory sites, which indirectly allows you to control a number of things, including replication traffic and which server a client connects to when using products such as Exchange Server. Sites are configured for the forest, with each domain in the forest sharing the same set of sites.

An Active Directory site is a collection of TCP/IP subnets. Sites allow you to define geographic locations for Active Directory on the basis of TCP/IP subnets. You can have multiple TCP/IP subnets in a site. You should put subnets together in a site where the hosts in that site have a high-bandwidth connection to each other. Usually, this means being in the same building, but it could also be multiple buildings with very-low-latency gigabit links between them.

For example, imagine that your organization has its head office in Melbourne and a branch office in Sydney. You can set up two sites: one site for Melbourne and the other for Sydney. This ensures that computers in the Melbourne location interact as much as possible with resources located in Melbourne, and computers in the

Technet24

Sydney location interact as much as possible with resources located in Sydney.

You associate the TCP/IP subnets in the head office with the Melbourne site and the TCP/IP subnets in the branch office with the Sydney site. After you do this, functionality such as replication topology is automatically configured.

The default name for the site you installed the first domain controller in is something like *Default-First-Site-Name*. You should change this as appropriate prior to adding subnets. When you add a subnet, you need to specify its associated site, so add sites first and then add subnets after your sites are present.

To add a new Active Directory Site, right-click the Sites node in the Active Directory Sites And Services console and click New Site. Specify the site name and choose Select A Site Link object. A site link object represents a connection between two sites. The default site link object is named *DEFAULTIPSITELINK*. You can change the site link object later. Figure 4-2 shows the creation of a site named *Sydney*.

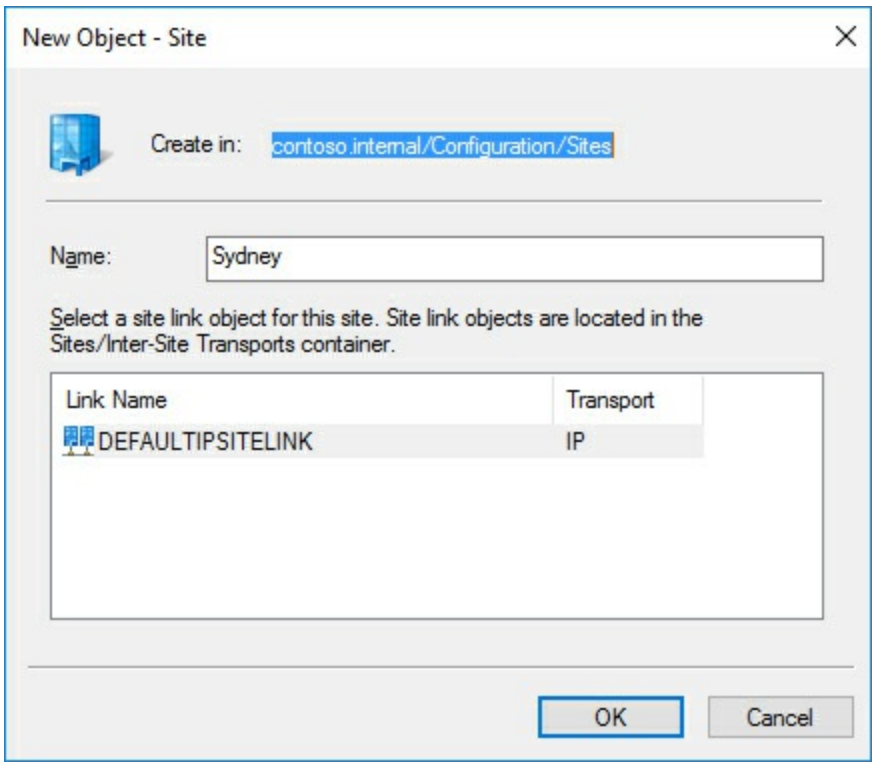


Figure 4-2 Create a new site.

To add a subnet, right-click the Subnets node in Active Directory Sites And Services and then click New Subnet. You can specify the new subnet in IPv4 or IPv6 format. After you've specified the subnet, you need to specify which site the subnet is associated with. Figure 4-3 shows the *10.10.10.0/24* subnet associated with the *Melbourne* site.

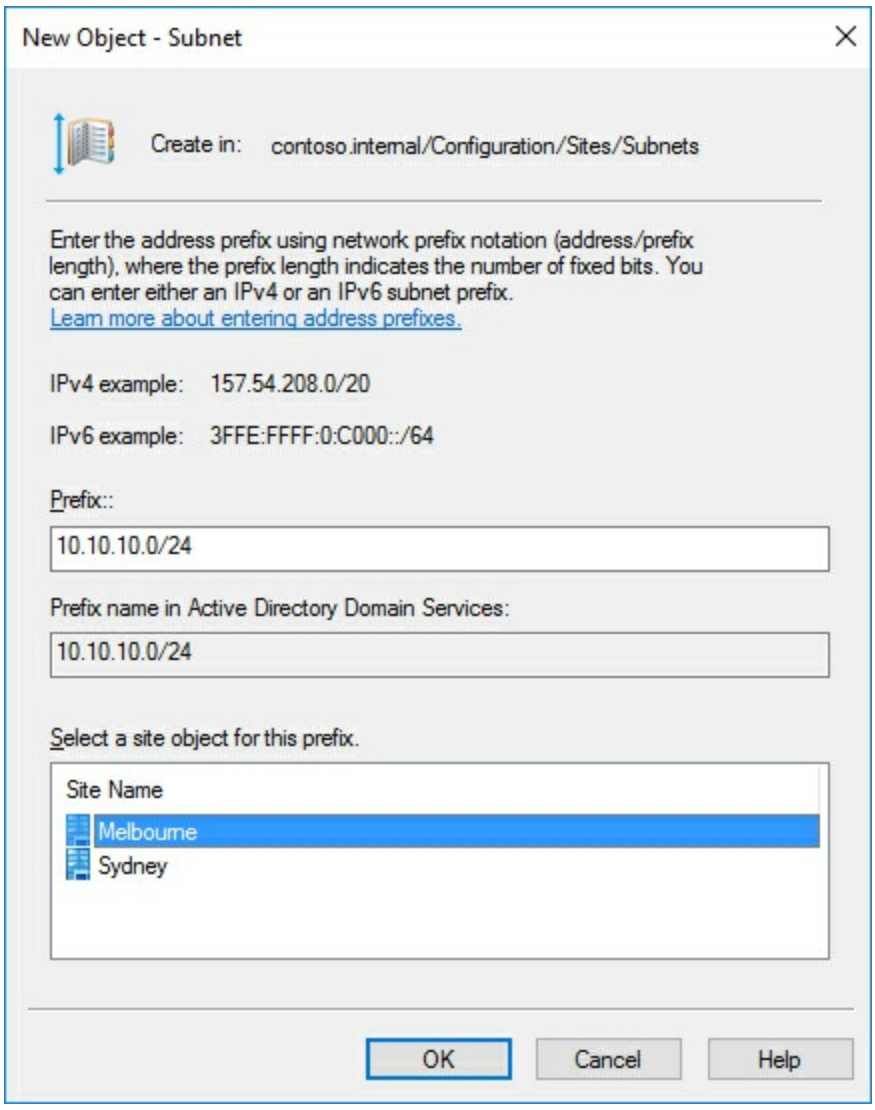


Figure 4-3 New subnet

Site links represent connections between sites. You can create IP or SMTP site links. You almost always use IP site links because these are authenticated and encrypted. You use SMTP site links when the connection between sites is unreliable. When configuring a site link, you can configure the replication

interval, with the default being once every 180 minutes. You can also configure the hours in which replication can occur, the default being 24 hours a day, 7 days a week. You create a site link by right-clicking the IP node under Inter-Site-Transports in the Active Directory Sites And Services console and by clicking New Site Link. And you specify two or more sites that are connected by the link. After the link is created, you can modify the replication schedule, cost, and the hours replication can occur in. Site links with lower replication costs are preferred over site links with higher replication costs.

A site link bridge is a collection of site links, but you only need to use them when your IP network is not fully routed. For example, you might create a site link bridge to connect the Sydney-Melbourne and Melbourne-Perth site links as a way of ensuring that replication traffic between Sydney and Perth passes through the Melbourne site.

After you've provided information about sites, subnets, site links, and site-link bridges, an Active Directory component called the Knowledge Consistency Checker automatically creates a replication topology using a least-cost spanning tree design.

Domain controllers automatically associate with the site that they are deployed in. If you didn't set up sites and subnets prior to deploying domain controllers in

separate locations, right-click the domain controller that you want to move within Active Directory Sites And Services, click Move, and then specify the destination site.

Active Directory Domains and Trusts console

You use the Active Directory Domains And Trusts console to configure and manage trust relationships. By default, all domains in a forest trust each other. Your primary use of this console is to create trust relationships between:

- Domains in separate forests
- Separate forests
- Kerberos V5 realms

When creating a trust, you can choose between the following types:

- **One-Way: Incoming.** In this trust relationship, your local domain or forest is trusted by a remote domain or forest.
- **One-Way: Outgoing.** In this trust relationship, your local domain or forest trusts a remote domain or forest.
- **Two-way.** In this trust relationship, your local domain or forest trusts, and is trusted by, a remote domain or forest.

When configuring a trust, you can determine whether you want to configure selective authentication. By default, the trust works for all users in the source or destination forest or domain. When you configure selective authentication, you can limit which security

principals are allowed access and which resources they can access. You do this by configuring those users with the Allowed To Authenticate permission on each resource computer in the trusting domain or forest.

DOMAIN CONTROLLERS

Domain controllers (DCs) are the heart of Active Directory. Their primary role is to host the Active Directory database, stored in the *ntds.dit* file. Given their importance, you should limit who can sign on locally to a domain controller.

> **Inside OUT**
>
> **Make sure you double-up!**
>
> Ensure that you have at least two domain controllers in each domain, and make sure you make regular backups. If one domain controller fails, the other one is still there to perform all necessary functionality. Imagine you only have one domain controller and you don't have a backup. If that DC fails, you'll need to rebuild your domain from scratch.

Deployment

Although the wizard for deploying a domain controller has been updated from what it looked like back in Windows Server 2008 R2, the process is still functionally the same. You choose whether you want to:

- Add a domain controller to an existing domain
- Add a new domain to an existing forest
- Add a new forest

If you are adding a domain controller to an existing domain, have the computer domain joined before promoting it. When you add a new domain to an existing forest, choose between adding a child domain or a tree domain. When you add a tree domain, you create a new namespace within an existing forest. For example, you can create the *Adatum.com* domain in the existing *Contoso.com* forest. *Contoso.com* remains the root domain of the forest.

Microsoft's recommendation is that you use a registered root domain name, such as *Contoso.com* for the domain name, rather than a non-externally resolvable domain name like *contoso.internal*. Having a registered, externally resolvable domain name simplifies the process when configuring synchronization with Azure AD Connect.

Windows Server 2019 supports split DNS, which means that you can configure zones so that a subset of records is resolvable for clients on external networks. Many organizations still use non-resolvable domain names, and you should only take the Microsoft advice into account if you are deploying a new forest or you are reconfiguring a domain in preparation for synchronizing with Azure Active Directory.

When you add a new child domain to an existing forest, you specify the parent domain. For example, you could

add the *australia.contoso.com* child domain to the *contoso.com* domain. After you've done that, you can add the *victoria.australia.contoso.com* child domain to the *australia.contoso.com* parent domain. Adding a child domain requires Enterprise Administrator credentials in the forest.

Domain controllers can also host the DNS service and function as global catalog servers. In Windows Server environments, DNS is almost always hosted on domain controllers. You'll learn more about DNS configuration and options in Chapter 5 "DNS, DHCP, and IPAM." Global Catalog servers are covered later in this chapter.

When deploying a domain controller in an environment with multiple sites configured, you can select which site you want the domain controller to belong to. You can change this after deployment using the Active Directory Sites And Services console.

Directory Services Restore Mode passwords

Directory Services Restore Mode (DSRM) allows you to perform an authoritative restore of deleted objects from the Active Directory database. You must perform an authoritative restore of deleted items because if you don't, the restored item is deleted the next time the AD database synchronizes with other domain controllers where the item is marked as deleted. Authoritative restores are covered later in this chapter. You configure

Technet24

the Directory Services Restore Mode password on the Domain Controller Options page of the Active Directory Domain Services Configuration Wizard, as shown in Figure 4-4. Note that even though a computer running Windows Server 2019 is being configured as a domain controller, the maximum Forest and Domain functional level is Windows Server 2016. This is because there is no Windows Server 2019 domain or forest functional level.

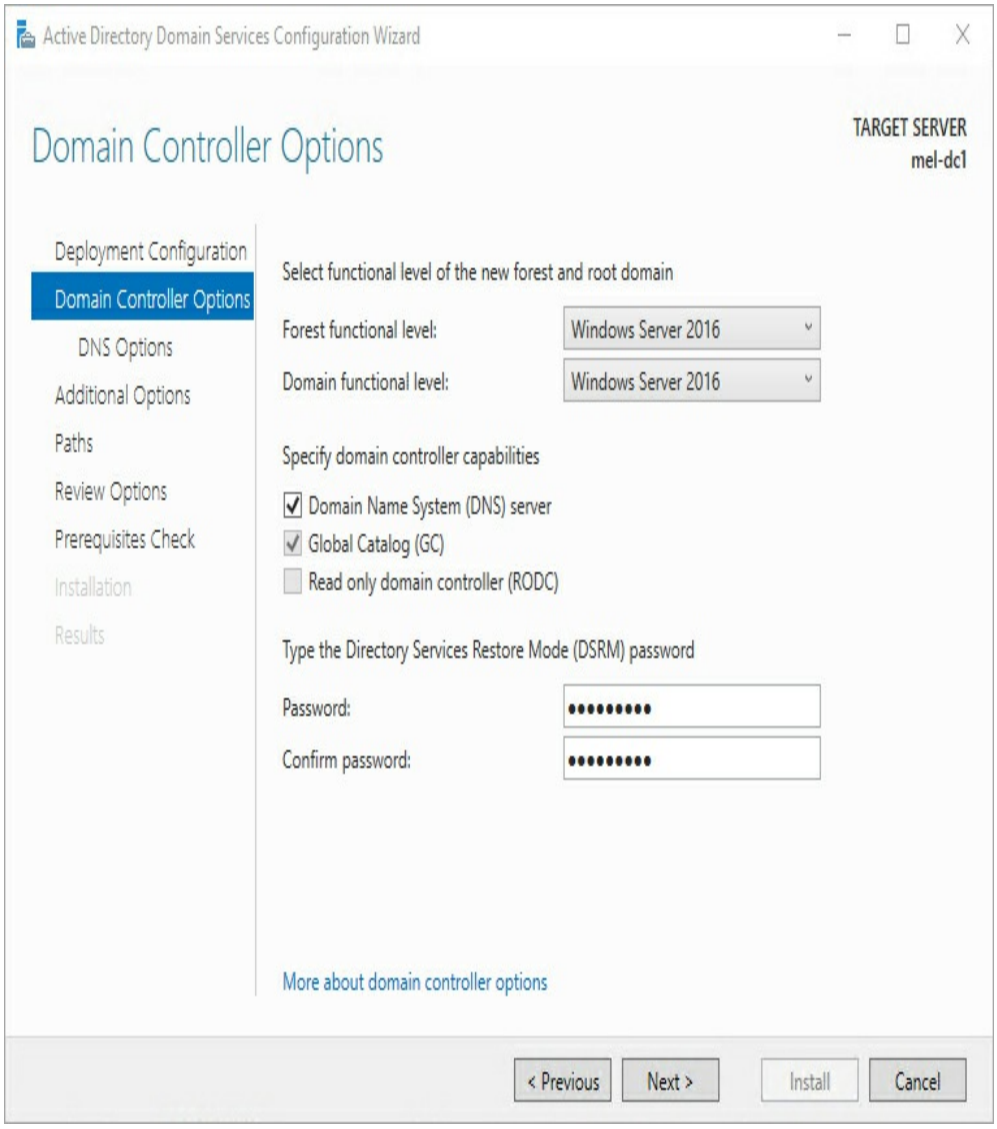


Figure 4-4 Configuring Domain Controller Options

In the event that you forget the DSRM password, which, in theory, should be unique for each domain controller in your organization, you can reset it by running *ntdsutil.exe* from an elevated command prompt and entering the following command at the ntdsutil.exe prompt, at which point you are prompted to enter a new

DSRM password:

```
Reset password on server null
```

Advanced installation options

One of the advanced installation options for domain controllers is to install from media. Installing from media gives you the option of pre-populating the AD DS database for a new DC from a backup of an existing DC's AD DS database, rather than having that database populated through synchronization from other domain controllers in your organization. This is very useful when you need to deploy a DC at a remote location that has limited WAN connectivity and you don't want to flood the WAN link with AD DS database synchronization traffic during domain controller deployment. Instead, you ship a backup of the AD DS database to the remote site and that backup is used to perform initial AD DS database population. After the newly installed DC connects to other domain controllers, it performs a synchronization, bringing the database up to date with a much smaller synchronization than what would be required when synchronizing from scratch.

Server Core

As the reduced attack surface area of Server Core deployments makes it more secure, domain controllers

—which don't have a GUI dependency—are one of several perfect workloads for Server Core deployments. Microsoft recommends deploying domain controllers using the Server Core deployment option and managing those servers remotely.

To configure a computer running Server Core as a domain controller, you can:

1. Remotely connect to the server using the Server Manager console.
2. Run the Add Roles And Features Wizard to remotely install the Active Directory binaries on the server.
3. Run the Active Directory Domain Services Configuration Wizard to promote the computer to a domain controller.

You can also use Windows Admin Center to deploy the AD DS feature, but at the time this book was written, you couldn't configure the feature once it was deployed. As an alternative to using the wizard or console, you can run the following PowerShell commands, either locally or remotely, to install the AD DS binaries and promote the server to domain controller:

Click here to view code image

```
Install-WindowsFeature AD-Domain-Services -IncludeManagementTools
Install-ADDSDomainController -DomainName contoso.internal -
InstallDNS:$True
-credential (Get-Credential)
```

In addition to running this set of commands, you also need to specify a Directory Service Restore Mode password before the computer running Server Core

completes the domain controller promotion process. You'll need to use the *Install-ADDSForest* cmdlet if you are installing the first domain in a new forest, as shown here:

Click here to view code image

```
Install-ADDSForest -DomainName contoso.internal -InstallDNS
```

Virtualized domain controllers

Domain controllers can be run on supported virtualization platforms, including the latest version of VMware and Hyper-V. With the Production Checkpoints feature available in Windows Server 2019 Hyper-V, domain controllers can be restored from a checkpoint without causing problems. Microsoft recommends that you run virtualized domain controllers as shielded VMs on a guarded virtualization fabric because this will minimize the chance that a nefarious or compromised virtualization administrator account could be used to access the contents of the DC VM. You can learn more about production checkpoints, shielded VMs, and guarded fabrics in Chapter 6, "Hyper-V."

Global catalog servers

Global catalog servers host a full copy of all objects stored in its host directory and a partial, read-only copy of all other objects in other domains in the same forest. They are used when it's necessary to perform a check of

other objects in the forest, such as when a check is performed of a universal group's membership, which could contain members from other domains in the forest.

You can use the Active Directory Sites And Services console to configure a server to function as a global catalog server by right-clicking the NTDS settings of the server. Alternatively, you can run the *Set-ADObject* cmdlet. For example, to configure the DC MEL-DC1 in the *Melbourne-Site* site as a global catalog server, run the following command:

Click here to view code image

```
Set-ADObject "CN=NTDS Settings,CN=MEL-
DC1,CN=Servers,CN=Melbourne-Site,CN=Sites,
CN=Configuration,DC=Contoso,DC=Internal" -Replace @{options='1'}
```

Read only domain controllers

Read Only Domain Controllers (RODC) are a special type of domain controller. They host a read-only copy of the AD DS database. Rather than storing the passwords for all users in the domain, RODCs only store the passwords of a specific set of accounts that you configure. The first domain controller in a new domain or forest cannot be an RODC.

The justification for RODCs is that DCs sometimes need to be located in places where servers have poor physical security and might be stolen. For example, many

organizations had branch offices where servers were located under someone's desk. A good rule of thumb is that you should consider a location insecure if it is accessible to anyone other than IT. If a janitor can pull out a computer's power cord to power a vacuum cleaner, the computer isn't in a secure location.

If a server that hosts a domain controller is stolen, the best practice is to reset the account passwords that might have been compromised because it's possible, with the correct tools, to extract passwords from the AD DS database if you have direct access to it. If an ordinary DC is stolen, you would, in theory, need to reset the passwords of every account in the domain because you could never be sure that someone hadn't gained access to the Active Directory database and found a way to extract the passwords of people in your organization.

With shielded VMs and shielded fabrics, it's possible to run a DC in a manner where the VM itself is protected by encryption. In the event that the host server is stolen, the AD DS database cannot be recovered because the contents of the virtualization server's storage is encrypted using BitLocker.

Inside OUT

RODC and Writeable DC Deployment

Concerns about the physical security of a DC are the primary reason to deploy an RODC, so it is extremely unlikely that you would have both an RODC and a writable DC at the same site. RODCs are for sites where the domain controller once was

placed in a location that wasn't secure. Though, if you do have concerns about the security of a location, it's probably not a great idea to deploy a domain controller at that location!

RODC password replication

One of the most important steps in configuring an RODC is limiting which passwords can be replicated down to the server from a writable domain controller. The default configuration of an RODC has it store almost everything from Active Directory except for user and computer account passwords. In the event that a user or computer account needs to be authenticated against Active Directory, the RODC acts as a proxy for a writable Windows Server DC. The authentication occurs but depends on the WAN link to be functional because if you could host a writable DC locally, you wouldn't need the RODC.

Although you can configure an RODC to not cache any passwords locally, you can configure an RODC to cache the passwords of select staff working at a branch office to speed their login. Caching passwords also allows branch office users to log in if the WAN link fails. If the WAN link fails and the user's credentials are not cached, the user is simply unable to log in to the domain.

You configure which accounts can authenticate using the RODC by using the Password Replication Policy, as shown in Figure 4-5. By default, only members of the

Allowed RODC Password Replication group can use the RODC to authenticate. This is only the case if the user account is not a member of the Account Operators, Administrators, Backup Operators, Server Operators, or Denied RODC Password Replication Group groups.

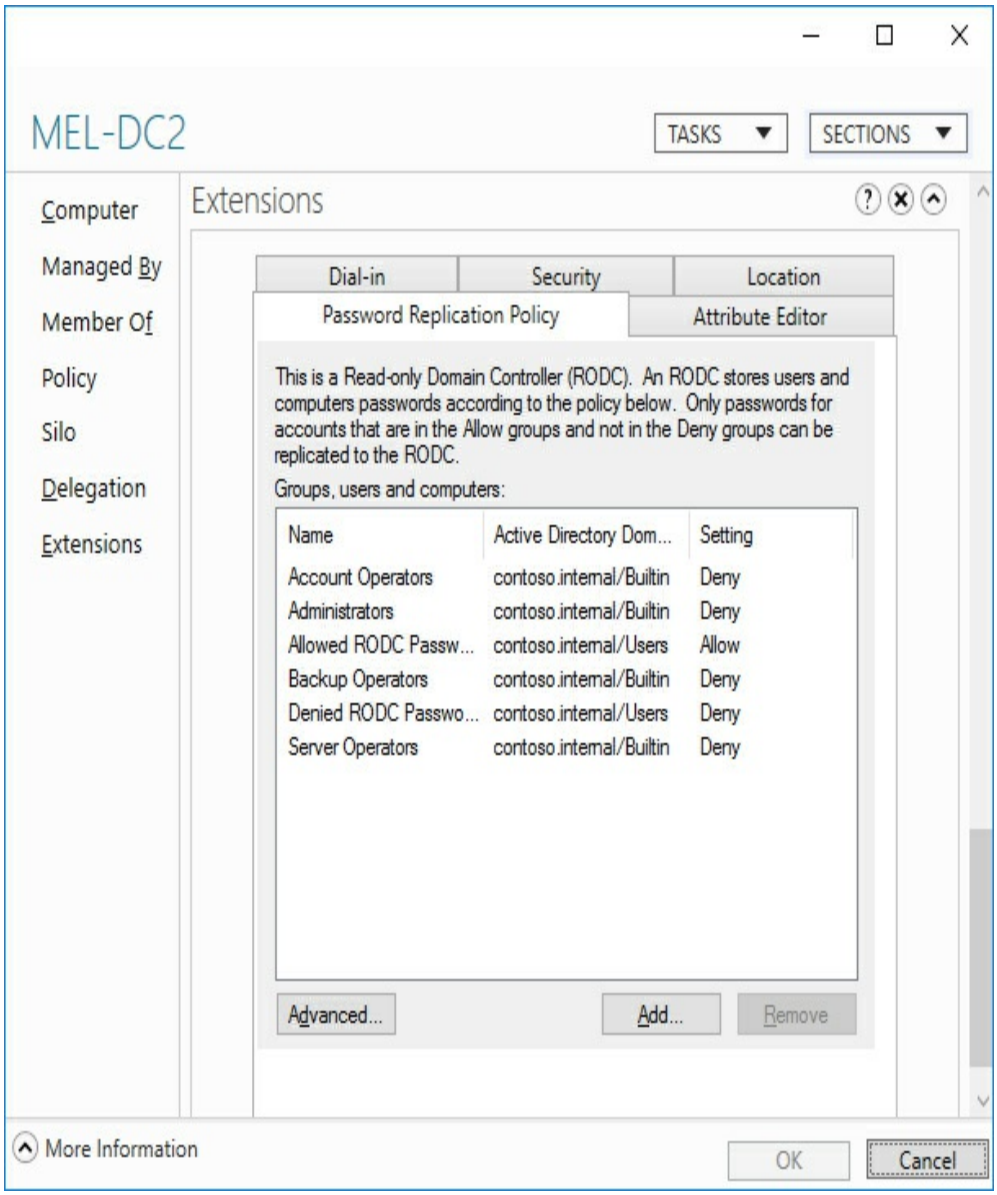


Figure 4-5 Password Replication Policy

RODC local administrators

Because you deploy RODCs as a security measure, they are almost always placed at branch office sites. Because resources at branch office sites are often sparse, it is also likely that you'll co-locate other services on the server hosting the RODC role. For example, a server that functions as an RODC can also function as a file server, DNS server, DHCP server, and local intranet server.

If the computer hosting the RODC role also needs to host other roles, you might need to grant administrator access to a user who works at the branch office (but who is not a member of your organization's usual IT staff) in case your normal remote administration techniques don't work. RODCs differ from normal domain controllers in that you can grant local administrator access without having to make the user a member of the Domain Admins group.

To configure a user to function as a local administrator on the computer that hosts the RODC role, edit the properties of the RODC's computer account and configure a security group for the Managed By setting.

In Chapter 19, "Hardening Windows Server and Active Directory," you'll learn about Just Enough Administration (JEA), a technology that you could use to allow a user to perform remote administrative tasks on an RODC without being made an RODC local

Technet24

administrator.

Decommissioning an RODC

If you suspect that an RODC has been compromised, you can delete the RODC's account from the Domain Controllers container in Active Directory Users And Computers. When you do this, you get the option of resetting all passwords for user and computer accounts that were cached on the RODC, as well as the option of exporting a list of all potentially compromised accounts.

Virtual domain controller cloning

Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019 support virtual domain controller cloning. Rather than redeploying domain controllers from scratch each time you need one, domain controller cloning allows you to take an existing virtual machine, make a copy of it, and deploy that copy.

Virtual domain controller cloning has the following prerequisites:

- Hypervisor must support VM-GenerationID. The version of Hyper-V included with Windows Server 2012 and later supports this technology, as does the most recent versions of VMware.
- The source domain controller needs to be running Windows Server 2012 or later.
- The domain controller that hosts the PDC Emulator role must be online and contactable by the cloned DC. The computer that hosts

the PDC Emulator role must also be running Windows Server 2012 or later.

- The source DC must be a member of the Cloneable Domain Controllers group.

You also need to create the *DCCloneConfig.xml* file. You can do this by using the *New-ADDCCloneConfig* cmdlet in PowerShell. When running this cmdlet, you need to specify the cloned DC's IPv4 address information and the site the cloned DC is deployed into. For example, to create the clone configuration file for a clone DC that has the IP address *10.10.10.42* with the subnet mask *255.255.255.0*, default gateway of *10.10.10.1*, DNS server address of *10.10.10.10*, and site name *MEL-SITE*, issue this command:

Click here to view code image

```
New-ADDCCloneConfigFile -IPv4Address 10.10.10.42 -
IPv4DefaultGateway 10.10.10.1
-IPv4SubnetMask 255.255.255.0 -IPv4DNSResolver 10.10.10.10 -
Static -SiteName MEL-SITE
```

After the clone configuration file is created, you import the VM using this file and specify a copy of the source DC's exported virtual hard disk.

AD DS STRUCTURE

Active Directory is made up of forests and domains. A forest is a collection of Active Directory domains that share a schema and some security principals. The

majority of organizations in the world have a single forest domain. Multiple domain forests are generally used by larger, geographically dispersed organizations.

Domains

For the majority of organizations in the world, a single domain would be sufficient. There are two general reasons for having multiple domains in a forest. The first is that your organization is geographically dispersed, and there are issues around domain replication traffic. The second is that your organization is very large. A single domain can hold a staggering number of objects. Unless your organization has tens of thousands of users, a single domain is usually more than enough.

A domain tree is a collection of domains that share a namespace in a parent–child relationship. For example, the domains *australia.contoso.com* and *tonga.contoso.com* would be child domains of the *contoso.com* domain.

You should always deploy at least two domain controllers per domain for redundancy purposes. Make sure that if you have a multi-domain forest, you are making regular backups of the domain controllers in the root domain. There has been more than one organization with a multidomain forest that has had the root domain's domain controllers fail irreparably,

making it necessary to redeploy the entire forest from scratch.

Domain functional levels

The domain functional level determines which Active Directory features are available and which operating systems can participate as domain controllers within the domain. The domain functional level determines the minimum domain controller operating system. For example, if the domain functional level is set to Windows Server 2012, domain controllers must run Windows Server 2012 or later. This rule does not apply to member servers. You can have a domain running at the Windows Server 2016 functional level that still has servers running the Windows Server 2008 R2 operating system. Unlike previous versions of the Windows Server operating system, which introduced new domain functional levels, there is no Windows Server 2019 domain functional level.

> **Inside OUT**
>
> **Run DC at highest OS level**
>
> Because domain controllers are so important to your network, Microsoft recommends that, from a security perspective, you always run domain controllers with the most recent operating system. Microsoft's argument is that the most recent operating system is the most secure and that you should pay special attention to the security of your organization's domain controllers.

Windows Server 2019 DCs support the following

functional levels:

- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2
- Windows Server 2016

You can introduce a DC running Windows Server 2019 to a domain at the Windows Server 2008 functional level as long as all the appropriate updates are installed. You can raise the functional level after you've retired existing domain controllers running older versions of the Windows Server operating system. If all your domain controllers are running Windows Server 2019, you should update your domain functional level to the Windows Server 2016 functional level.

> **Inside OUT**
>
> **Check your functional level**
>
> Even when older Windows Server DCs have long been removed, it's not unusual for people to have forgotten to elevate the functional level. Because you get greater functionality by raising functional levels and there is no downside to doing so beyond not being able to introduce a domain controller running an operating system below that of the functional level, you should raise functional levels as high as possible.

Forests

A forest is a collection of domains that share a schema. There are automatic trust relationships between all

domains in a forest. Accounts in one domain in a forest can be granted rights to resources in other domains. As mentioned earlier in this chapter, forests don't need to have a contiguous namespace. For example, a forest can contain both the *contoso.com* and *Adatum.com* domains.

There are several reasons why an organization might have multiple forests with trust relationships configured between those forests. The most common reason is that one organization has acquired another organization, and multiple forests exist until such a time that the users and resources hosted in the forest of the acquired organization are moved to the forest of the acquiring organization. A less common one is that an organization is splitting off a part of itself and users need to be migrated out of the existing forest and into a new forest prior to the split occurring.

Forest functional levels are determined by the minimum domain functional level in the forest. After you've raised the domain functional levels in the forest, you can raise the forest functional level. Unlike in previous versions of Active Directory, it is possible to lower both the forest and domain functional levels after they have been raised. An important caveat with this is that the forest functional level must be lowered so that it is never higher than the lowest intended domain functional level.

Technet24

Account and resource forests

Organizations are increasingly deploying what are known as Enhanced Security Administrative Environment (ESAE) forests. In an ESAE forest design, all the accounts used for administrative tasks in the production forest are hosted in a second forest known as the ESAE, bastion, or administrative forest. The ESAE forest is configured with one-way trust relationships with the production forest.

In their native administrative forest, the accounts used for administrative tasks in the production forest are traditional unprivileged user accounts. These accounts and the groups that they are members of in the administrative forest are delegated privileges in the production forest.

The advantage of this approach is that should one of these accounts become compromised, it can't be used to alter any permissions or settings in the administrative forest because the account only has privileges in the production forest.

Organizational units

The benefit of a comprehensive OU structure is that it simplifies the delegation of permissions over objects and allows you to apply GPOs closer to the accounts that they influence. While it's possible to filter GPOs based

on security groups or WMI queries, it's important to remember that each GPO within a computer and user account's scope must be checked to see if it applies. If you have all of your GPOs applied at the domain level and filtered on the basis of security group or WMI query, your Group Policy processing time at startup and sign-on will be substantially longer than if you apply GPOs as close to the accounts they influence as possible.

Better privilege delegation targeting is also a great reason to have a good OU structure. When delegating privileges, you want to ensure that you only delegate privileges over an appropriate set of objects. For example, imagine that you want to allow each department's administrative assistant to have the ability to reset passwords for users within their respective department. If all the user accounts for your organization are located in one OU, you won't be able to do this because, at best, you'll provide each administrative assistant with the ability to reset every user's password in the organization. If user accounts are located in department-specific OUs, you can then delegate the ability to reset passwords on a per-OU basis. This way, the administrative assistant for the Marketing department can be granted the ability to change passwords for user accounts located in the Marketing OU without granting him or her the ability to

change passwords for user accounts located in the Accounting OU.

> **Inside OUT**
>
> **OU structures don't always relate to organizational structure**
>
> Many organizations have had an AD deployment for more than a decade and a half. A big problem for many is that while the company goes through regular reorganizations, the OU structure doesn't go through a similar reorganization. This might mean that the OU structures don't in any way represent the organizational structure. You should periodically review your organization's OU structure to ensure that it maps in a sensible way to the organizational structure. While you might not implement GPOs at the OU level or delegate privileges, it doesn't mean that you won't do so at some point in the future.

When you create OUs in Windows Server 2019, the OUs are automatically configured so that accidental deletion protection is enabled.

Flexible Single Master Operations roles

The Flexible Single Master Operations (FSMO) roles are five special roles present on domain controllers. Two of these roles—the schema master and the domain naming master—are unique within each forest. The other three roles—PDC Emulator, infrastructure master, and RID master—must be present within each domain in the forest. For example, in a three-domain forest, there is only one schema master and domain naming master, but each domain in the forest has its own PDC emulator, infrastructure master, and RID master.

By default, FSMO roles are allocated to the first domain

controller in a domain. After you have more than one domain controller in each domain, you should manually start to move FSMO roles to other domain controllers. This protects you from a situation where the first domain controller deployed in each domain goes offline and all FSMO roles become unavailable. When you do need to take a domain controller offline for an extended period, ensure that you transition any FSMO roles that it hosts to another domain controller in the same domain.

Schema master

The schema master is the single server in the forest that is able to process updates to the Active Directory schema. The Active Directory schema defines the functionality of Active Directory. For example, by modifying the schema, you can increase the available attributes for existing objects as well as enable Active Directory to store new objects. Products including Exchange Server and Configuration Manager require that the default Active Directory schema be extended prior to product installation so that each product can store important data in Active Directory.

The domain controllers that host the schema master role should be located in the root domain of the forest. If you need to extend the schema prior to installing products such as Exchange Server, do so either on the

computer that hosts the schema master role or on a computer in the same site. The account used to extend the schema needs to be a member of the Schema Admins security group.

You can determine which computer hosts the schema master role by running the following PowerShell command:

Click here to view code image

```
Get-ADForest contoso.internal | FT SchemaMaster
```

Domain naming master

The domain naming master is a forest level role that is responsible for managing the addition and removal of domains from the forest. The domain naming master also manages references to domains in trusted forests. In a multi-domain environment, the domain controller that hosts this forest level role should be deployed in the root forest.

You can determine which server hosts the domain naming master role by running the following PowerShell command:

Click here to view code image

```
Get-ADForest contoso.internal | FT DomainNamingMaster
```

PDC emulator

The PDC emulator role is a domain level FSMO role that is responsible for handling both changes to account passwords as well as domain time synchronization. PDC emulators in child domains synchronize time against the PDC emulator in the forest root domain. You should ensure that the PDC emulator in the forest root domain synchronizes against a reliable external time source.

To determine which domain controller in a specific domain hosts the PDC emulator role, run the following PowerShell command:

Click here to view code image

```
Get-ADDomain contoso.internal | FT PDCEmulator
```

Infrastructure master

The computer that hosts the infrastructure master role keeps track of changes that occur in other domains in the forest as they apply to objects in the local domain. You should avoid having the infrastructure master role co-located with a domain controller that hosts the global catalog server role unless all domain controllers in the domain are configured as global catalog servers. You can determine which computer in a domain hosts the infrastructure master role by running the following PowerShell command:

Click here to view code image

```
Get-ADDomain contoso.internal | FT InfrastructureMaster
```

RID master

The RID (relative identifier) master processes relative ID requests from domain controllers in a specific domain. Relative IDs and domain Security IDs are combined to create a unique Security ID (SID) for the object. There is a RID master in each domain in the forest. You can use the following PowerShell command to determine which computer hosts the RID Master role:

Click here to view code image

```
Get-ADDomain contoso.internal | FT RidMaster
```

Seizing FSMO roles

In some cases, a domain controller hosting a FSMO role fails and you need to seize the FSMO role to move it to another domain controller. For example, to move the RID Master, Infrastructure Master, and Domain Naming Master roles to a domain controller named MEL-DC2, run the following command:

Click here to view code image

```
Move-ADDirectoryServerOperationMasterRole -Identity MEL-DC2 -
OperationMasterRole
RIDMaster,InfrastructureMaster,DomainNamingMaster -Force
```

ACCOUNTS

Accounts represent the identities of security principals

in an Active Directory environment. The most common type of account is a user account, which represents a person as he or she interacts with the Windows environment. IT Ops personnel also need to regularly deal with computer accounts, group accounts, and service accounts.

User accounts

User accounts almost always represent real people in an Active Directory environment, with the caveat that some user accounts are used for services rather than traditional users signing in to their desktop computers.

In many organizations, a user account is nothing more than a username, a password, and a collection of group memberships. User accounts can contain substantial additional information, including:

- First name
- Last name
- Middle initial
- Full name
- Office information
- Email address
- Webpage
- Job title
- Department
- Company

- Manager
- Phone numbers (main, home, mobile, fax, pager, IP phone)
- Address
- User profile location
- Log on script
- Home folder
- Remote desktop service profile
- Dial-in permissions
- Published certificates
- Remote Desktop Sessions settings
- Remote Control settings

You should configure important accounts to be protected from deletion by enabling the Protect From Accidental Deletion option, as shown in Figure 4-6. Unless this option is removed, the account cannot be deleted. This setting doesn't stop an account from being removed deliberately, but it does stop the account from being deleted accidentally.

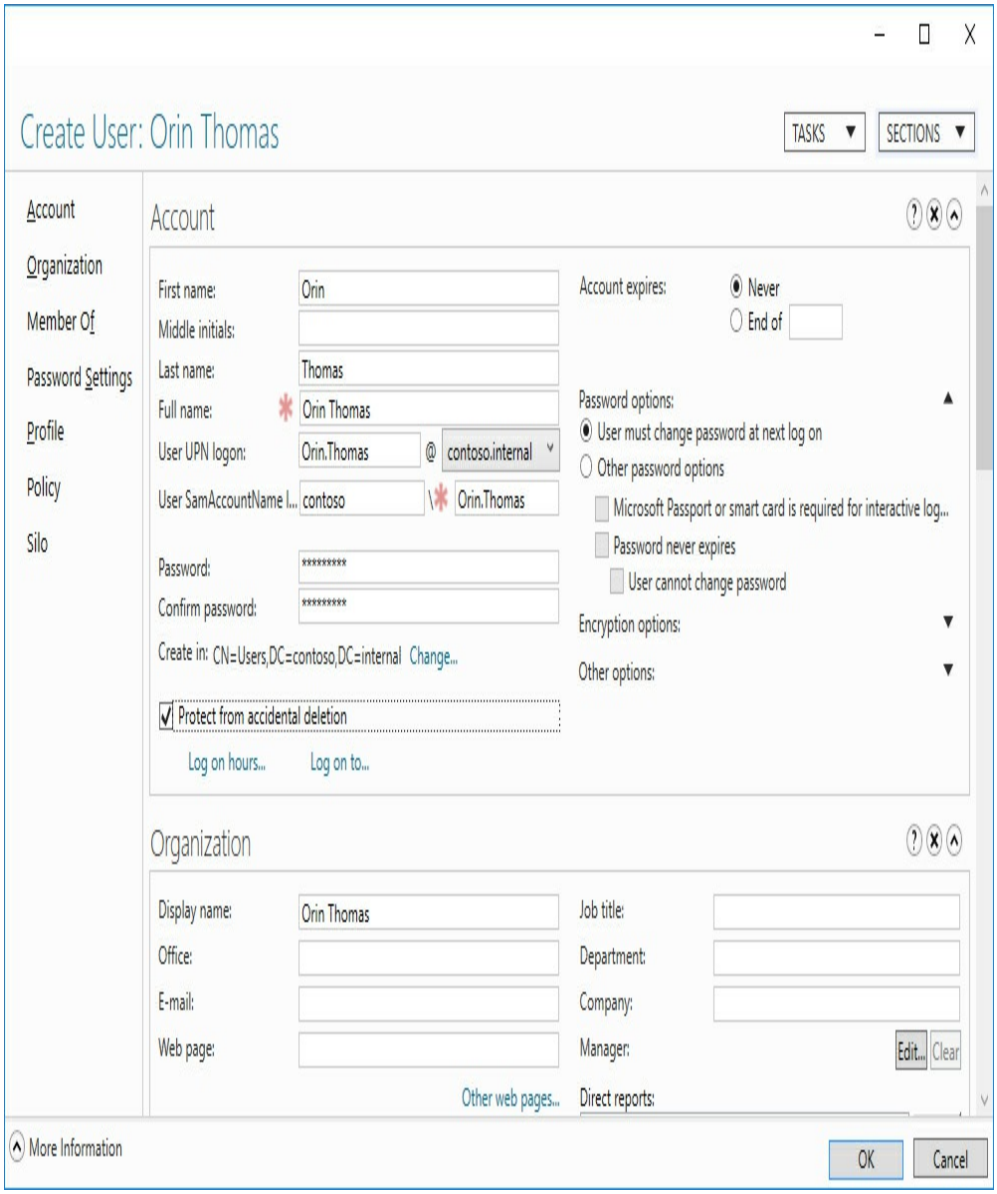


Figure 4-6 User account properties

You learn about securing accounts with log on restrictions, password policies, and authentication silos in Chapter 19, "Hardening Windows Server and Active Directory."

Technet24

Computer accounts

Computer accounts represent the computer within Active Directory. You often move computer accounts to specific OUs and then apply group policies to those OUs as a way of configuring the computer. When you join a computer to the domain, the computer account is automatically placed in the default Computers container.

If a computer account becomes desynchronized from the domain of which it is a member and loses its trust relationship, you can repair the relationship by signing in with the local Administrator account and running the following PowerShell command:

Click here to view code image

```
Test-ComputerSecureChannel -Credential Domain\<AdminAccount> -
Repair
```

If you don't know the local Administrator account password but suspect that cached domain administrator credentials might be present, disconnect the computer from the network, either by physically removing the Ethernet connection or disconnecting the virtual network adapter, sign in using those credentials, and then run the PowerShell command previously mentioned. The Local Administrator Password Solution, which ensures that you don't forget the local Administrator account password, is covered in Chapter

19, “Hardening Windows Server and Active Directory.”

Group accounts

Active Directory supports three different types of group account scopes: domain local, global (the default), and universal. It also supports two group types: Security and Distribution. You use security groups to control access to resources, delegate permissions, and for email distribution. Distribution groups can only be used for email distribution. If your organization uses Exchange Server, you manage distribution groups through Exchange. Best practice is to place users within global groups, add those global groups to domain local groups, and assign permissions and rights directly to the domain local groups. By default, members of the Account Operators, Domain Admins, or Enterprise Admins groups can modify the membership of groups. You can also delegate the right to manage groups, which is covered in Chapter 19, “Hardening Windows Server and Active Directory.”

Universal

A universal group can hold accounts and groups from the same forest. Universal groups are stored in the global catalog. If you change the membership of a universal group, this change replicates to all global catalog servers in the forest. Universal groups are great in single forest environments where replication traffic

Technet24

isn't an issue or where there are few changes to universal group membership. Universal groups can be nested within other universal groups, or they can be added to domain local groups.

Global

Global groups can contain user accounts from its home domain. Global groups can also contain other global groups from its home domain. Global groups can be members of universal groups and domain local groups.

Domain local

Domain local groups can have universal groups, global groups, and domain local groups as members. Domain local groups can host accounts from any domain in the forest and accounts from domains in trusted forests. Domain local groups can only be added to domain local groups within its own domain. A domain local group can only be assigned rights and permissions within its own domain. Domain local groups cannot be added to global or universal groups.

Default groups

The groups listed in Table 4-1 are included in a Windows Server 2019 domain by default.

Table 4-1 Important groups

Group	Function
Access Control Assistance Operators	Members can query authorization attributes and permissions for resources on the computer.
Account Operators	Members can manage domain user and group accounts.
Administrators	This is the built-in administrators group. On a DC, this provides unlimited access to the computer and domain.
Allowed RODC Password Replication Group	Members can have their passwords replicated to RODCs.
Backup Operators	Members can override security restrictions to back up or restore files.
Cert Publishers	Members can publish certificates to AD DS.
Certificate Services DCOM Access	Members can connect to Certification Authorities.
Cloneable Domain Controllers	Domain controllers that can be cloned.

Cryptographic Operators	Can perform cryptographic operations.
Denied RODC Password Replication Group	Members cannot have their passwords replicated to RODCs.
Distributed COM Users	Can launch, activate, and use Distributed COM objects.
DNSAdmins	DNS administrators group.
DNSUpdateProxy	Group for DNS clients who are able to perform dynamic updates on behalf of services, such as DHCP servers.
Domain Admins	Members can administer the domain. Members of this group are automatically members of the local Administrators group on every computer in the domain. Default owner of all objects created in Active Directory.
Domain Computers	All computers that are members of the domain.
Domain Controllers	All domain controllers in the domain.

Domain Guests	Hosts the built-in domain guest account.
Domain Users	Hosts all user accounts in the domain.
Enterprise Admins	Only present in the root domain of a forest. Can make forest-wide changes, such as raising forest functional level or adding domains.
Enterprise Key Admins	Members can perform administrative tasks on key objects at the forest level.
Enterprise Read-Only Domain Controllers	Group for enterprise read-only domain controllers.
Event Log Readers	Can view the contents of event logs on the computer.
Group Policy Creator Owners	Members can create and modify GPOs.
Hyper-V Administrators	Can manage Hyper-V on the local machine.

IIS_IUSRS	Internet Information Services built-in group.
Incoming Forest Trust Builders	Members can create incoming, one-way trusts to the forest.
Key Admins	Members can perform administrative tasks on key objects at the domain level.
Network Configuration Operators	Members can configure networking features.
Performance Log Users	Members can schedule logging of performance counters and collect event traces.
Performance Monitor Users	Members can access performance counter data locally and remotely.
Pre-Windows 2000 Compatible Access	A group that exists for backward compatibility.
Print Operators	Members can manage printers deployed on domain controllers.

Protected Users	Members of this group have stricter security controls applied to their accounts.
RAS and IAS Servers	Members in this group are able to access the remote access settings of user accounts.
RDS Endpoint Servers	Servers in this group host VMs used with RemoteApp programs and personal virtual desktops.
RDS Management Servers	Servers in this group can perform administrative actions on other servers running the Remote Desktop Services role.
RDS Remote Access Servers	Servers in this group enable users of RemoteApp programs and personal virtual desktop to access resources.
Read-Only Domain Controllers	Read-only domain controllers in the domain.
Remote Desktop Users	Granted the right to sign on using Remote Desktop.
Remote	Members can access WMI resources over management

Management Users	protocols.
Replicator	Group supporting file replication at the domain level.
Schema Admins	Root domain group. Members are authorized to make changes to Active Directory Schema.
Server Operators	Members can administer domain member servers.
Storage Replica Administrators	Members can manage storage replica.
Terminal Server License Servers	Members can update Active Directory information about terminal services licensing.
Users	Built-in Users group.
Windows Authorization Access Group	Have access to computed *tokenGroupGlobalAndUniversal* attribute on User objects.

Service accounts

Service accounts are functionally user accounts used by

services to interact with the operating system and resources on the network. By assigning rights to the service account, you can limit what the service can or cannot do.

In Chapter 19, "Hardening Windows Server and Active Directory," you'll learn how to lock down service accounts. One of the key things to remember about a service account is that it should not have the right to log on locally to a computer but should have the log on as a service right. This is important because administrators in many organizations have bad habits where they grant service accounts unnecessary rights and even go as far as to give all service accounts the same non-expiring password. Sophisticated attackers know this and use this to compromise service accounts and use them to gain privileged access.

Local System

The Local System (*NT AUTHORITY\SYSTEM*) account is a built-in account. It has privileges equivalent to a local administrator account on the local computer. It acts with the computer account's credentials when interacting on the network. This is the most powerful service account, and generally you should be reluctant to assign this service account manually to a service given its extensive privileges.

Local Service

Technet24

The Local Service (*NT AUTHORITY\LocalService*) account has the same level of privilege as user accounts that are members of the local Users group on a computer. This account has fewer local privileges than the Local System account. Any services assigned to this account access network resources as a null session without credentials. Use this account when the service doesn't require network access or can access network resources as an anonymous user and requires only minimal privileges on the computer it is being used on.

Network Service

The Network Service (*NT AUTHORITY\NetworkService*) account is similar to the Local Service account in that it has privileges on the local computer equivalent to those assigned to a member of the local Users group. The primary difference is that this account interacts with the computer account's credentials with resources on the network.

GMSA

A group managed service account (GMSA) is a service account that is managed by the domain. This means that rather than having to update the service account's password manually, Active Directory updates the password in line with the domain password policy. Many organizations that use regular user accounts as

service accounts tend to set a static password for these accounts, which is often simple rather than complex.

Group Managed Service Accounts require the forest functional level to be Windows Server 2012 or higher. Forests at the Windows Server 2008 level support a version of a GMSA called an MSA, but this is more limited than a GMSA, with each MSA only being able to be installed on a single machine.

Prior to using GMSAs, you'll need to create a Key Distribution Services (KDS) root key. You can do this with the following command:

Click here to view code image

```
Add-KdsRootKey -EffectiveTime ((get-date).addhours(-10))
```

You manage GMSAs using Windows PowerShell. To create a GMSA, specify the name of the account, a DNS hostname associated with the account, and the name of the security principals that are allowed to use the account. For example, to create a GMSA named *MEL-SQL-GMSA* for the *contoso.internal* domain that can be used by servers in the MEL-SQL-Servers security group, enact the following command:

Click here to view code image

```
New-ADServiceAccount MEL-SQL-GMSA -DNSHostname MEL-SQL-
GMSA.contso.internal
-PrincipalsAllowedToRetrieveManagedPassword MEL-SQL-Servers
```

To install the account on a specific server, so that you can use it, run the *Install-ADServiceAccount* cmdlet. For example, to install the *MEL-SQL-GMSA* account on a server so that you can assign the account to services, enact the command:

Click here to view code image

```
Install-ADServiceAccount MEL-SQL-GMSA
```

Prior to running this command, you may need to install the RSAT ADDS Tools on the local server. You can do this by running the command:

Click here to view code image

```
Install-WindowsFeature RSAT-ADDS-Tools
```

When assigning the account to a service, you should clear the Password and Confirm Password textbox. You'll need to append a *$* to the account name when configuring it as shown in Figure 4-7. Windows Server 2019 recognizes that the account is a GMSA and manages the password settings.

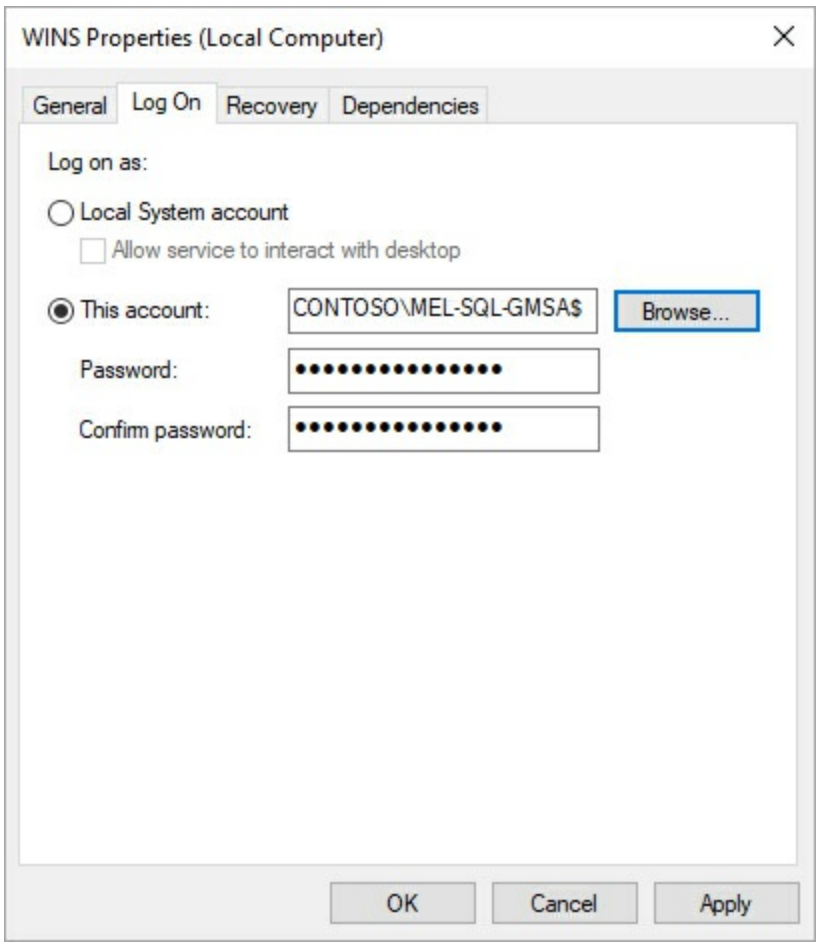


Figure 4-7 GMSA configuration

Virtual account

A virtual account is the local equivalent of a group managed service account. Virtual accounts are supported by products such as SQL Server 2012 as an alternative to the default built-in accounts. You can create virtual service accounts by editing the property of a service and setting the account name to **NT Service\<ServiceName>**.

Technet24

GROUP POLICY

Group policy provides a central way of managing user and computer configuration. You can use Group Policy to configure everything from password and auditing policies to software deployment, desktop background settings, and mappings between file extensions and default applications.

Inside OUT

Not the only configuration management option

It's fair to say that Group Policy is really starting to show its age. On the server side, Microsoft is pushing Desired State Configuration as the primary method of managing server configuration. On the client side, you often get a better result using tools such as Configuration Manager for on-premises deployments and Intune for clients that aren't domain joined.

GPO management

The Group Policy Management console is the primary method that the Group Policy is managed through in a standard Windows Server environment. This tool allows you to view Group Policy items from the forest level down. It provides information on which GPOs are linked at the domain, site, and OU level. You can see a list of all GPOs in a domain under the Group Policy Objects node. There is also a WMI Filters node that contains all WMI filters configured for the domain, a Starter GPOs node that contains all starter GPOs, a Group Policy Modeling node, and a Group Policy Results node. By right-clicking

a GPO and selecting Edit, you can edit a GPO in the Group Policy Management Editor. You can output information about a GPO, including which settings are configured and where the GPO is linked, to a file in HTML format using the *Get-GPOReport* cmdlet.

Advanced Group Policy Management

Advanced Group Policy Management (AGPM) is available to Software Assurance customers through the Microsoft Desktop Optimization Pack. AGPM integrates with the existing Group Policy Management console, extending its functionality. AGPM has several benefits over the standard Group Policy Management console.

The biggest advantage of AGPM is that it enables you to apply a workflow to the management of GPOs, allowing you to use versioning and offline editing. This is an advantage over the standard Group Policy Management console where any changes made to an active Group Policy are automatically propagated to clients that are influenced by that policy. Having offline functionality for editing GPOs allows for GPOs that are already active to be more extensively tested during the revision process than might be the case in a traditional GPO deployment.

AGPM offers a more advanced model for editing GPOs because it allows administrators to track precisely which person in an organization made a specific change to a

GPO. In a traditional Windows Server deployment, even though it is possible to control who has the permissions required to edit and apply a GPO, if multiple people have these permissions, it can be almost impossible to determine who made a specific change to an existing GPO. Being able to track the changes made to a GPO is very important in a high-security environment.

AGPM uses the following roles in its delegation architecture:

- **AGPM Administrator.** Users assigned this role can perform all tasks in AGPM. This includes configuring domain-wide options as well as delegating permissions.
- **Approver.** Users assigned this role can deploy GPOs into the domain. Approvers are also able to approve or reject requests from users assigned the Editor role. Users assigned this role are also able to create and delete GPOs but cannot edit GPOs unless assigned the Editor role.
- **Reviewer.** Users assigned the Reviewer role can view the contents of GPOs and compare the contents of GPOs, but they do not have the ability to modify or deploy GPOs. The Reviewer role is designed to review and audit GPO management.
- **Editor.** Users assigned this role can view and compare GPOs as Reviewers can. They are also able to check GPOs out of the archive, edit them offline, and then check those GPOs back into the archive. They are unable to deploy GPOs but can request GPO deployment.

Backup and restore

In any environment where you make a change, you need to have the option to roll back that change. Even if you have the AGPM, it is a good idea to regularly back up

Group Policy Objects. While GPOs are backed up when you perform a system state backup, they can be troublesome to restore if you delete or modify them in an unintended way. Using a tool like AGPM can reduce the chance that you can't roll back to a previous configuration, but if you don't have AGPM, you should make a backup of a GPO prior to making substantial, or even insubstantial, policy modifications.

Backing up a GPO is straightforward. In the Group Policy Management console, right-click the policy and then click Back Up. On the Back Up Group Policy Object dialog box, enter the folder to which you are backing up the GPO and provide a description. You can back up all GPOs in your organization by right-clicking the Group Policy Objects node within the Group Policy Management console and clicking Back Up All.

You can manage the backups that you have taken of Group Policy Objects by right-clicking the Group Policy Objects node and then clicking Manage Backups. This opens the Manage Backups dialog box. Using this dialog box, you can:

- Restore a backed-up GPO
- Verify what settings are applied in a backup GPO
- View where the GPO was linked and how it was filtered
- View the delegation settings on the GPO
- Delete a specific GPO backup

You can also restore an individual GPO by right-clicking the GPO and then clicking Restore From Backup. When you do this, you have the option of restoring from any backup you've taken of that GPO.

> **Inside OUT**
>
> **Restoring GPO doesn't restore links**
>
> Note that restoring a GPO does not restore where the GPO was linked. You can, however, configure that prior to restoring the GPO by viewing information about the backed-up GPO.

You can use the Import Settings Wizard to import settings from a backed-up GPO into an existing GPO. You can use this technique to quickly duplicate GPOs by backing one up, creating a new "blank" GPO, and then importing the settings into the new GPO. To run the Import Settings Wizard, right-click a GPO under the Group Policy Objects in the Group Policy Management console, and then click Import Settings. Importing settings into a GPO overwrites all current settings in the GPO, effectively replacing the current GPO with the backed-up GPO. You do get the option to back up the GPO that you are going to overwrite when running the wizard; even if you are completely sure of what you are doing, this is a good option to take.

Policy processing

Group policy is processed in the following order, with

subsequent policies overriding previously applied policies where conflicts exist:

- Local
- Site
- Domain
- Organizational unit

For example, if one policy configures a setting at the AD site–level and another policy configures the same setting at the organizational unit–level, the policy that configures the setting at the organization unit–level is applied.

The exception to this is when you configure one of the following settings:

- **Block Inheritance.** The Block Inheritance option allows you to block upstream policies from applying. You can configure this option at the Domain or OU level. It blocks all upstream settings except those configured through a policy that has the No Override setting.
- **No Override.** You use the No Override setting to stop the Block setting. Use this on policies that you need to have enforced.

Resultant Set of Policy tool

The Resultant Set of Policy tool allows you to generate a model of Group Policy application, allowing you to figure out which policies apply to particular objects within the domain. Resultant Set of Policy allows you to figure out why Group Policy application isn't behaving

in the way that you expect and allows you to resolve Group Policy conflicts.

There are two ways you can calculate Resultant Set of Policy. The first is to use Group Policy Modeling. The second is to use Group Policy Results. The difference between these is as follows:

- Group Policy Modeling allows you to view the impact of altering site membership, security group membership, filtering, slow links, loopback processing, and the movement of accounts to new OUs on the application of policy.
- Group Policy Results allows you to troubleshoot the application of policy by telling you which settings apply to a specific user or computer account.

By default, members of the Domain Admins and Enterprise Admins groups can generate Group Policy Modeling or Group Policy Results information. You can delegate permissions so users can perform these tasks at the OU or domain level.

Inside OUT

Policy in multidomain forests

Policies linked at the site level in multidomain forests are stored in the root domain. This can cause challenges with Group Policy application in scenarios in which a site-linked policy applies but where no DC with a root domain membership is present in the local site.

Filtering

GPO filtering allows you to determine whether a linked GPO applies based on group membership or the results

of a WMI query. For example, you could write a WMI query that performs a check so that a GPO applies only if the computer has more than 4 GB of RAM. If the computer has more than 4 GB of RAM, the policy applies normally. If the computer has less than 4 GB of RAM, the GPO does not apply.

Inside OUT

Site-linked GPOs

If Group Policy processing takes too long because of site-linked GPOs, the simplest solution is to not link polices at the site level; instead, you apply the links at the domain level and use Group Policy filtering through a WMI query to ensure location-specific GPOs are applied. An alternative is to place domain controllers that have root domain membership at each site, although this is likely to be more expensive than the policy filtering option.

Group Policy preferences

Group Policy preferences work around the idea of eliminating (or at least substantially reducing) the need for traditional startup and login scripts. Login scripts have a way of becoming convoluted over time. Group Policy preferences allow simplification of common login and startup script tasks such as drive mappings and setting environment variables.

By reducing or eliminating some of the complexity of login scripts, you can use Group Policy preferences to reduce configuration errors. You can use Group Policy preferences to configure the following:

- Applications

- Drive mappings
- Environment variables
- File updates
- Folders
- *Ini* files
- Registry settings
- Shortcuts
- Data sources
- Devices
- Folder options
- Internet settings
- Local users and groups
- Network options
- Network shares
- Power options
- Printer settings
- Regional options
- Scheduled tasks
- Start Menu settings

Some of these items can also be configured using a traditional Group Policy. If an item is configured in the same GPO using both policy and preferences, the traditional setting takes precedence. The difference between a Group Policy preference and a normal Group Policy setting is that users can change a Group Policy

preference if they have the appropriate permissions. For example, users can un-map a mapped network drive. The drive would remain unmapped until the user logged in again, at which point, it would be remapped. Generally, if you want to enforce a setting, use a standard Group Policy. If you want to apply the setting and allow users to change it, use a Group Policy preference. The closest you can come to enforcing a Group Policy preference is to disable the Apply Once And Do Not Reapply setting in the policy item's configuration. This way, the preference is applied each time Group Policy refreshes.

You can target Group Policy preferences so that different preferences can apply to the same item types within a single GPO. You can use the following items to restrict how a Group Policy preference applies:

- The computer has a battery.
- The computer has a specific name.
- The computer has a specific CPU speed.
- Apply by or after a specific date.
- The computer has a certain amount of disk space.
- The computer is a member of a domain.
- The computer has a particular environment variable set.
- A certain file is present on the computer.
- The computer is within a particular IP address range.
- The computer uses specific language settings.

- The computer meets the requirements of an LDAP query.
- The computer has a MAC address within a specific range.
- The computer meets the requirements of a WMI query.
- The computer uses a specific type of network connection.
- The computer is running a specific operating system.
- The computer is a member of a specific OU.
- The computer has PCMCIA present.
- The computer is portable.
- The computer uses a specific processing mode.
- The computer has a certain amount of RAM.
- The computer has a certain registry entry.
- The user or computer is a member of a specific security group.
- The computer is in a specific Active Directory site.
- The computer has a Remote Desktop Setting.
- A specific time range is present.
- The user has a specific name.

Administrative templates

Group Policy Administrative Templates allow you to extend Group Policy beyond the settings available in GPOs. Common software packages, such as Microsoft Office, often include Administrative Templates that you can import to manage software-specific settings. In early versions of Windows Server, Administrative Templates were available as files in ADM format. Since the release of Windows Server 2008, Administrative Templates are

available in a standards-based XML file format called ADMX.

To be able to use an Administrative Template, you can import it directly into a GPO using the Add/Remove Templates option when you right-click the Administrative Templates node. A second option is to copy the Administrative Template files to the Central Store, located in the *c:\Windows\Sysvol\sysvol\<domainname>\Policies\PolicyDefinitions* folder on any domain controller. You might need to create this folder if it does not already exist. After the folder is present, the template is then replicated to all domain controllers, and you can access the newly imported Administrative Templates through the Administrative Templates node of a GPO.

RESTORING DELETED ITEMS

Sometimes an Active Directory account, such as a user account or even an entire OU, is accidentally—or on occasion, maliciously—deleted. Rather than go through the process of re-creating the deleted item or items, it's possible to restore the items. Deleted items are retained within the AD DS database for a period of time specified as the *tombstone lifetime*. You can recover a deleted item without having to restore the item from a backup of Active Directory as long as the item was deleted within the Tombstone Lifetime window.

The default tombstone lifetime for an Active Directory environment at the Windows Server 2008 forest functional level or higher is 180 days. You can check the value of the tombstone lifetime by issuing the following command from an elevated command prompt, (substituting *dc=Contoso,dc=Internal* for the suffix of your organization's forest root domain):

Click here to view code image

```
Dsquery * "cn=Directory Service,cn=Windows
NT,cn=Services,cn=Configuration,dc=Contoso,
dc=Internal" -scope base -attr tombstonelifetime
```

For most organizations the 180-day default is fine, but some administrators might want to increase or decrease this value to give them a greater or lesser window for easily restoring deleted items. You can change the default tombstone lifetime by performing the following steps:

1. From an elevated command prompt or PowerShell session, type **ADSIEdit.msc**.
2. From the Action menu, click Connect To. In the Connection Settings dialog box, click Configuration, as shown in Figure 4-8, and then click OK.

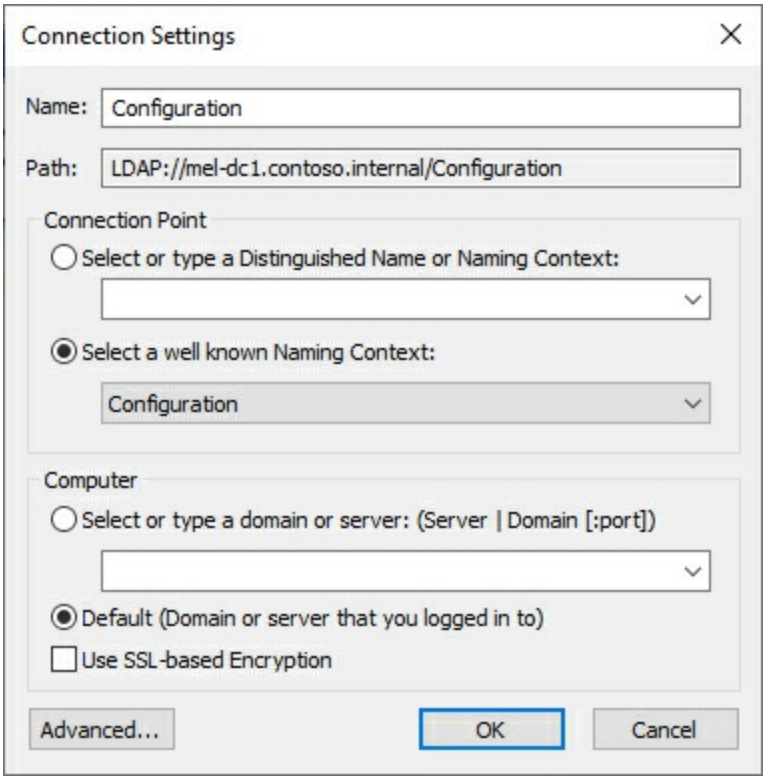


Figure 4-8 Connection Settings

3. Right-click the CN=Services, CN=Windows NT, CN=Directory Service node and click Properties.
4. In the list of attributes, select *tombstoneLifetime*, as shown in Figure 4-9, and click Edit.

Technet24

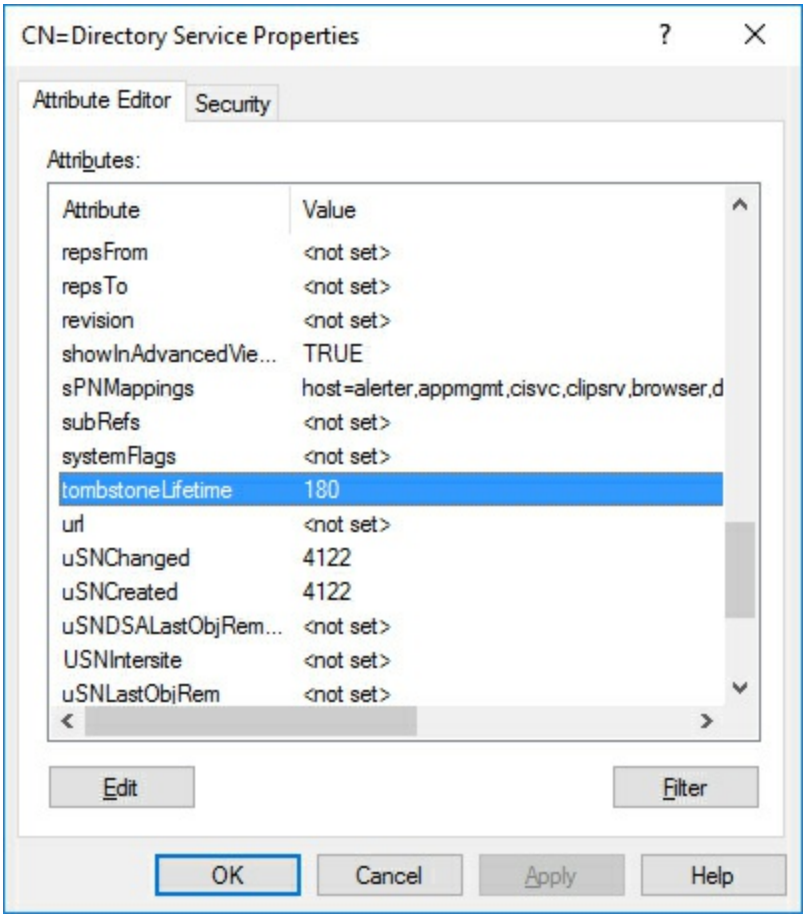


Figure 4-9 Tombstone lifetime

5. On the Integer Attribute Editor dialog box, enter the new value for the tombstone lifetime and click OK twice; close the ADSI Edit console.

Active Directory Recycle Bin

Active Directory Recycle Bin allows you to restore items that have been deleted from Active Directory but are still present within the database because the tombstone lifetime has not been exceeded. Active Directory Recycle Bin requires that the domain functional level be set to Windows Server 2008 R2 or higher.

Inside OUT

Enable AD Recycle Bin first

You can't use the Active Directory Recycle Bin to restore items that were deleted before you enabled Active Directory Recycle Bin.

Once activated, you can't deactivate the Active Directory Recycle Bin. There isn't any great reason to want to deactivate AD Recycle Bin once it is activated. You don't have to use it to restore deleted items should you still prefer to go through the authoritative restore process.

To activate the Active Directory Recycle Bin, perform the following steps:

1. Open the Active Directory Administrative Center and select the domain that you want to enable.
2. In the Tasks pane, click Enable Recycle Bin, as shown in Figure 4-10.

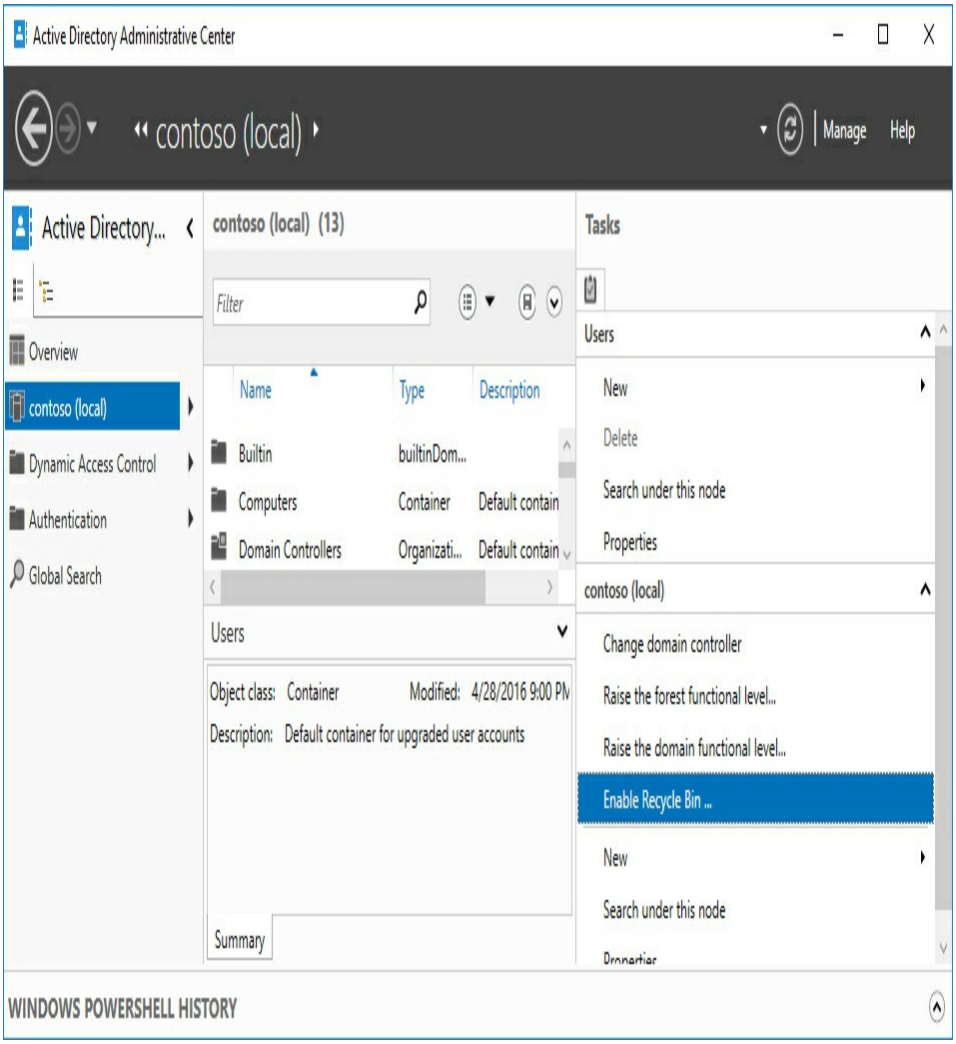


Figure 4-10 Enable Recycle Bin

After you have enabled the AD Recycle Bin, you can restore an object from the newly available Deleted Objects container. This is, of course, assuming that the object was deleted after the Recycle Bin was enabled and assuming that the tombstone lifetime value has not been exceeded. To recover the object, select the object in the Deleted Items container and then click Restore or Restore To. Figure 4-11 shows a deleted item being

selected, which can then be restored to its original location. The Restore To option allows you to restore the object to another available location, such as another OU.

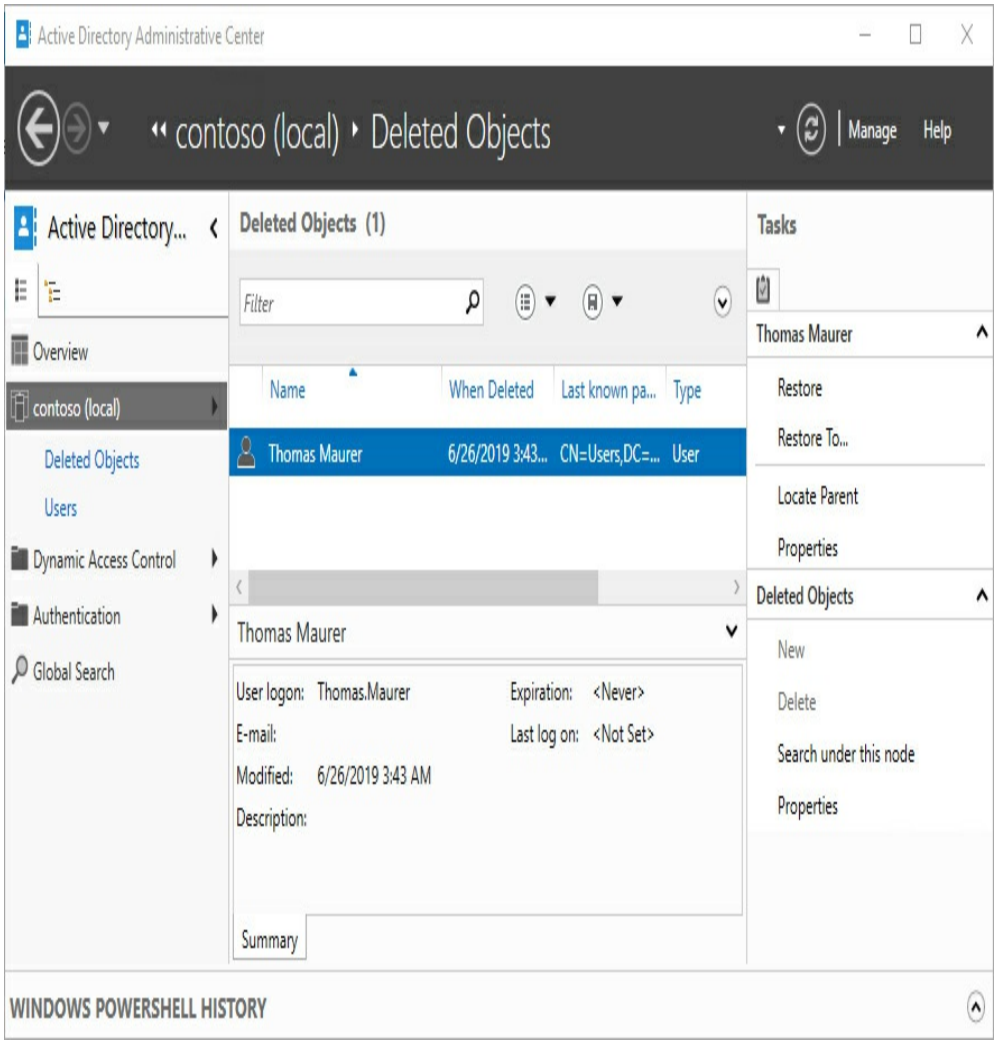


Figure 4-11 Deleted Objects container

Authoritative restore

An authoritative restore is performed when you want

the items you are recovering to overwrite items that are in the current Active Directory database. If you don't perform an authoritative restore, Active Directory assumes that the restored data is simply out of date and overwrites it when it is synchronized from another domain controller. If you perform a normal Restore on an item that was backed up last Tuesday when it was deleted the following Thursday, the item is deleted the next time the Active Directory database is synchronized. You do not need to perform an authoritative restore if you only have one domain controller in your organization because there is no other domain controller that can overwrite the changes.

Authoritative restore is useful in the following scenarios:

- You haven't enabled Active Directory Recycle Bin.
- You have enabled Active Directory Recycle Bin, but the object you want to restore was deleted before you enabled Active Directory Recycle Bin.
- You need to restore items that are older than the tombstone lifetime of the AD DS database.

To perform an authoritative restore, you need to reboot a DC into Directory Services Restore Mode. If you want to restore an item that is older than the tombstone lifetime of the AD DS database, you also need to restore the AD DS database. You can do this by restoring the system state data on the server. You'll likely need to

take the DC temporarily off the network to perform this operation simply because if you restore a computer with old system state data and the DC synchronizes, all the data that you wish to recover will be deleted when the domain controller synchronizes.

You can configure a server to boot into Directory Services Restore Mode from the System Configuration utility. To do this, select the Active Directory Repair option on the Boot tab as shown in Figure 4-12. After you've finished with Directory Services Restore Mode, use this same utility to restore normal boot functionality.

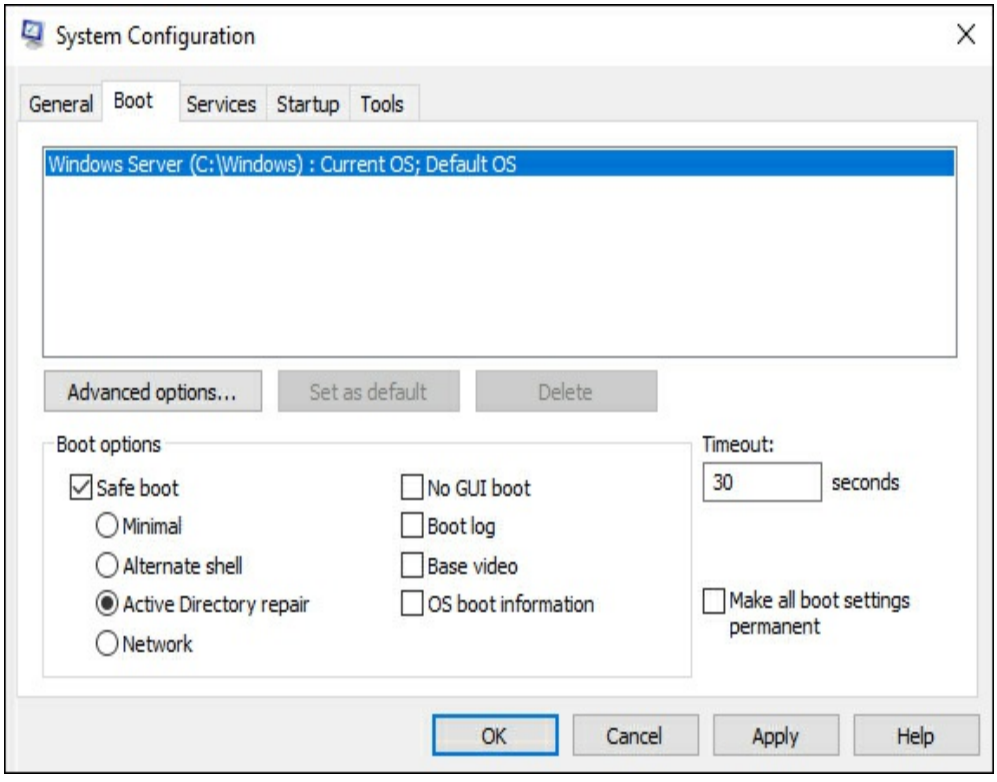


Figure 4-12 System Configuration

To enter Directory Services Restore Mode, you need to enter the Directory Services Restore Mode password.

To perform an authoritative restore, perform the following general steps:

1. Choose a computer that functions as a global catalog server. This DC functions as your restore server.
2. Locate the most recent system state backup that contains the objects that you want to restore.
3. Restart the restore server in DSRM mode. Enter the DSRM password.
4. Restore the system state data.
5. Use the following command to restore items (where *Mercury* is the object name, *Planets* is the OU that it is contained in, and *contoso.com* is the host domain):

   Click here to view code image

   ```
   Ntdsutil "authoritative restore" "restore
   object cn=Mercury,ou=Planets,dc=contoso,dc=com"
   q q
   ```

6. If an entire OU is deleted, you can use the Restore Subtree option. For example, if you deleted the Planets OU and all the accounts that it held in the *contoso.com* domain, you could use the following command to restore it and all the items it contained:

   Click here to view code image

   ```
   Ntdsutil "authoritative restore" "restore
   subtree OU=Planets,
   DC=contoso,DC=com" q q
   ```

Active Directory snapshots

You can use the Active Directory database mounting tool, *dsamain.exe*, to mount the contents of the AD DS database as it exists in snapshots or in backups so that you can interact with it using the Active Directory Users and Computers console. This provides you with a quick way of checking the state of a snapshot or backup without actually having to restore it to a production or development environment. For example, you might do this if you want to check what the AD DS database looked like at a specific point in time to determine the state of particular accounts or organizational units.

If you restored the AD DS database file *ntds.dit* to the location *c:\restore*, you could mount this file using the *dsamain.exe* utility on port 51389 by entering the following command:

Click here to view code image

```
Dsamain.exe /dbpath C:\restore\ntds.dit /ldapport 51389
```

You can then use Active Directory Users And Computers to connect to the specified port—in this case, port 51389—to view the mounted version of the AD DS database.

You can also create snapshots of the AD DS database that you can interact with without having to make a backup. A snapshot is a copy of the AD DS database as it exists at a specific point in time. You create a snapshot of the AD DS database by running the following

command from an elevated command prompt on a domain controller:

Click here to view code image

```
Ntdsutil "Activate Instance NTDS" snapshot create quit quit
```

You can list all current snapshots of the AD DS database on a domain controller by running the following command:

Click here to view code image

```
Ntdsutil snapshot "list all" quit quit
```

Each snapshot has an odd number associated with it next to the date the snapshot was taken. To select and mount a snapshot, use the command:

Click here to view code image

```
Ntdsutil snapshot "list all" "mount 1" quit quit
```

This mounts snapshot *1*. Make a note of the path on which the snapshot is mounted. You can then use the *dsadmin.exe* command to mount the snapshot so that it is accessible as an LDAP server. For example, to make the snapshot at path *C:\$SNAP_201606100306_VOLUMEC$* accessible on port 51389, use the command:

Click here to view code image

```
Dsamain.exe /dbpath
C:\$SNAP_201606100306_VOLUMEC$\WINDOWS\NTDS\ntds.dit /ldapport
```

```
51389
```

You can then use Active Directory Users And Computers to connect to the specified port to view the mounted snapshot. Only members of the Domain Admins and Enterprise Admins groups can view snapshots. While you can't directly copy objects from a snapshot to the production version of the AD DS database, you can use utilities such as CSVDE and LDIFDE to export information from a snapshot and then later import it into the production AD DS database.

MANAGING AD DS WITH POWERSHELL

Many Active Directory administrative tasks are repetitive. If you're likely to perform a task more than once, it's better to script it in PowerShell than work your way through the appropriate wizard multiple times.

There are three PowerShell modules related to Active Directory. The *Active Directory PowerShell* module (Table 4-2) is the one you're likely to use on a regular basis when managing Active Directory. The *GroupPolicy* module (Table 4-3) allows you to manage Group Policy from PowerShell. You use the *ADDSDeployment* (Table 4-4) module when performing deployment tasks.

Table 4-2 Active Directory module cmdlets

Technet24

Noun	Verbs	Function
ADAccount	Unlock, Search, Enable, Disable	Allows you to find, unlock, enable, or disable a user, computer, or service account.
ADAccount Authentication PolicySilo	Set	Allows you to configure the authentication policy or authentication policy silo of an account.
ADAccount Authorization Group	Get	Gets the security groups for a specified user, computer, or service account based on its token. Uses the global catalog to determine this information.
ADAccount Control	Set	Modifies the user account control values of an Active Directory user or computer account.
ADAccount Expiration	Set, Clear	Configure account expiration.
ADAccount Password	Set	Configure the password of a user, computer, or service account.
ADAccount Resultant PasswordReplicationPolicy	Get	Gets the password replication policy for a user, computer, or service account on a specific RODC.
ADAuthenti	Set, Remove,	Manipulate the properties of the AD DS

cationPolicy	New, Get	authentication policy.
ADAuthenticationPolicy Expression	Show	Displays Edit Access Control Conditions windows update or SSDL security descriptors.
ADAuthenticationPolicySilo	New, Remove, Get, Set	Manipulate Active Directory Domain Services authentication policy silos.
ADAuthenticationPolicySiloAccess	Revoke, Grant	Manage membership of authentication policy silos.
ADCentralAccessPolicy	Remove, Get, Set, New	Manage central access rules and policies.
ADCentralAccessPolicyMember	Remove, Add	Add and remove rules from a central access policy.
ADCentralAccessRule	New, Set, Remove, Get	Manage central access rules.
ADClaimTransformLink	Set, Clear, Remove	Manage claims transforms from being applied to one or more cross-forest trust relationships.
ADClaimTransformPol	New, Set, Get	Manage claim transformation policy objects from Active Directory.

icy		
ADClaimType	New, Get, Remove, Set	Manage Active Directory claim types.
ADComputer	Remove, New, Set, Get	Manage Active Directory computer accounts.
ADComputerServiceAccount	Remove, Add, Get	Add service accounts from Active Directory to a local computer.
ADDCCloneConfigFile	New	Generates a clone configuration file for a domain controller.
ADDCCloningExcludedApplicationList	Get	Manage which Active Directory applications are excluded when cloning the configuration of a domain controller.
ADDefaultDomainPasswordPolicy	Set, Get	Manage the default password policy for a domain.
ADDirectoryServer	Move	Use this cmdlet to move a DC to another AD site.
ADDirectoryServerOperationMasterRole	Move	Move an operations master role to another computer.

ADDomain	Set, Get	View and manage the properties of a domain.
ADDomainController	Get	View the properties of a domain controller.
ADDomainControllerPasswordReplicationPolicy	Remove, Get, Add	Manage which accounts can be replicated to an RODC.
ADDomainMode	Set	Set the domain functional level.
ADFineGrainedPasswordPolicy	Remove, Get, Set, New	Manage AD fine grained password policy.
ADFineGrainedPasswordPolicySubject	Get, Remove, Add	Manage the application of fine-grained password policies.
ADForest	Set, Get	Manage forest properties.
ADForestMode	Set	Configure forest functional level.
ADGroup	Get, Set, Remove, New	Manage AD groups.

ADGroupMember	Get, Remove, Add	Manage AD group membership.
ADObject	Get, Restore, Rename, Set, Move, Remove, Sync, New	Manage AD objects.
ADOptionalFeature	Disable, Get, Enable	Configure AD optional features.
ADOrganizationalUnit	Set, Get, New, Remove	Manage AD OUs.
ADPrincipalGroupMembership	Remove, Add, Get	Manage group membership on the basis of user account.
ADReplicationAttributeMetadata	Get	View replication metadata for AD object attributes.
ADReplicationConnection	Get, Set	Manage the properties of an AD replication connection.
ADReplicationFailure	Get	View information about AD replication failure.
ADReplicationPartner	Get	View information about AD replication partners.

Metadata		
ADReplicationQueueOperation	Get	View all operations in the AD replication queue.
ADReplicationSite	Set, Get, Remove, New	Manage AD replication sites.
ADReplicationSiteLink	Set, New, Get, Remove	Manage AD replication site links.
ADReplicationSiteLinkBridge	Get, Remove, New, Set	Manage AD replication site link bridges.
ADReplicationSubnet	New, Get, Set, Remove	Manage AD replication subnets.
ADReplicationUpToDatenessVectorTable	Get	Displays Update Sequence Numbers (USNs) for domain controllers.
ADResourceProperty	Set, New, Remove, Get	Manage Active Directory resource properties.
ADResourcePropertyList	Remove, Set, New, Get	Manage Active Directory resource property list.
ADResourc	Remove, Add	Add and remove resource properties

Technet24

ePropertyListMember		from an Active Directory resource property list.
ADResourcePropertyValueType	Get	View a resource property value type.
ADRootDSE	Get	View the root of a Directory Server information tree.
ADServiceAccount	Get, Test, Set, Install, New, Remove, Uninstall	Manage AD Managed Service Account.
ADServiceAccountPassword	Reset	Reset AD Managed Service Account password.
ADTrust	Get	View the properties of an AD Trust
ADUser	New, Set, Get, Remove	Manage an Active Directory user.
ADUserResultantPasswordPolicy	Get	Use this cmdlet to determine the resultant password policy for an account that has multiple fine-grained password policies applied to it.

Table 4-3 Group Policy module cmdlets

No	Verbs	Function

un		
GPInheritance	Get, Set	View and manage which GPOs are applied and whether inheritance is blocked.
GPLink	Remove, New, Set	Manage whether a GPO is linked.
GPO	Restore, Import, New, Remove, Rename, Backup, Get, Copy	Manage GPOs, including backup restore and copy.
GPOReport	Get	Generate a report on a GPO.
GPPermission	Set, Get	Manage permissions on policies.
GPPrefRe	Remove, Set, Get,	Manage registry-based policy preference settings. Microsoft maintains spreadsheets that map Group Policy settings to registry settings. To use this

gistryValue		cmdlet to set registry settings, you need to consult the spreadsheet.
GPRegistryValue	Remove, Get, Set,	Manage registry-based policy settings.
GPResultantSetOfPolicy	Get	View the resultant set of policy information.
GPStarterGPO	New, Get	Manage starter GPO.
GPUpdate	Invoke	Triggers a Group Policy update.

Table 4-4 AD Deployment module cmdlets

Noun	Verbs	Function
ADDSDomain	Install	Installs a new Active Directory Domain Services domain.
ADDSDomainController	Install, Uninstall	Use to add or remove a domain controller.
ADDSDomainControllerInstallation	Test	Runs a prerequisite check prior to installing a domain controller.
ADDSDomainControllerUninstallation	Test	Runs a prerequisite check prior to removing a domain controller.
ADDSDomainInstallation	Test	Checks the prerequisites for a new Active Directory Domain services domain.
ADDSForest	Install	Allows you to install a new Active Directory forest configuration.
ADDSForestInstallation	Test	Allows you to perform a prerequisite check prior to performing an Active Directory forest installation.
ADDSReadOnlyDomainControllerAccount	Add	Use this cmdlet to create a RODC account in the AD DS database.
ADDSReadOnlyDo	Test	Allows you to check that the necessary

Technet24

mainControllerAccountCreation	prerequisites are in place before you create an RODC account.

Active Directory module

As already mentioned, the *Active Directory PowerShell* module (see Table 4-2) is the one you're likely to use on a regular basis when managing Active Directory.

Group Policy module

As mentioned earlier, the *GroupPolicy* module (see Table 4-3) allows you to manage Group Policy from PowerShell.

ADDSDeployment module

As previously mentioned, you use the *ADDSDeployment* (see Table 4-4) module when performing deployment tasks.

Chapter 5

DNS, DHCP, and IPAM

DNS

DHCP

IPAM

DNS and DHCP servers are a part of all networks' core infrastructure. It doesn't matter if your organization has hosts running on Windows, Linux, macOS, or Internet of Things (IoT) devices because DHCP and DNS are likely managing the name and address information for those hosts. While new features have been introduced with each successive server operating system, the core functionality of Windows Server's DHCP and DNS server services hasn't changed much since the release of Windows 2000.

What *did* change with the release of Windows Server 2012, however, was the inclusion of the IPAM feature. IPAM allows you to manage all your DNS and DHCP

Technet24

infrastructure centrally. Instead of managing scopes and zones on individual DHCP and DNS servers, after you deploy IPAM, you can manage all your organization's zones and scopes from a single console.

DNS

At the most basic level, DNS servers translate host names to IP addresses and translate IP addresses to host names. By querying special records on DNS servers, it's possible to locate mail servers and name servers, verify domain ownership, and locate servers such as domain controllers. While DNS servers are usually deployed on a domain controller on a Windows Server network, it's also possible to deploy them on standalone computers.

DNS zone types

Zones store DNS resource record information. The DNS Server service in Windows Server 2019 supports several zone types, each of which is appropriate for a different set of circumstances. These zone types include primary, secondary, stub, and GlobalNames zones. You can integrate zones into Active Directory, or you can use the traditional primary or secondary architecture.

Active Directory–integrated zones

Active Directory–integrated zones can be replicated to all domain controllers in a domain, all domain

controllers in a forest, or all domain controllers enrolled in a specific Active Directory partition. You can create an Active Directory–integrated zone only on a writable domain controller, but you can also configure primary and stub zones as Active Directory–integrated zones. Domain controllers with DNS servers that host Active Directory–integrated zones can process updates to those zones.

When you choose to make a zone Active Directory–integrated, you have the option to configure a replication scope. You can configure the zone to be replicated, so that it is present on all domain controllers in the domain, in the forest, or within the scope of a custom Active Directory partition. When determining the appropriate replication scope, consider which clients need regular, direct access to the zone and which clients only require occasional, indirect access.

You can use custom directory partitions to replicate a zone to some, but not all, domain controllers. You can only select this option if there is an existing custom DNS application directory partition. You can use the *Add-DNSServerDirectoryPartition* cmdlet to create a directory partition. For example, to create a DNS Server directory partition called *Tasmania* on a subset of your organization's domain controllers, execute this command:

Technet24

Click here to view code image

```
Add-DNSServerDirectoryPartition –Name Tasmania
```

Active Directory–integrated zones use multimaster replication, meaning any domain controller that hosts the DNS zone can process updates to the zone. When creating a DNS zone, you must specify whether the zone supports dynamic updates, which allow clients to update DNS records. This is useful in environments in which clients change IP addresses on a regular basis. When a client gets a new IP address, it can update the record associated with its host name in the appropriate DNS zone. There are three dynamic update configuration options:

- **Allow Only Secure Dynamic Updates.** You can use this option only with Active Directory–integrated zones. Only authenticated clients can update DNS records.
- **Allow Both Non-Secure And Secure Dynamic Updates.** With this option, any client can update a record. Although this option is convenient, it is also insecure because any client can update the DNS zone, which could potentially redirect clients that trust the quality of the information stored on the DNS server.
- **Do Not Allow Dynamic Updates.** When you choose this option, all DNS updates must be performed manually. This option is very secure, but it is also labor-intensive.

An Active Directory–integrated zone can replicate to a read-only domain controller (RODC). With this configuration, the RODC-hosted zone is read only and the RODC cannot process updates to the zone, which is

also the case with a traditional writable domain controller. When you replicate an Active Directory–integrated zone to an RODC, the RODC forwards any client that wants to update the zone to a writable domain controller.

You can create an Active Directory–integrated primary zone by using the *Add-DnsServerPrimaryZone* cmdlet with the *ReplicationScope* parameter. For example, to create the Active Directory–integrated zone *cpandl.com* so that it replicates to all domain controllers in the forest, issue this command:

Click here to view code image

```
Add-DnsServerPrimaryZone –Name cpandl.com –ReplicationScope
Forest
```

When you first install Active Directory, the installation process ensures that the DNS zone associated with the root domain is automatically configured as an Active Directory–integrated zone. This root domain zone is automatically replicated to all domain controllers in the forest.

Primary and secondary zones

In traditional DNS implementations, a single server hosts a primary zone, which processes all zone updates, and a collection of secondary servers replicate zone data from the primary zone. One drawback to this model is

Technet24

that if the primary server fails, no zone updates can occur until the primary zone is restored.

Windows Server 2019 supports two types of primary zones: Active Directory–integrated zones and standard primary zones. Active Directory–integrated zones can only be hosted on computers that also function as domain controllers. Computers running Windows Server 2019 that do not function as domain controllers, however, can host standard primary zones. When you create a primary zone on a computer that is not a domain controller, the wizard does not enable you to specify a replication scope for the zone.

> **Inside OUT**
>
> **All zone types on a domain controller**
>
> The DNS server service on a domain controller supports all zone types. This means that you can choose to deploy a standard or Active Directory–integrated primary zone, a stub zone, a reverse lookup zone, or a secondary zone on a domain controller.

A secondary zone is a read-only copy of a primary zone. Secondary zones cannot process updates and can only retrieve updates from a primary zone. Also, secondary zones cannot be Active Directory—integrated zones, but you can configure a secondary zone for an Active Directory–integrated primary zone. Prior to configuring a secondary zone, you need to configure the primary zone that you want it to replicate from to enable

transfers to that zone. You can do this on the Zone Transfers tab of the Zone Properties dialog box. Secondary zones work best when the primary zone that they replicate from does not update frequently. If the primary zone does update frequently, it is possible that the secondary zone might have out-of-date records.

Use the *Add-DNSServerSecondaryZone* cmdlet to create a secondary zone. For example, to create a secondary zone for the *australia.contoso.com* zone, where a DNS server hosting the primary zone is located at *192.168.15.100*, issue the following command:

Click here to view code image

```
Add-DnsServerSecondaryZone -Name "australia.contoso.com" -
ZoneFile "australia.contoso.com.dns" -MasterServers
192.168.15.100
```

Reverse lookup zones

Reverse lookup zones translate IP addresses into FQDNs. You can create IPv4 or IPv6 reverse lookup zones, and you can also configure reverse lookup zones as Active Directory–integrated zones, standard primary zones, secondary zones, or stub zones. The domain controller promotion process automatically creates a reverse lookup zone based on the IP address of the first domain controller promoted in the organization.

Reverse lookup zones are dependent on the network ID of the IP address range that they represent. IPv4 reverse

lookup zones can only represent /8, /16, or /24 networks, which are the old Class A, Class B, and Class C networks. You can't create a single reverse lookup zone for IP subnets that doesn't fit into one of these categories, and the smallest reverse lookup zone you can create is for subnet mask /24, or 255.255.255.0.

Use the *Add-DNSPrimaryZone* cmdlet and the *NetworkID* parameter to create a reverse lookup zone. For example, to create an Active Directory–integrated reverse lookup zone for the *192.168.15.0/24* subnet, issue the following command:

Click here to view code image

```
Add-DnsServerPrimaryZone -NetworkID "192.168.15.0/24" -
ReplicationScope "Forest"
```

Inside OUT

Reverse lookup zones

Few applications actually require that you configure reverse lookup zones. In most organizations, the only reverse lookup zone is the one that is automatically created when Active Directory is installed. One of the few times where reverse lookup zones seem necessary is when you configure Simple Mail Transfer Protocol (SMTP) gateways. This is because some anti-SPAM checks perform a reverse IP address lookup to verify the identity of the SMTP gateway. Often, the SMTP gateway's IP address, being a public address, belongs to the ISP, which means creating the reverse lookup zone entry is beyond your direct control as a systems administrator.

Zone delegation

Zone delegations function as pointers to the next subordinate DNS layer in the DNS hierarchy. For

example, if your organization uses the *contoso.com* DNS zone and you want to create a separate *australia.contoso.com* DNS zone, you can perform a zone delegation so that the DNS servers for the contoso.com DNS zone point to the DNS servers for the *australia.contoso.com* DNS zone. When you create a new child domain in an Active Directory forest, zone delegation occurs automatically. When you are performing a manual delegation, create the delegated zone on the target DNS server prior to performing the delegation from the parent zone.

Although you can delegate several levels, remember that the maximum length of a FQDN is 255 bytes, and the maximum length of a FQDN for an Active Directory Domain Services domain controller is 155 bytes.

Use the *Add-DnsServerZoneDelegation* cmdlet to add a delegation. For example, to add a delegation for the *australia.contoso.com* zone to the *contoso.com* zone pointing at the name server *ausdns.australia.contoso.com*, which has the IP address *192.168.15.100*, use the following command:

Click here to view code image

```
Add-DnsServerZoneDelegation -Name "contoso.com" -ChildZoneName
"australia" -NameServer "ausdns.australia.contoso.com" -IPAddress
192.168.15.100 -PassThru
-Verbose
```

Forwarders and conditional forwarders

Forwarders and conditional forwarders enable your DNS server to forward traffic to specific DNS servers when a lookup request cannot be handled locally. For example, you might configure a conditional forwarder to forward all traffic for resource records in the *tailspintoys.com* zone to a DNS server at a specific IP address. If you don't configure a forwarder or if a configured forwarder can't be contacted, the DNS Server service forwards the request to a DNS root server and the request is resolved normally.

Forwarders

When you want to have a specific DNS server on the Internet handle your organization's DNS resolution traffic, you are likely to use a DNS forwarder rather than have your DNS server just use the root server. Most organizations configure their ISP's DNS server as a forwarder. When you do this, the ISP's DNS server performs all the query work, returning the result to your organization's DNS server. Your organization's DNS server then returns the result of the query back to the original requesting client.

You configure forwarders on a per-DNS server level. You can configure a forwarder using the DNS Manager. You can do this by editing the properties of the DNS server and then configuring the list of forwarders on the

Forwarders tab.

You can create a DNS forwarder by using the *Add-DnsServerForwarder* cmdlet. For example, to create a DNS forwarder for a DNS server with *10.10.10.111* as the IP address, issue this command:

Click here to view code image

```
Add-DnsServerForwarder 10.10.10.111
```

You can't create a forwarder on one DNS server and then have it replicate to all other DNS servers in the forest or the domain; however, this is possible with conditional forwarders and stub zones.

Conditional forwarders

Conditional forwarders only forward address requests from specific domains rather than forwarding all requests that can't be resolved by the DNS server. When configured, a conditional forwarder takes precedence over a forwarder. Conditional forwarders are useful when your organization has a trust relationship or partnership with another organization. You can configure a conditional forwarder that directs all traffic to host names within that organization instead of having to resolve those host names through the standard DNS-resolution process.

You can create conditional forwarders using the *Add-DnsServerConditionalForwarderZone* PowerShell

cmdlet. For example, to create a conditional forwarder for the DNS domain *tailspintoys.com* that forwards DNS queries to the server at IP address *10.10.10.102* and also replicates that conditional forwarder to all DNS servers within the Active Directory forest, issue this command:

Click here to view code image

```
Add-DnsServerConditionalForwarderZone –MasterServers 10.10.10.102
–Name tailspintoys.com –ReplicationScope Forest
```

Stub zones

A stub zone is a special zone that stores authoritative name server records for a target zone. Stub zones have an advantage over forwarders when the address of a target zone's authoritative DNS server changes on a regular basis. Stub zones are often used to host the records for authoritative DNS servers in delegated zones. Using stub zones in this way ensures that delegated zone information is up to date. If you create the stub zone on a writable domain controller, it can be stored with Active Directory and replicated to other domain controllers in the domain or forest.

You can add a stub zone using the *Add-DnsServerStubZone* cmdlet. For example, to add a DNS stub zone for the *fabrikam.com* zone using the DNS server at *10.10.10.222* and to also replicate that stub zone to all DNS servers in the forest, execute this command:

Click here to view code image

```
Add-DnsServerStubZone –MasterServers 10.10.10.222 –Name
fabrikam.com –ReplicationScope Forest –LoadExisting
```

GlobalNames zones

GlobalNames zones are a single-label, name-resolution replacement for WINS that can utilize existing DNS infrastructure. Deploying GlobalNames zones enables organizations to retire their existing WINS servers. You can use GlobalNames zones as long as your organization's DNS servers are running Windows Server 2008 or later.

Your organization should consider deploying GlobalNames zones instead of WINS in the following situations:

- Your organization is transitioning to IPv6. WINS does not support IPv6, and you need to support single-label name resolution.
- Single-label name resolution is limited to a small number of hosts that rarely change. GlobalNames zones must be updated manually.
- You have a large number of suffix search lists because of a complex naming strategy or disjoined namespace.

Entries in the GlobalNames zones must be populated manually. GlobalNames zone entries are alias (CNAME) records to existing DNS A or AAAA records. The existing DNS A and AAAA records can be dynamically updated, with these updates flowing on to records in the GlobalNames zone.

Technet24

To deploy a GlobalNames zone in a forest, perform the following steps:

1. On a domain controller that is configured as a DNS server, create a new Active Directory–integrated forward lookup zone and configure it to replicate to every domain controller in the forest using the New Zone Wizard.
2. On the Zone Name page, enter **GlobalNames** as the zone name. You can also accomplish the same task by running the following Windows PowerShell command:

   Click here to view code image

   ```
   Add-DnsServerPrimaryZone –Name GlobalNames –
   ReplicationScope Fores
   ```

3. Next, activate the GlobalNames zone on each authoritative DNS server hosted on a domain controller in the forest. To do this, execute the following Windows PowerShell command. (where *DNSServerName* is the name of the domain controller hosting DNS):

   Click here to view code image

   ```
   Set-DnsServerGlobalNameZone –ComputerName
   DNSServerName –Enable $True
   ```

4. To populate the GlobalNames zone, create alias (CNAME) records in the GlobalNames zone that point to A or AAAA records in existing zones.

Peer Name Resolution Protocol

Peer Name Resolution Protocol (PNRP) provides IPv6 with a peer-to-peer name resolution method. Devices that are connected to the Internet and are assigned an IPv6 address can publish their peer name/address combination and provide it to peers. These name

combinations include both FQDN and single-label names. Other devices query peers to learn IPv6 addressing information. For example, imagine that there are computers named *Sydney*, *Melbourne*, *Canberra*, *Adelaide*, *Brisbane*, *Hobart*, and *Perth*. Computer *Sydney* needs to know the IPv6 address of *Perth*. *Sydney* is near *Canberra*, *Melbourne*, and *Brisbane*. *Sydney* queries *Canberra*, *Melbourne*, and *Brisbane*, one after the other, to determine whether any of them knows *Perth's* IPv6 address. Rather than using a central server, such as DNS or WINS, PNRP uses the information that each individual computer knows about all the other computers to determine address information. PNRP has the following properties:

- Does not require a centralized infrastructure. Servers are required only for bootstrapping.
- Can scale to billions of names and is fault tolerant. Multiple computers can host copies of the same PNRP record.
- Names are updated in real time and PNRP is designed to not return stale addresses.
- Can be used to name services rather than just computers.
- Names can be published in a secure or insecure manner. When published in a secure manner, PNRP uses public key cryptography to validate records.
- Installed as a feature.

PNRP peer groups are termed *clouds*. When installed, PNRP can use the following clouds:

- **Global cloud.** Uses the global IPv6 address scope and represents all computers and devices connected to the Internet. There is a single global cloud.
- **Link-local cloud.** Uses all addresses in the link-local IPv6 address scope. It roughly corresponds to the local subnet. An organization can have multiple link-local clouds.

Resource records

DNS supports a large number of resource records. The most basic resource record maps an FQDN to an IP address. More complex resource records provide information about the location of services, such as SMTP servers and domain controllers.

Host records

Host records are the most common form of record and can be used to map FQDNs to IP addresses. There are two types of host records:

- The A record, which is used to map FQDNs to IPv4 addresses
- AAAA records, which are used to map FQDNs to IPv6 addresses

You can add a new host record to a zone by right-clicking the zone in DNS Manager and then clicking New Host (A or AAAA). You also have the option of creating a pointer (PTR) record in the appropriate reverse lookup zone, if one exists. You can add host records with the *Add-DnsServerResourceRecordA* cmdlet and AAAA records with the *Add-DnsServerResourceRecordAAAA* cmdlet.

Alias (CNAME)

An alias—or CNAME—record enables you to provide an alternate name when there is an existing host record. You can create as many aliases for a particular record as you need. To create a new alias in a zone, right-click the zone in DNS Manager and then click New Alias (CNAME). When you create an alias, you must point the alias to an existing host record. You can use the Browse button to navigate to the target host record or enter it manually. You can also add an alias record to a zone from Windows PowerShell by using the *Add-DnsServerResourceRecordCName* cmdlet.

Mail exchanger

Mail exchanger (MX) records are used to locate mail gateways. For example, when a remote mail gateway wants to forward an email message to an email address associated with your organization's DNS zone, it performs an MX lookup to determine the location of the mail gateway. After that determination has been made, the remote mail gateway contacts the local gateway and transmits the message. MX records must map to existing host records. To create an MX record, right-click the zone in DNS Manager, click New Mail Exchanger (MX), and enter information in the New Resource Record dialog box. The Mail Server Priority field is available to allow for the existence of more than one MX

record in a zone and is often used when organizations have multiple mail gateways. This is done so that if an organization's primary mail gateway fails, remote mail servers can forward message traffic to other mail gateways. You can add MX records using the *Add-DnsServerResourceRecordMX* PowerShell cmdlet.

Pointer record

Pointer (PTR) records enable you to connect IP addresses to FQDNs and are hosted in reverse lookup zones. When you create a host record, a PTR record is automatically created by default if an appropriate reverse lookup zone exists. To create a PTR record, right-click the reverse lookup zone in DNS Manager, click New Pointer (PTR), and enter the PTR record information in the New Resource Record dialog box. You can create a PTR record from Windows PowerShell by using the *Add-DnsServerResourceRecordPtr* cmdlet.

Inside OUT

Unknown Records

Windows Server 2019 DNS supports "unknown records," which are resource records whose RDATAformat is unknown to the DNS server. "Unknown records" are defined in RFC 3597 and can be added to Windows Server DNS zones when they are hosted on computers running the Windows Server 2016 operating system.

Zone aging and scavenging

Aging and scavenging provide a technique to reduce the

incidence of stale resource records in a primary DNS zone. Stale records are records that are out of date or no longer relevant. If your organization has zones that relate to users with portable computers, such as laptops and tablets, those zones might end up accumulating stale resource records. This can lead to the following problems:

- DNS queries return stale rather than relevant results.
- Large zones can cause DNS server performance problems.
- Stale records might prevent DNS names from being reassigned to different devices.

To resolve these problems, you can configure the DNS Server service to do the following:

- Time stamp resource records that are dynamically added to primary zones. This occurs when you enable aging and scavenging.
- Age resource records based on a refresh time period.
- Scavenge resource records that are still present beyond the refresh period.

After you configure the DNS Server service, aging and scavenging occurs automatically. It is also possible to trigger scavenging by right-clicking the DNS server in DNS Manager and then clicking Scavenge Stale Resource Records. You can also configure aging and scavenging using the *Set-DnsServerScavenging* cmdlet. For example, to enable stale resource record scavenging on all zones of a DNS server and to set the Scavenging and Refresh Intervals to 10 days, issue this command:

Click here to view code image

```
Set-DnsServerScavenging -ApplyOnAllZones -RefreshInterval
10.0:0:0 -ScavengingInterval 10.0:0:0 -ScavengingState $True
```

DNSSEC

Domain Name System Security Extensions (DNSSEC) add security to DNS by enabling DNS servers to validate the responses given by other DNS servers. DNSSEC enables digital signatures to be used with DNS zones. When the DNS resolver issues a query for a record in a signed zone, the authoritative DNS server provides both the record and a digital signature, enabling validation of that record.

To configure DNSSEC, perform the following steps:

1. Right-click the zone in DNS manager, click DNSSEC, and then click Sign The Zone.
2. On the Signing Options page, select Use Default Settings To Sign The Zone.

When you configure DNSSEC, three new resource records are available. These records have the following properties:

- **Resource Record Signature (RRSIG) record.** This record is stored within the zone, and each is associated with a different zone record. When the DNS server is queried for a zone record, it returns the record and the associated RRSIG record.
- **DNSKEY.** This is a public key resource record that enables the validation of RRSIG records.

- **Next Secure (NSEC/NSEC3) record.** This record is used as proof that a record does not exist. For example, the *contoso.com* zone is configured with DNSSEC. A client issues a query for the host record *tasmania.contoso.com* and if there is no record for *tasmania.contoso.com*, the NSEC record returns, which informs the host making the query that no such record exists.

In addition to the special resource records, a DNSSEC implementation has the following components:

- **Trust anchor.** This is a special public key associated with a zone. Trust anchors enable a DNS server to validate DNSKEY resource records. If you deploy DNSSEC on a DNS server hosted on a domain controller, the trust anchors can be stored in the Active Directory forest directory partition. This replicates the trust anchor to all DNS servers hosted on domain controllers in the forest.
- **DNSSEC Key Master.** This is a special DNS server that you use to generate and manage signing keys for a DNSSEC-protected zone. Any computer running Windows Server 2012 or later that hosts a primary zone, whether standard or integrated, can function as a DNSSEC Key Master. A single computer can function as a DNSSEC Key Master for multiple zones. The DNSSEC Key Master role can also be transferred to another DNS server that hosts the primary zone.
- **Key Signing Key (KSK).** You use the KSK to sign all DNSKEY records at the zone root. You create the KSK by using the DNSSEC Key Master.
- **Zone Signing Key (ZSK).** You use the ZSK to sign zone data, such as individual records hosted in the zone. You create the ZSK by using the DNSSEC Key Master.

You can configure group policy to ensure that clients only accept records from a DNS server for a specific zone if those records have been signed using DNSSEC.

You do this by configuring the Name Resolution Policy Table, which is located in the *Computer Configuration\Policies\Windows Settings\Name Resolution Policy* node of a GPO. You create entries in the table—for example, requiring that all queries against a specific zone require DNSSEC validation. You can configure the NRPT by using Group Policy or through Windows PowerShell.

DNSSEC is appropriate for high-security environments, such as those where IPSec and authenticating switches are in use. DNSSEC protects against attacks where clients are fed false DNS information. In many small- to medium-sized environments, the likelihood of such an attack is minimal; however, in high-security environments, enabling DNSSEC is a prudent precaution.

DNS event logs

The DNS server log is located in the Applications And Services Logs folder in Event Viewer. Depending on how you configure event logging on the Event Logging tab of the DNS server's properties, this event log can record information including:

- Changes to the DNS service, such as when the DNS server service is stopped or started.
- Zone loading and signing events.
- Modifications to the DNS server configuration.

- DNS warning and error events.

By default, the DNS server records all these events. It's also possible to configure the DNS server to only log errors, or you can have it log errors and warning events. The key with any type of logging is that you should only enable logging for information that you might need to review at some point in time. Many administrators log everything "just in case," even though they are only ever interested in a specific type of event.

In the event that you need to debug how a DNS server performs, you can enable Debug Logging on the Debug Logging tab of the DNS server's Properties dialog box. Debug logging is resource intensive, and you should only use it when you have a specific problem related to the functionality of the DNS server. You can configure debug logging to use a filter, so that only traffic from specific hosts is recorded, instead of recording traffic from all hosts that interact with the DNS server.

DNS options

In high-security environments, there are a number of steps that you can take to make a DNS server more secure from attackers who attempt to spoof it. Spoofing can cause the server to provide records that redirect clients to malicious sites. While DNSSEC provides security for zones hosted on the server, most DNS server traffic involves retrieving information from

remote DNS servers and passing that information on to clients. In this section, you learn about settings that you can configure to ensure that the information relayed to clients retains its integrity in the event that a nefarious third party attempts to spoof your organization's DNS servers.

DNS socket pool

DNS socket pool is a technology that makes cache-tampering and spoofing attacks more difficult by using source port randomization when issuing DNS queries to remote DNS servers. To spoof the DNS server with an incorrect record, the attacker needs to guess which randomized port was used as well as the randomized transaction ID that was issued with the query. By default, the DNS server running on Windows Server 2012 or later uses a socket pool of 2,500. You can use the *dnscmd* command-line tool to change the socket pool size to a value between *0* and *10,000*. For example, to set the socket pool size to *4,000*, issue the following command:

Click here to view code image

```
Dnscmd /config /socketpoolsize 4000
```

You must restart the DNS service before you can use the reconfigured socket pool size.

DNS cache locking

DNS cache locking provides you with control over when information stored in the DNS server's cache can be overwritten. For example, when a recursive DNS server responds to a query for a record that is hosted on another DNS server, it caches the results of that query so that it doesn't have to contact the remote DNS server if the same record is queried again within the Time to Live (TTL) value of the resource record. DNS cache locking prevents record data in a DNS server's cache from being overwritten until a configured percentage of the TTL value has expired. By default, the cache-locking value is set to 100 percent, but you can reset it by using the *Set-DNSServerCache* cmdlet with the *LockingPercent* option. For example, to set the cache-locking value to 80 percent, issue the following command and then restart the DNS server service:

Click here to view code image

```
Set-DNSServerCache –LockingPercent 80
```

DNS recursion

By default, DNS servers on Windows Server 2012 and later perform recursive queries on behalf of clients. This means when the client asks the DNS server to find a record that isn't stored in a zone hosted by the DNS server, the DNS server then goes out and finds the result of that query and passes it back to the client. It's possible however, for nefarious third parties to use

recursion as a denial-of-service attack vector, which could slow a DNS server to the point where it becomes unresponsive. You can disable recursion on the Advanced Tab of the DNS server's properties, or you can configure it by using the *Set-DNSServerRecursion* cmdlet. For example, to configure the recursion retry interval to *3* seconds, use the following command:

Click here to view code image

```
Set-DNSServerRecursion -RetryInterval 3 -PassThru
```

Netmask ordering

Netmask ordering ensures that the DNS server returns the host record on the requesting client's subnet, if such a record exists. For example, imagine that the following host records exist on a network that uses 24-bit subnet masks:

10.10.10.105 wsus.contoso.com

10.10.20.105 wsus.contoso.com

10.10.30.105 wsus.contoso.com

If netmask ordering is enabled and a client with the IP address 10.10.20.50 performs a lookup of *wsus.contoso.com*, the lookup always returns the record 10.10.20.105 because this record is on the same subnet as the client. If netmask ordering is not enabled, the DNS server returns records in a round-robin fashion. If

the requesting client is not on the same network as any of the host records, the DNS server also returns records in a round-robin fashion. Netmask ordering is useful for services, such as WSUS, that you might have at each branch office. When you use netmask ordering, the DNS server redirects the client in the branch office to a resource on the local subnet when one exists.

Netmask ordering is enabled by default on Windows Server 2012 and later DNS servers. You can verify that netmask ordering is enabled by viewing the Advanced Properties of the DNS server.

Response rate limiting

Distributed denial-of-service attacks against DNS servers are becoming increasingly common. One method of ameliorating such attacks is by configuring response rate limiting; this setting determines how a DNS server responds to clients sending an unusually high number of DNS queries. You can configure the following response rate limiting settings:

- **Responses Per Second.** Determines the maximum number of times an identical response can be returned to a client per second.
- **Errors Per Second.** Determines the maximum number of error responses returned to a client per second.
- **Window.** The timeout period for any client that exceeds the maximum request threshold.
- **Leak Rate.** Determines how often the DNS server responds to queries if a client is in the suspension window. The leak rate is the

number of queries it takes before a response is sent. A leak rate of 42 means that the DNS server only responds to one query out of every 42 when a client is in the suspension window period.

- **TC Rate.** Tells the client to try connecting with TCP when the client is in the suspension window. The TC rate should be below the leak rate to give the client the option of attempting a TCP connection before a leak response is sent.
- **Maximum Responses.** The maximum number of responses that a DNS server issues to a client during the timeout period.
- **White List Domains.** Domains that are excluded from response rate limiting settings.
- **White List Subnets.** Subnets that are excluded from response rate limiting settings.
- **White List Server Interfaces.** DNS server interfaces that are excluded from response rate limiting settings.

DANE

DNS-based Authentication of Named Entities (DANE) allows you to configure records that use Transport Layer Security Authentication (TLSA) to inform DNS clients of which certificate authorities (CA) they should accept certificates from for your organization's domain name DNS zone. DANE blocks attackers from using attacks where a certificate is issued from a rogue CA and is used to provide validation for a rogue website when a DNS server is compromised. When DANE is implemented, a client requesting a TLS connection to a website, such as *www.tailspintoys.com*, knows that it should only accept a TLS certificate from a specific CA. If the TLS certificate

for *www.tailwindtraders.com* is from a different CA, the client knows that the website that appears to be at *www.tailwindtraders.com* might have been compromised or that the client might have been redirected by an attacker corrupting a DNS cache.

DNS Policies

You use DNS policies to control how DNS servers respond to queries based on the properties of that query. You can configure DNS policies to achieve the following:

- **Application load balancing.** Allows you to weight which endpoint is returned when multiple endpoints exist for an address. For example, returning one IP address 75 percent of the time and another IP address 25 percent of the time. This differs from DNS round robin where multiple IP addresses associated with a specific addresse would be returned to clients on an equal basis (netmask ordering being the exception).
- **Traffic management.** The DNS server provides a response to the client that directs them to the datacenter closest to them. For example, clients making DNS requests from Australia are redirected to an application endpoint in Australia, and clients making DNS requests from Canada are redirected to an application endpoint in Canada.
- **Split Brain DNS.** DNS records can be segmented into zone scopes. Records can be placed in different scopes. This allows external clients to query the DNS zone and be only able to access a specific number of DNS records that are members of a scope designed for external clients. Internal clients have access to a more comprehensive internal scope.
- **Filtering.** Allows you to block DNS queries from specific IP addresses, IP address ranges, or FQDNs.

Technet24

- **Forensics.** Redirects malicious clients to a sinkhole host rather than to the host they are trying to reach. For example, if clients appear to be attempting a denial-of-service attack, redirect them to a static page rather than the site they are attempting to attack.
- **Time of day–based redirection.** Allows you to configure a policy where clients are redirected to specific hosts based on the time of day when the query is received.

To accomplish most of these tasks, you create a zone scope, which is a collection of records within a zone. For example, in a split-brain scenario, you'd create a zone scope for the internal records. You'd then use the *Add-DNSServerQueryResolutionPolicy* cmdlet to create a policy that allowed all queries from the DNS server's internal network interface to have access to records in the internal zone scope; queries that came through any other interface on the DNS server would only have access to the Default scope, which would only be composed of those records in the zone to which you wanted public clients to have access. Similarly, in the time of day–based redirection scenario, you'd create scopes for each period of time and then create policies that were triggered based on the time of day that would direct clients to records within the particular scope associated with that period of time.

Delegated administration

In some larger environments, you might want to separate administrative privileges, so that the people

responsible for managing your organization's DNS servers don't have other permissions, such as the ability to create user accounts or reset passwords. By default, members of the Domain Admins group can perform all DNS administration tasks on DNS servers within a domain. Members of the Enterprise Admins group can perform all DNS administration tasks on any DNS server in the forest.

You can use the DNSAdmins domain local group to grant users the ability to view and modify DNS data, as well as view and modify server configurations of DNS servers within a domain. Add users to this group when you want to allow them to perform DNS-administration tasks without giving them additional permissions. You can assign permissions that allow users or security groups to manage a specific DNS server on the Security tab of the server's Properties dialog box.

You can also configure permissions at the zone level. To do this, assign security principal permissions on the Security tab of the zone's Properties dialog box. You might do this when you want to allow a specific person to manage host records without assigning them any other permissions. Today, most organizations allow DNS records to be updated dynamically. This means that the only zones where you might need to configure special permissions to allow manual management are ones for which security is especially important, such as

those that are accessible to clients on the Internet.

Managing DNS with PowerShell

You can manage almost all aspects of DNS server operations using the DNSServer module of Windows PowerShell. Table 5-1 lists the DNS server PowerShell cmdlets.

Table 5-1 DNS PowerShell cmdlets

Noun	Verbs	Function
DnsServer	Set, Get, Test	View and manage DNS server configuration
DnsServerCache	Get, Set, Show, Clear	Manage DHS server cache configuration
DnsServerClientSubnet	Set, Get, Add, Remove	Manage DHS server client subnets for use in DNS policies
DnsServerConditionalForwarderZone	Set, Add	Manage conditional forwarders
DnsServerDiagnostics	Set, Get	View DNS event log information
DnsServerDirectoryPartition	Unregister, Remove, Get, Register, Add	Manage DNS server directory partitions

DnsServerDnsSecPublicKey	Export	Export the DS and DNSKEY data for a DNSSEC signed zone
DnsServerDnsSecZoneSetting	Get, Set, Test	Configure DNSSEC settings for a zone
DnsServerDsSetting	Set, Get	Configure Active Directory server settings
DnsServerEDns	Get, Set	Manage CacheTimeout, EnableProbes, and EnableReception settings
DnsServerForwarder	Add, Set, Get, Remove	Manage DNS server forwarders
DnsServerGlobalNameZone	Set, Get	Manage the global names zone
DnsServerGlobalQueryBlockList	Get, Set	Manage the global query block list
DnsServerKeyStorageProvider	Show	View key storage provider information
DnsServerPolicy	Disable, Enable	Enable or disable DNS server policies
DnsServerPrimaryZone	Restore, Set, Add, CovertTo	Manage DNS server primary zones

DnsServerQueryResolutionPolicy	Add,Get,Remove, Set	Configure DNS server query resolution policy
DnsServerRecursion	Set, Get	Manage DNS server recursion settings
DnsServerRecursionScope	Remove, Set, Get, Add	Configure DNS server recursion scopes
DnsServerResourceRecord	Add, Set, Remove, Get	Manage DNS server resource records
DnsServerResourceRecordA	Add	Add an IPv4 host record
DnsServerResourceRecordAAAA	Add	Add an IPv6 host record
DnsServerResourceRecordAging	Set	Configure resource record aging
DnsServerResourceRecordCName	Add	Add a CNAME record
DnsServerResourceRecordDnsKey	Add	Add a DNSKey resource record
DnsServerResourceRecordDS	Import, Add	Manage delegation signer (DS) resource records
DnsServerResourceRecordMX	Add	Add a mail exchanger (MX) resource record

DnsServerResourceRecordPtr	Add	Add a pointer (PTR) resource record
DnsServerResponseRateLimiting	Get, Set	Configure response rate limiting
DnsServerResponseRateLimitingExceptionlist	Add, Remove, Get, Set	Configure response rate limit exceptions
DnsServerRootHint	Remove, Get, Add, Set, Import	Manage DNS server root hints
DnsServerScavenging	Start, Set, Get	Manage DNS server scavenging settings
DnsServerSecondaryZone	Add, Restore, ConvertTo, Set	Manage DNS server secondary zones
DnsServerSetting	Get, Set	Configure DNS server settings
DnsServerSigningKey	Get, Add, Set, Remove	Manage DNS Server signing keys
DnsServerSigningKeyRollover	Invoke, Enable, Step, Disable	Configure signing key rollover
DnsServerStatistics	Clear, Get	Manage DNS server statistics
DnsServerStubZo	Set, Add	Configure stub zones

ne		
DnsServerTrustAnchor	Get, Import, Remove, Add	Configure DNSSEC trust anchor
DnsServerTrustPoint	Get, Update	Configure DNSSET trust point
DnsServerVirtualizationInstance	Remove, Get, Set, Add	Manage virtualization instances
DnsServerZone	Get, Export, Sync, Suspend, Remove, Resume	Manage zones, including export and import
DnsServerZoneAging	Set, Get	Configure zone record aging
DnsServerZoneDelegation	Remove, Add, Get, Set	Configure zone delegation
DnsServerZoneKeyMasterRole	Reset	Transfers key master role for a DNS zone
DnsServerZoneScope	Remove, Get, Add	Manage zone scopes
DnsServerZoneSign	Invoke	Sign a DNS zone (DNSSEC)
DnsServerZoneTransfer	Start	Manage zone transfers

DnsServerZoneTransferPolicy	Add, Get, Set, Remove	Configure zone transfer policies
DnsServerZoneUnsign	Invoke,	Unsign a DNS zone (DNSSEC)

DHCP

DHCP is a network service that most administrators barely pay attention to after they've configured it. The main concern that most administrators have with DHCP is up until the release of Windows Server 2012, it had been difficult to configure as a highly available service. While DNS became highly available through its ability to be hosted on any domain controller, the problem with making DHCP highly available was ensuring that when multiple DHCP servers were in play, that duplicate addresses were not being assigned to separate clients. Windows Server 2012 introduced highly available DHCP servers, allowing two servers to share a scope, rather than having to split the scope according to the old 80/20 rule.

Scopes

A DHCP scope is a collection of IP address settings that a client uses to determine its IP address configuration. You configure a DHCP scope for every separate IPv4 subnet to which you want the DHCP servers to provide

IP address configuration information. When configuring an IPv4 scope, specify the following information:

- **Scope Name.** The name of the scope. The name should be descriptive enough that you recognize which hosts the scope applies to. For example, Level 2, Old Arts Building.
- **IP Address Range.** This is the range of IP addresses encompassed by the scope, and it should be a logical subnet.
- **IP Address Exclusions.** This includes which IP addresses within the IP address range you do not want the DHCP server to lease. For example, you could configure exclusions for several computers that have statically assigned IP addresses.
- **Delay.** This setting determines how long the DHCP server waits before offering an address. This provides you with a method of DHCP server redundancy, with one DHCP server offering addresses immediately and another only offering addresses after a certain period of time has elapsed.
- **Lease Duration.** The lease duration determines how long a client has an IP address before the address returns to the lease pool.
- **DHCP Options.** There are 61 standard options that you can configure. These include settings, such as DNS server address and default gateway address, that apply to leases in the scope.

You can add a DHCP server scope using the *Add-DHCPServer4Scope* or *Add-DHCPServer6Scope* PowerShell cmdlets. For example, to create a new scope for the *192.168.10.0/24* network named *MelbourneNet*, enact the following PowerShell command:

Click here to view code image

```
Add-DHCPServerv4Scope -Name "MelbourneNet" -StartRange
192.168.10.1 -Endrange 192.168.10.254 -SubnetMask 255.255.255.0
```

Inside OUT

IPv6 DHCP

Most organizations are still using IPv4 with their DHCP server. For the sake of brevity, all PowerShell examples are for IPv4, with the cmdlets for managing IPv6 DHCP scopes listed later in the chapter.

Server and scope options

DHCP options include additional configuration settings, such as the address of a network's default gateway or the addresses of DNS servers. You can configure DHCP options at two levels: the scope level and the server level. Options configured at the scope level override options configured at the server level.

Common scope options include:

- **003 Router.** The address of the default gateway for the subnet.
- **006 DNS Servers.** The address of DNS servers that the client should use for name resolution.
- **015 DNS Domain Name.** This is the DNS suffix assigned to the client.
- **058 Renewal time.** This determines how long after the lease is initially acquired that the client waits before it attempts to renew the lease.

You can configure IPv4 scope options using the *Set-DHCPServer4OptionValue* cmdlet. For example, to configure the scope with ID *192.168.10.0* on DHCP server *dhcp.contoso.com* with the value for the DNS server set to *192.168.0.100*, issue the following

command:

Click here to view code image

```
Set-DhcpServerv4OptionValue -Computername dhcp.contoso.com -
ScopeID 192.168.10.0
-DNSServer 192.168.0.100
```

Reservations

Reservations allow you to ensure that a particular computer always receives a specific IP address. You can use reservations to allow servers to always have the same address, even when they are configured to retrieve that address through DHCP. If you don't configure a reservation for a computer, it can be assigned any available address from the pool. Although you can configure DHCP to update DNS to ensure that other hosts can connect using the client's hostname, it is generally a good idea to ensure that a server retains the same IP address.

There are two ways to configure a DHCP reservation. The simplest way is to locate an existing IP address lease in the DHCP console, right-click it, and then select Add To Reservation. This then ties that specific network adapter address to that IP address. The only drawback of this method is that the reservation is for the assigned IP address, and you cannot customize it.

If you know the network adapter address of the computer for which you want to configure a reservation,

you can right-click the Reservations node under the scope that you want to configure the reservation under, and then click New Reservation. In the New Reservation dialog box, enter the Reservation Name, IP Address, and Network Adapter Address, and then click Add to create the reservation. You can also configure a reservation using the *Add-DhcpServerv4Reservation* cmdlet. For example, to configure a reservation, issue the following command:

Click here to view code image

```
Add-DhcpServerv4Reservation -ScopeId 192.168.10.0 -IPAddress
192.168.10.8 -ClientId A0-FE-C1-7A-00-E5 -Description "A4
Printer"
```

DHCP filtering

You can configure DHCP filtering on the Filters tab of the IPv4 Properties dialog box in the DHCP console. DHCP filtering is used in high-security environments to restrict the network adapter addresses that can utilize DHCP. For example, you can have a list of all network adapter addresses in your organization and use the Enable Allow List option to lease only IP addresses to this list of authorized adapters. This method of security is not entirely effective because it is possible for unauthorized people to fake an authorized adapter address as a way of gaining network access. You can add a specific network adapter to the allowed list of MAC

addresses using the *Add-DHCPServerv4Filter* cmdlet. For example, to allow a laptop with the MAC address *D1-F2-7C-00-5E-F0* to the authorized list of adapters, issue the command:

Click here to view code image

```
Add-DhcpServerv4Filter -List Allow -MacAddress D1-F2-7C-00-5E-F0
Description
"Rooslan's Laptop"
```

Superscopes

A superscope is a collection of individual DHCP scopes. You might create a superscope when you want to bind existing scopes for administrative reasons. For example, imagine that you have a subnet in a building that is close to fully allocated. You can add a second subnet to the building and then bind them into a superscope. The process of binding several separate logical subnets on the same physical network is known as multinetting.

There needs to be at least one existing scope present on the DHCP server before you can create a superscope. After you have created a superscope, you can add new subnets or remove subnets from that scope. It's also possible to deactivate some subnets within a scope while keeping others active. You might use this technique when migrating clients from one IP address range to another. For example, you could have both the source and destination scopes as part of the same

superscope and activate the new scope and deactivate the original scope as necessary when performing the migration. You can create superscopes using the *Add-DHCPServerv4Superscope* cmdlet. For example, to add the scopes with ID *192.168.10.0* and *192.168.11.0* to a superscope named *MelbourneCBD*, run the following command:

Click here to view code image

```
Add-DhcpServerv4Superscope -SuperscopeName "MelbourneCBD" -
ScopeId 192.168.10.0, 192.168.11.0
```

Multicast scopes

A multicast address is an address that allows one-to-many communications on a network. When you use multicast, multiple hosts on a network listen for traffic on a single multicast IP address. Multicast addresses are in the IPv4 range of 224.0.0.0 through to 239.255.255.255. Multicast scopes are collections of multicast addresses. You can configure a Windows Server 2016 DHCP server to host multicast scopes. Multicast scopes are also known as Multicast Address Dynamic Client Allocation Protocol (MADCAP) scopes because applications that require access to multicast addresses support the MADCAP API.

The most common use for multicast addresses using the default Windows Server 2019 roles is for Windows Deployment Services. You can, however, configure the

WDS server with its own set of multicast addresses, and you don't need to configure a special multicast scope in DHCP to support this role.

Split scopes

A split scope is one method of providing fault tolerance for a DHCP scope. The idea behind a split scope is that you host one part of the scope on one DHCP server and a second, smaller part of the scope on a second DHCP server. Usually, this split has 80 percent of the addresses on the first DHCP server and 20 percent of the addresses on the second server. In this scenario, the DHCP server that hosts the 20-percent portion of the address space is usually located on a remote subnet. In the split-scope scenario, you also use a DHCP relay agent configured with a delay so that most addresses are leased from the DHCP server that hosts 80 percent of the address space. Split scopes are most likely to be used in scenarios where your DHCP servers aren't running on Windows Server 2012 or later operating systems. If you want to provide fault tolerance for scopes hosted on servers running Windows Server 2019, you should instead implement DHCP failover.

Name protection

DHCP name protection is a feature that allows you to ensure that the host names a DHCP server registers

with a DNS server are not overwritten in the event that a non-Windows operating system has the same name. DHCP name protection also protects names from being overwritten by hosts that use static addresses that conflict with DHCP-assigned addresses.

For example, imagine that in the *tailwindtraders.com* domain there is a computer running the Windows 10 operating system that has the name *Auckland* and receives its IP address information from a Windows Server 2019 DHCP server. The DHCP server registers this name in DNS and creates a record in the *contoso.com* DNS zone that associates the name *Auckland.tailwindtraders.com* with the IP address assigned to the computer running Windows 10. A newly installed computer running a distribution of Linux is also assigned the name *Auckland*. Because name protection has been enabled, this new computer is unable to overwrite the existing record with a new record that associates the name *Auckland.tailwindtraders.com* with the Linux computer's IP address. If name protection had not been enabled, it is possible that the record would have been overwritten.

You can enable name protection on a scope by clicking Configure on the DNS tab of the IPv4 or IPv6 Properties dialog box. You can also do this by using the *Set-DhcpServerv4DnsSetting* or the *Set-*

Technet24

DhcpServerv6DnsSetting cmdlet. For example, to configure the DHCP server on computer *MEL-DC* so that name protection is enabled on all IPv4 scopes, issue the following command:

Click here to view code image

```
Set-DhcpServerv4DnsSetting –Computer MEL-DC –NameProtection $true
```

DHCP failover

DHCP failover allows you to configure DHCP to be highly available without using split scopes. You have two options when configuring DHCP failover:

- **Hot Standby mode.** This relationship is a traditional failover relationship. When you configure this relationship, the primary server handles all DHCP traffic unless it becomes unavailable. You can configure DHCP servers to be in multiple separate relationships, so it's possible that a DHCP server can be the primary server in one relationship and a hot standby server in another relationship. When configuring this relationship, you specify a percentage of the address ranges to be reserved on the standby server. The default value is 5 percent. These 5 percent of addresses are instantly available after the primary server becomes unavailable. The hot standby server takes control of the entire address range when the figure specified by the state switchover interval is reached. The default value for this interval is 60 minutes.
- **Load Sharing mode.** This is the default mode when you create a DHCP failover relationship. In this mode, both servers provide IP addresses to clients according to the ratio defined by the load balance percentage figure. The default is for each server to share 50 percent of the load. The maximum client lead time figure enables one of the partners to take control of the entire scope in the event that the other partner is not available for the specified time.

Prior to configuring DHCP failover, you need to remove any split scopes between the potential partners. You can also choose a shared secret to authenticate replication traffic, but you don't have to enter this secret on the partner DHCP server.

Administration

DHCP servers need to be authorized in Active Directory before they start leasing addresses. DHCP servers on Microsoft operating systems must always be authorized before they start working, simply because this reduces the chance that a rogue DHCP server might start leasing IP addresses. Only the following can authorize DHCP servers:

- Members of the Enterprise Admins group
- Members of the Domain Admins group in the forest root domain
- Accounts delegated the ability to add, modify, or delete new objects of the DHCP class type on the NetServices folder in Active Directory

When you deploy the DHCP server role, the following local groups are created. You can use membership of these groups to delegate DHCP server–management permissions to users or groups:

- **DHCP Administrators.** Members of this group can view or modify any DHCP server settings, including configuring DHCP options, failover, and scopes.
- **DHCP Users.** Members of this group only have read-only access to the DHCP server. They can view, but not modify, DHCP server settings, including scope information and data.

Managing DHCP with PowerShell

You can manage the DHCP role by using the PowerShell cmdlets listed in Table 5-2.

Table 5-2 DHCP PowerShell cmdlets

Noun	Verbs	Function
DhcpServer	Restore, Export, Backup, Import	Backup, recover, export, and import DHCP server configuration
DhcpServerAuditLog	Get, Set	Manage the DHCP server audit log
DhcpServerDatabase	Get, Set	Manage the DHCP server database
DhcpServerDnsCredential	Get, Remove, Set	Configure account that DHCP server uses to manage DHS server records
DhcpServerInDC	Add, Get, Remove	Manage DHCP server authorization in Active Directory
DhcpServerSecurityGroup	Add	Add security groups to a DHCP server
DhcpServerSetting	Set, Get	Manage the DHCP server database configuration parameters

DhcpServerv4Binding	Get, Set	Manage the network interfaces to which the DHCP server services are bound
DhcpServerv4Class	Add, Remove, Get, Set	Manage IPv4 vendor or user classes for the DHCP server service
DhcpServerv4DnsSetting	Set, Get	Configure DNS settings at the scope, reservation, or server level
DhcpServerv4ExclusionRange	Add, Get, Remove	Manage IPv4 exclusion ranges
DhcpServerv4Failover	Set, Add, Remove, Get	Configure IPv4 failover relationships
DhcpServerv4FailoverReplication	Invoke	Replicate scope configuration between failover partners
DhcpServerv4FailoverScope	Add, Remove	Configure DHCP IPv4 failover scopes
DhcpServerv4Filter	Remove, Add, Get	Manage DHCP IPv4 MAC address filters
DhcpServerv4FilterList	Set, Get	Manage the filter list status
DhcpServerv4FreeIPAddress	Get	Find free IP addresses in a scope
DhcpServerv4IP	Repair	Reconcile inconsistent records in

Record		the DHCP database.
DhcpServerv4Lease	Add, Remove, Get	Manage IPv4 DHCP leases
DhcpServerv4MulticastExclusionRange	Get, Remove, Add	Manage multicast exclusions
DhcpServerv4MulticastLease	Get, Remove	Manage IPv4 multicast lease
DhcpServerv4MulticastScope	Remove, Get, Add, Set	Manage IPv4 multicast scopes
DhcpServerv4MulticastScopeStatistics	Get	View IPv4 multicast scope statistics
DhcpServerv4OptionDefinition	Add, Set, Remove, Get	Manage IPv4 option definitions
DhcpServerv4OptionValue	Get, Remove, Set	Configure IPv4 option values
DhcpServerv4Policy	Add, Remove, Get, Set	Manage IPv4 policies
DhcpServerv4PolicyIPRange	Get, Add, Remove	Manage bindings between IP ranges and policies
DhcpServerv4Re	Add, Set, Get,	Configure IPv4 reservations

servation	Remove	
DhcpServerv4Scope	Get, Remove, Set, Add	Manage IPv4 DHCP scopes
DhcpServerv4ScopeStatistics	Get	View IPv4 DHCP scope statistics
DhcpServerv4Statistics	Get	View IPv4 DHCP statistics
DhcpServerv4Superscope	Remove, Rename, Get, Add	Manage IPv4 DHCP superscopes
DhcpServerv4SuperscopeStatistics	Get	View IPv4 DHCP superscope statistics
DhcpServerv6Binding	Set, Get	Manage binding of DHCP server to IPv6 interfaces
DhcpServerv6Class	Remove, Set, Get, Add	Manage IPv6 vendor or user classes
DhcpServerv6DnsSetting	Get, Set	Configure IPv6 DNS settings at the scope, reservation, or server
DhcpServerv6ExclusionRange	Get, Add, Remove	Manage IPv6 exclusions
DhcpServerv6FreeIPAddress	Get	Locate free IPv6 addresses

DhcpServerv6Lease	Add, Remove, Get	Manage IPv6 leases
DhcpServerv6OptionDefinition	Add, Remove, Get, Set	Manage IPv6 option definitions
DhcpServerv6OptionValue	Set, Remove, Get	Manage IPv6 option values
DhcpServerv6Reservation	Remove, Add, Get, Set	Manage IPv6 DHCP reservations
DhcpServerv6Scope	Remove, Add, Get, Set	Manage IPv6 scope
DhcpServerv6ScopeStatistics	Get	View IPv6 scope statistics
DhcpServerv6StatelessStatistics	Get	View IPv6 stateless statistics
DhcpServerv6StatelessStore	Set, Get	Manage the properties of the stateless store
DhcpServerv6Statistics	Get	View IPv6 server statistics
DhcpServerVersion	Get	View DHCP server version information

IPAM

IPAM allows you to centralize the management of DHCP and DNS servers. Rather than managing each server separately, you can use IPAM to manage them from a single console. You can use a single IPAM server to manage up to 150 separate DHCP servers and up to 500 individual DNS servers. Also, it can manage 6,000 separate DHCP scopes and 150 separate DNS zones. You can perform tasks, such as creating address scopes, configuring address reservations, and managing DHCP and DNS options globally, instead of having to perform these tasks on a server-by-server basis. IPAM supports both Active Directory and file-based DNS zones. IPAM can also be integrated into System Center Virtual Machine Manager (VMM) and can keep track of VMM IP address pools.

Deploy IPAM

You can only install the IPAM feature on a computer that is a member of an Active Directory domain. You can have multiple IPAM servers within a single Active Directory forest. You are likely to do this if your organization is geographically dispersed. It's important to note that IPAM cannot manage a DHCP or DNS server if the role is installed on the same computer that hosts IPAM. For this reason, you should install the IPAM feature on a server that doesn't host the DNS or

DHCP roles. IPAM is also not supported on computers that host the Domain Controller server role. Additionally, if you want to use the IPAM server to manage IPv6 address ranges, you need to ensure that IPv6 is enabled on the computer hosting the IPAM server.

Configure server discovery

Server discovery is the process by which the IPAM server checks with Active Directory to locate Domain Controllers, DNS servers, and DHCP servers. You select which domains to discover on the Configure Server Discovery dialog box. After you've completed server discovery, you need to run a special PowerShell cmdlet that creates and provisions group policy objects that allow these servers to be managed by the IPAM server. When you set up the IPAM server, you choose a GPO name prefix. Use this prefix when executing the *Invoke-IpamGpoProvisioning* Windows PowerShell cmdlet that creates the appropriate GPOS.

If you use the GPO prefix IPAM, the three GPOS are named:

- IPAM_DC_NPS
- IPAM_DHCP
- IPAM_DNS

Until these GPOs apply to the discovered servers, these

servers are listed as having an IPAM Access status of blocked. After the GPOs are applied to the discovered servers, the IPAM Access status is set to Unblocked. After the discovered service has an IPAM Access Status set to Unblocked, you can edit the properties of the server and set it to Managed. After you do this, you can use IPAM to manage the selected services on the server.

IPAM supports managing DNS and DHCP servers in forests where a two-way trust relationship has been configured. You perform DHCP and DNS server discovery in a trusted forest on a per-domain basis.

IPAM Administration

You can delegate administrative permissions by adding user accounts to one of five local security groups on the IPAM server. By default, members of the Domain Admins and Enterprise Admins groups can perform all tasks on the IPAM server. The nine following local security groups are present on the IPAM server after you deploy the role, and they allow you to delegate the following permissions:

- **DNS Record Administrator.** Members of this role can manage DNS resource records.
- **IP Address Record Administrator.** Members of this role can manage IP addresses, but not address spaces, ranges, subnets, or blocks.
- **IPAM Administrators.** Members of this role can perform all tasks on the IPAM server, including viewing IP Address tracking

information.

- **IPAM ASM Administrators.** ASM stands for Address Space Administrator. Users added to this group can perform all tasks that can be performed by members of the IPAM Users group, but they are also able to manage the IP address space. They cannot perform monitoring tasks and are unable to perform IP address tracking tasks.
- **IPAM DHCP Administrator.** Members of this role can manage DHCP servers using IPAM.
- **IPAM DHCP Reservations Administrator.** Members of this role can manage DHCP reservations using IPAM.
- **IPAM DHCP Scope Administrator.** Members of this role can manage DHCP scopes.
- **IPAM MSM Administrators.** MSM stands for Multi-Server-Management. Members of this role can manage all DNS and DHCP servers managed by IPAM.
- **IPAM DNS Administrator.** Members of this role can manage DHCP servers using IPAM.

Managing the IP address space

The benefit of IPAM is that it allows you to manage all the IP addresses in your organization. IPAM supports the management of IPv4 public and private addresses, whether they are statically or dynamically assigned. IPAM allows you to detect whether there are overlapping IP address ranges defined in DHCP scopes on different servers; determine IP address utilization and whether there are free IP addresses in a specific range; create DHCP reservations centrally without having to configure them on individual DHCP servers;

and create DNS records based on IP address lease information.

IPAM separates the IP address space into blocks, ranges, and individual addresses. An IP address block is a large collection of IP addresses that you use to organize the address space that your organization uses at the highest level. An organization might only have one or two address blocks: one for its entire internal network and another smaller block that represents the public IP address space used by the organization. An IP address range is part of an IP address block. An IP address range cannot map to multiple IP address blocks. Generally, an IP address range corresponds to a DHCP scope. An individual IP address maps to a single IP address range. IPAM stores the following information with an IP address:

- Any associated MAC addresses
- How the address is assigned (static or DHCP)
- Assignment date
- Assignment expiration
- Which service manages the IP address
- DNS records associated with the IP address

IP address tracking

One of the most important features of IPAM is its ability to track IP addresses by correlating DHCP leases with

user and computer authentication events on managed domain controllers and NPS servers. IP address tracking allows you to figure out which user was associated with a specific IP address at a particular point in time, which can be important when trying to determine the cause of unauthorized activity on the organizational network.

You can search for IP address records using one of the following four parameters:

- **Track by IP address.** You can track by IPv4 address, but IPAM does not support tracking by IPv6 address.
- **Track by client ID.** You can track by client ID in IPAM, which allows you to track IP address activity by MAC address.
- **Track by host name.** You can track by the computer's name as registered in DNS.
- **Track by user name.** You can track a user name, but to do this, you must also provide a host name.

Naturally, you can only track data that has been recorded since IPAM has been deployed. So, while it is possible to store several years of data in the Windows Internal Database that IPAM uses, you are limited to only being able to retrieve events recorded after IPAM was configured.

Microsoft estimates that the Windows Internal Database IPAM uses is able to store 3 years of IP address utilization data for an organization that has 100,000 users before data must be purged.

You can use IPAM to locate free subnets and free IP address ranges in your existing IP address scheme. For example, if you need 50 consecutive free IP addresses for a project, you can use the *Find-IpamFreeRange* PowerShell cmdlet to determine whether such a range is present in your current environment.

Managing IPAM with PowerShell

You can manage IPAM with Windows PowerShell cmdlets that are contained in the IPAMServer module. Table 5-3 outlines the functionality of these cmdlets.

Table 5-3 IPAM PowerShell cmdlets

Noun	Verbs	Function
IpamAccessScope	Set	Configure an IPAM access scope
IpamAddress	Export, Remove, Get, Import, Add, Set	Use this set of cmdlets to manage IPAM addresses
IpamAddressSpace	Set, Remove, Get, Add	Use this set of cmdlets to manage an IPAM address space
IpamAddressUtilizationThreshold	Set, Get	Configure the IPAM address utilization thresholds

Technet24

IpamBlock	Remove, Get, Add, Set	Manage IPAM address blocks
IpamCapability	Disable, Get, Enable	Manage optional IPAM capabilities
IpamConfiguration	Set, Get	Manage the configuration of the IPAM server
IpamConfigurationEvent	Get, Remove	View IPAM configuration events stored in the IPAM database
IpamCustomField	Rename, Add, Remove, Get	Manage IPAM custom fields
IpamCustomFieldAssociation	Remove, Add, Set, Get	Configure association between custom field values
IpamCustomValue	Remove, Add, Rename	Manage values assigned to custom fields in IPAM
IpamDatabase	Move, Get, Set	Manage the IPAM database
IpamDhcpConfigurationEvent	Remove, Get	Manage DHCP configuration events stored in the IPAM database
IpamDhcpScope	Get	View IPAM DHCP scope information

IpamDhcpServer	Get	View IPAM DHCP server information
IpamDhcpSuperscope	Get	View IPAM DHCP superscope information
IpamDiscoveryDomain	Set, Remove, Add, Get	Configure IPAM discovery domains
IpamDnsConditionalForwarder	Get	View information about IPAM DNS conditional forwarders
IpamDnsResourceRecord	Get	View information about IPAM DNS resource records
IpamDnsServer	Get	View information about IPAM DNS servers
IpamDnsZone	Get	View information about IPAM DNS zones
IpamFreeAddress	Find	Locate available IP addresses from the IPAM database
IpamFreeRange	Find	Locate ranges of free IP addresses from the IPAM database
IpamFreeSubnet	Find	Locate free subnets in the IPAM database

IpamGpoProvisioning	Invoke	Configures GPOs necessary to prepare DNS and DHCP servers managed by the IPAM server
IpamIpAddressAuditEvent	Remove, Get	Manage IPAM audit events in the IPAM database
IpamRange	Export, Import, Add, Get, Remove, Set	Manage IPAM IP address ranges
IpamServer	Update	Updates an IPAM server after you update the host server's operating system
IpamServerInventory	Remove, Add, Get, Set	Manage the properties of an infrastructure server stored in the IPAM server inventory
IpamServerProvisioning	Invoke	Use this cmdlet to automate the provisioning of an IPAM server
IpamSubnet	Add, Get, Set, Import, Export, Remove	Manage an IP subnet stored in the IPAM server database
IpamUtilizationData	Remove	Remove existing IPAM IP address utilization information

Chapter 6

Hyper-V

Dynamic memory

Smart paging

Resource metering

Guest integration services

Generation 2 VMs

Enhanced Session Mode

Discrete Device Assignment

Nested virtualization

PowerShell Direct

HVC for Linux

Virtual hard disks

Managing checkpoints

Virtual Fibre Channel adapters

Storage QoS

- Hyper-V storage optimization
- Hyper-V virtual switches
- Virtual machine network adapters
- Optimizing network performance
- Virtual machine MAC addresses
- Network isolation
- Hyper-V replica
- Hyper-V failover clusters
- Hyper-V guest clusters
- Live migration
- Storage migration
- Exporting, importing, and copying VMs
- VM Network Health Detection
- VM drain on shutdown
- Domain controller cloning
- Shielded virtual machines
- Managing Hyper-V using PowerShell

Hyper-V is a virtualization platform that is built into Windows Server. Not only can you use Hyper-V to host virtual machines and a special type of containers,

Hyper-V is integrated into the very fabric of Microsoft's Azure cloud. Hyper-V is also available on some editions of Windows 10. This means that it's possible to transfer a virtual machine from Windows 10 to Windows Server to Azure and back without needing to alter the virtual machine's format.

In this chapter, we look at many topics, including memory configuration, generation 2 VMs, nested virtualization, virtual hard disks, replicas and failovers, live migration, storage migration, guarded virtualization fabrics, and shielded virtual machines.

DYNAMIC MEMORY

You have two options when assigning memory to VMs. You can assign a static amount of memory, or you can configure dynamic memory. When you assign a static amount of memory, the amount of memory assigned to the VM remains the same, whether the VM is starting up, is currently running, or is in the process of shutting down.

When you configure dynamic memory, you can configure the following values in Windows Admin Center:

- **Startup Memory.** This is the amount of memory allocated to the VM during startup. This can be the same as the minimum amount of memory, or it can be as large as the maximum amount of allocated memory. Once the VM has started, it will instead use the

amount of memory configured as the Minimum Memory.

- **Minimum Memory.** This is the minimum amount of memory that the VM will be assigned by the virtualization host when dynamic memory is enabled. When multiple VMs are demanding memory, Hyper-V may reallocate memory away from the VM until this Minimum Memory value is met. You can reduce the Minimum Memory setting while the VM is running, but you cannot increase it while the VM is running.
- **Maximum Memory.** This is the maximum amount of memory that the VM will be allocated by the virtualization host when dynamic memory is enabled. You can increase the Maximum Memory setting while the VM is running, but you cannot decrease it while the VM is running.
- **Memory Buffer.** This is the percentage of memory that Hyper-V should allocate to the VM as a buffer.
- **Memory Weight.** This setting allows you to configure how memory should be allocated to this particular VM as compared to other VMs running on the same virtualization host.

Generally, when you configure dynamic memory, the amount of memory used by a VM will fluctuate between the Minimum Memory and Maximum Memory values. You should monitor VM memory utilization and tune these values so that they accurately represent the VM's actual requirements. If you allocate a Minimum Memory value below what the VM would actually need to run, this shortage might cause the virtualization host to reduce the amount of memory allocated to this Minimum Memory value, which will cause the VM to stop running.

SMART PAGING

Smart paging is a special technology in Hyper-V that functions in certain conditions when a VM is restarting. Smart paging uses a file on the disk to simulate memory to meet Startup Memory requirements when the Startup Memory setting exceeds the Minimum Memory setting. Startup Memory is the amount of memory allocated to the VM when it starts, but not when it is in a running state. For example, you could set Startup Memory to 2,048 MB and the Minimum Memory to 512 MB for a specific virtual machine. In a scenario where 1,024 MB of free memory was available on the virtualization host, smart paging would allow the VM to access the required 2,048 MB of memory.

Because it uses disk to simulate memory, smart paging is only active if the following three conditions occur at the same time:

- The VM is being restarted.
- There is not enough memory on the virtualization host to meet the Startup Memory setting.
- Memory cannot be reclaimed from other VMs running on the same host.

Smart paging doesn't allow a VM to perform a "cold start" if the required amount of Startup Memory is not available but the Minimum Memory amount is available. Smart paging is only used when a VM that was

already running restarts and the above three conditions have been met.

You can configure the location of the smart paging file on a per-VM basis. By default, smart paging files are written to the *C:\ProgramData\Microsoft\Windows\Hyper-V* folder. The smart paging file is created only when needed and is deleted within 10 minutes of the VM restarting.

RESOURCE METERING

Resource metering allows you to track the consumption of processor, disk, memory, and network resources by individual VMs. To enable resource metering, use the *Enable-VMResourceMetering* Windows PowerShell cmdlet. You can view metering data using the *Measure-VM* Windows PowerShell cmdlet. Resource metering allows you to record the following information:

- Average CPU use
- Average memory use
- Minimum memory use
- Maximum memory use
- Maximum disk allocation
- Incoming network traffic
- Outgoing network traffic

Average CPU use is measured in megahertz (MHz). All

other metrics are measured in megabytes. Although you can extract data using the *Measure-VM* cmdlet, you need to use another solution to output this data into a visual form, such as a graph.

GUEST INTEGRATION SERVICES

Integration services allow the virtualization host to extract information and perform operations on a hosted VM. By default, Windows Server 2019, Windows Server 2016, Windows Server 2012 R2, Windows 10, and Windows 8.1 include Hyper-V integration services. If you are running a Linux guest VM on your Hyper-V Server, you can download the Linux Integration Services for Hyper-V and Azure from Microsoft's website. Integration services installation files are available for all operating systems that are supported on Hyper-V. You can enable the following integration services:

- **Operating System Shutdown.** This integration service allows you to shut down the VM from the virtualization host, rather than from within the VM's OS.
- **Time Synchronization.** Synchronizes the virtualization host's clock with the VM's clock. Ensures that the VM clock doesn't drift when the VM is started, stopped, or reverted to a checkpoint.
- **Data Exchange.** Allows the virtualization host to read and modify specific VM registry values.
- **Heartbeat.** Allows the virtualization host to verify that the VM OS is still functioning and responding to requests.
- **Backup (Volume Checkpoint).** For VMs that support Volume

Shadow Copy, this service synchronizes with the virtualization host, allowing backups of the VM while the VM is in operation.

- **Guest Services.** Guest services allow you to copy files from the virtualization host to the VM using the *Copy-VMFile* Windows PowerShell cmdlet.

GENERATION 2 VMS

Generation 2 VMs are a special type of VM that differ in configuration from the VMs that are now termed "generation 1 VMs," which could be created on Hyper-V virtualization hosts running the Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012 operating systems. Generation 2 VMs are supported on Windows Server 2012 R2 and later operating systems. Generation 2 VMs are also now supported in Azure.

Generation 2 VMs provide the following functionality:

- Can boot from a SCSI virtual hard disk
- Can boot from a SCSI virtual DVD
- Supports UEFI firmware on the VM
- Supports VM Secure Boot
- PXE boot using standard network adapter

There are no legacy network adapters with generation 2 VMs, and most legacy devices, such as COM ports and the Diskette Drive, are no longer present. Generation 2 VMs are "virtual first" and are not designed to simulate hardware for computers that have undergone physical-

to-virtual (P2V) conversion. If you need to deploy a VM that requires an emulated component such as a COM port, you'll need to deploy a generation 1 VM.

You configure the generation of a VM during the VM creation. Once a VM is created, Hyper-V doesn't allow you to modify the VM's generation. Windows Server 2016 and Windows Server 2019 support both generation 1 and generation 2 VMs.

Generation 2 VMs boot more quickly and allow the installation of operating systems more quickly than generation 1 VMs. Generation 2 VMs have the following limitations:

- You can only use generation 2 VMs if the guest operating system is running an x64 version of Windows 10, Windows 8.1, Windows 8, Windows Server 2016, Windows Server 2012 R2, or Windows Server 2012.
- Generation 2 VMs only support virtual hard disks in VHDX format.

ENHANCED SESSION MODE

Enhanced Session Mode allows you to perform actions including cutting and pasting, audio redirection, and volume and device mapping when using Virtual Machine Connection windows. You can also sign in to a VM with a smart card through Enhanced Session Mode. You enable Enhanced Session Mode on the Hyper-V server by selevting Allow Enhanced Session Mode in the Enhanced Session Mode Policy section of the Hyper-V

server's Properties dialog box.

You can only use Enhanced Session Mode with guest VMs running Windows Server 2019, Windows Server 2016, Windows Server 2012 R2, Windows 10, or Windows 8.1. To utilize Enhanced Session Mode, you must have permission to connect to the VM using Remote Desktop through the account you use to sign in to the guest VM. You can grant permission to the VM by adding the user to the Remote Desktop Users group. A user who is a member of the local Administrators group also has this permission. The Remote Desktop Services service must be running on the guest VM.

DISCRETE DEVICE ASSIGNMENT

Discrete Device Assignment (DDA) is available in Windows Server 2019 and Windows Server 2016 that allows you to directly assign a physical GPU or an NVMe storage device to a specific virtual machine. Each physical GPU or NVMe storage device can only be associated with one VM. DDA involves installing the device's native driver in the VM associated with that GPU. This process works for both Windows and Linux VMs as long as the drivers are available for the VM's operating system. DDA is supported for generation 1 or 2 VMs running Windows Server 2012 R2 or later, Windows 10, and some Linux guest operating systems.

Before assigning the physical GPU or storage device to a specific VM, you need to dismount the device from the Hyper-V host. Some device vendors provide partitioning drivers for the Hyper-V host. Partitioning drivers are different from the standard device drivers and improve the security of the DDA configuration. If a partitioning driver is available, you should install this driver prior to dismounting the device from the Hyper-V host. If no driver is available, you'll have to use the *-Force* option with the *Dismount-VMHostAssignableDevice* cmdlet.

Once you've dismounted the physical device from the Hyper-V host, you can assign it to a specific guest VM using the *Add-VMAssignableDevice* cmdlet.

Enabling DDA requires that you disable the VM's automatic stop action. You can do this using the following command:

click here to view code image

```
Set-VM -Name VMName -AutomaticStopAction TurnOff
```

In addition to disabling the automatic stop action, the following functionality is also not available to VMs configured with DDA:

- VM Save and Restore
- VM live migration
- Use of dynamic memory
- Deployment of VM to a high availability cluster

More Info

Direct Device Assignment

You can learn more about Direct Device Assignment at https://docs.microsoft.com/en-us/windows-server/virtualization/hyper-v/plan/plan-for-deploying-devices-using-discrete-device-assignment.

Inside OUT

RemoteFX

RemoteFX is a technology available in Windows Server 2016 that performed a similar function to Discrete Device Assignment in Windows Server 2019. RemoteFX provides a 3D virtual adapter and USB redirection support for VMs. You can only use RemoteFX if the virtualization host has a compatible GPU. RemoteFX allows one or more compatible graphics adapters to perform graphics processing tasks for multiple VMs. Like Discrete Device Assignment, you can use RemoteFX to provide support for graphically intensive applications, such as CAD, in VDI scenarios. RemoteFX is deprecated from version 1803 of Windows Server onward.

NESTED VIRTUALIZATION

Hyper-V on Windows Server 2019, Windows Server 2016, and Windows 10 client OS support nested virtualization. Nested virtualization allows you to enable Hyper-V and host virtual machines on a virtual machine running under Hyper-V as long as that VM is running Windows Server 2019, Windows Server 2016, or Windows 10. Nested virtualization can be enabled on a per-VM basis by running the following PowerShell command:

Click here to view code image

```
Set-VM Processor -VMName NameOfVM -ExposeVirtualiationExtensions
$true
```

Nested virtualization dynamic memory

Once you've run this command, you'll be able to enable Hyper-V on the VM. You won't be able to adjust the memory of a virtual machine that is enabled for nested virtualization while that virtual machine is running. While it is possible to enable dynamic memory, the amount of memory allocated to a virtual machine configured for nested virtualization will not fluctuate while the VM is running.

Nested virtualization networking

To route network packets through the multiple virtual switches required during nested virtualization, you can either enable MAC address spoofing or configure network address translation (NAT).

To enable MAC address spoofing on the virtual machine that you have configured for nested virtualization, run the following PowerShell command:

Click here to view code image

```
Get-VMNetworkAdapter -VMName NameOfVM | Set-VMNEtworkAdapter -
MacAddressSpoofing On
```

To enable NAT, create a virtual NAT switch in the VM that has been enabled for nested virtualization using the following PowerShell commands:

Click here to view code image

```
New-VMSwitch -name VMNAT -SwitchTypeInternal
```

```
New-NetNAT -Name LocalNAT -InternalIPInterfaceAddressPrefix
“192.168.15.0/24”
```

```
Get-NetAdapter “vEthernet (VmNat)” | New-NetIPAddress -IPAddress
192.168.15.1 -AddressFamily IPv4 -PrefixLength 24
```

Once you’ve done this, you’ll need to manually assign IP addresses to VMs running under the VM enabled for nested virtualization, using the default gateway of 192.168.15.1. You can use a separate internal addressing scheme other than 192.168.15.0/24 by altering the appropriate PowerShell commands in the example above.

More Info

Nested Virtualization

You can learn more about Hyper-V nested virtualization at https://docs.microsoft.com/en-us/virtualization/hyper-v-on-windows/user-guide/nested-virtualization.

POWERSHELL DIRECT

PowerShell Direct allows you to create a remote PowerShell session directly from a Hyper-V host to a virtual machine hosted on that Hyper-V host without requiring the virtual machine to be configured with a network connection. PowerShell Direct requires that both the Hyper-V host and the VM be running Windows Server 2019, Windows Server 2016, or Windows 10;

PowerShell Direct is not supported for earlier versions of the Windows Server or Windows client operating systems.

To use PowerShell Direct, you must be signed in locally to the Hyper-V host with Hyper-V Administrator privileges. You also must have access to valid credentials for the virtual machine. If you do not have credentials for the virtual machine, you will not be able to establish a PowerShell Direct connection.

To establish a PowerShell Direct connection, use the command:

Click here to view code image

```
Enter-PSSession -vmname NameOfVM
```

You exit the PowerShell Direct session by using the *Exit-PSSession* cmdlet.

More Info

PowerShell Direct

You can learn more about PowerShell Direct at https://docs.microsoft.com/en-us/virtualization/hyper-v-on-windows/user-guide/powershell-direct.

HVC FOR LINUX

HVC.exe, which is included with Windows Server 2019 and Windows 10, allows you to make a remote SSH connection from a Hyper-V host to a Linux virtual

machine guest without requiring the Linux VM to have a functioning network connection. It provides similar functionality to PowerShell Direct. For HVC.exe to work, you'll need to ensure that the Linux virtual machine has an updated kernel and that the Linux integration services are installed. You'll also need the SSH server on the Linux VM to be installed and configured before you'll be able to use HVC.exe to initiate an SSH connection.

VIRTUAL HARD DISKS

Hyper-V supports two separate virtual hard disk formats. Virtual hard disk files in .vhd format are limited to 2040 GB. Virtual hard disks in this format can be used on all supported versions of Hyper-V. Other than the size limitation, the important thing to remember is that you cannot use virtual hard disk files in .vhd format with generation 2 VMs.

Virtual hard disk files in .vhdx format are an improvement over virtual hard disks in .vhd format. The main limitation of virtual hard disks in .vhdx format is that they cannot be used with Hyper-V on Windows Server 2008 or Windows Server 2008 R2, which shouldn't be as much of a problem now as these operating systems no longer have mainstream support. Virtual hard disks in .vhdx format have the following benefits:

- Can be up to 64 TB in size
- Have larger block size for dynamic and differential disks
- Provide 4-KB logical sector virtual disks
- Have an internal log that reduces chance of corruption
- Support trim to reclaim unused space

You can convert hard disks between .vhd and .vhdx format. You can create virtual hard disks at the time you create the VM by using the New Virtual Hard Disk Wizard or the *New-VHD Windows* PowerShell cmdlet.

Fixed-sized disks

Virtual hard disks can either by dynamic, differencing, or fixed. When you create a fixed-size disk, all space used by the disk is allocated on the hosting volume at the time of creation. Fixed disks increase performance if the physical storage medium does not support Windows Offloaded Data Transfer. Improvements in Windows Server reduce the performance benefit of fixed-size disks when the storage medium supports Windows Offloaded Data Transfer. The space to be allocated to the disk must be present on the host volume when you create the disk. For example, you can't create a 3-TB fixed disk on a volume that has only 2 TB of space.

Dynamically expanding disks

Dynamically expanding disks use an initial small file and then grow as the VM allocates data to the virtual

hard disk. This means you can create a 3-TB dynamic virtual hard disk on a 2-TB volume because the entire 3 TB will not be allocated at disk creation. However, in this scenario, you would need to ensure that you extend the size of the 2-TB volume before the dynamic virtual disk outgrows the available storage space.

Differencing disks

Differencing disks are a special type of virtual hard disk that has a child relationship with a parent hard disk. Parent disks can be fixed size or dynamic virtual hard disks, but the differencing disk must be the same type as the parent disk. For example, you can create a differencing disk in .vhdx format for a parent disk that uses .vhdx format, but you cannot create a differencing disk in .vhd format for a parent disk in .vhdx format.

Differencing disks record the changes that would otherwise be made to the parent hard disk by the VM. For example, differencing disks are used to record Hyper-V VM checkpoints. One parent virtual hard disk can have multiple differencing disks associated with it.

For example, you can create a specially prepared parent virtual hard disk by installing Windows Server on a VM by running the sysprep utility within the VM and then shutting down the VM. You can use the virtual hard disk created by this process as a parent virtual hard disk. In this scenario, when you create new Windows Server

VMs, you would configure the VMs to use a new differencing disk that uses the sysprepped virtual hard disk as a parent. When you run the new VM, it will write any changes that it would normally make to the full virtual hard disk to the differencing disk. In this scenario, deploying new Windows Server VMs becomes a simple matter of creating new VMs that use a differencing disk that uses the sysprepped Windows Server virtual hard disk as a parent.

You can create differencing hard disks using the New Virtual Hard Disk Wizard or the *New-VHD* Windows PowerShell cmdlet. You need to specify the parent disk during the creation process.

The key to using differencing disks is to ensure that you don't make changes to the parent disk; doing so will invalidate the relationship with any child disks. Generally, differencing disks can provide storage efficiencies because only changes are recorded on child disks. For example, rather than storing 10 different instances of Windows Server 2019 in its entirety, you could create one parent disk and have 10 much smaller differencing disks to accomplish the same objective. If you store VM virtual hard disks on a volume that has been deduplicated, these efficiencies are reduced.

Modifying virtual hard disks

You can perform the following tasks to modify existing

virtual hard disks:

- Convert a virtual hard disk in .vhd format to .vhdx format
- Convert a virtual hard disk in .vhdx format to .vhd format
- Change the disk from fixed size to dynamically expanding or from dynamically expanding to fixed size
- Shrink or enlarge the virtual hard disk

You convert virtual hard disk types (.vhd to .vhdx, .vhdx to .vhdx, dynamic to fixed, or fixed to dynamic) either using the Edit Virtual Hard Disk Wizard or by using the *Convert-VHD* Windows PowerShell cmdlet. When converting from .vhdx to .vhd, remember that virtual hard disks in .vhd format cannot exceed 2,040 GB in size. So, while it is possible to convert virtual hard disks in .vhdx format that are smaller than 2,040 GB to .vhd format, you will not be able to convert virtual hard disks that are larger than 2,040 GB.

You can only perform conversions from one format to another and from one type to another while the VM is powered off. You must shrink the virtual hard disk using the disk manager in the VM operating system prior to shrinking the virtual hard disk using the Edit Virtual Hard Disk Wizard or the *Resize-VHD* cmdlet. You can resize a virtual hard disk while the VM is running under the following conditions:

- The virtualization host is running Windows Server 2019, Windows Server 2016, or Windows Server 2012 R2.

- The virtual hard disk is in .vhdx format.
- The virtual hard disk is attached to a virtual SCSI controller.
- The virtual hard disk must have been shrunk. You must shrink the virtual hard disk using disk manager in the host operating system prior to shrinking the virtual hard disk using the Edit Virtual Hard Disk Wizard or the *Resize-VHD* cmdlet.

Pass-through disks

Pass-through disks, also known as directly attached disks, allow a VM to directly access the underlying storage rather than accessing a virtual hard disk that resides on that storage. For example, with Hyper-V, you normally connect a VM to a virtual hard disk file hosted on a volume formatted with NTFS or ReFS. With pass-through disks, the VM instead accesses the disk directly, and there is no virtual hard disk file.

Pass-through disks allow VMs to access larger volumes than are possible when using virtual hard disks in .vhd format. In earlier versions of Hyper-V—such as the version available with Windows Server 2008—pass-through disks provide performance advantages over virtual hard disks. The need for pass-through disks has diminished with the availability of virtual hard disks in .vhdx format because .vhdx format allows you to create much larger volumes.

Pass-through disks can be directly attached to the virtualization host, or they can be attached to Fibre

Channel or iSCSI disks. When adding a pass-through disk, you will need to ensure that the disk is offline. You can use the Disk Management console or the diskpart.exe utility on the virtualization host to set a disk to be offline.

To add a pass-through disk using Windows PowerShell, use the *Get-Disk* cmdlet to get the properties of the disk that you want to add as a pass-through disk. Next, pipe the result to the *Add-VMHardDiskDrive* cmdlet. For example, to add physical disk 3 to the VM named Alpha-Test, execute the following command:

Click here to view code image

```
Get-Disk 3 | Add-VMHardDiskDrive –VMName Alpha-Test
```

A VM that uses pass-through disks will not support VM checkpoints. Pass-through disks also cannot be backed up with backup programs that use the Hyper-V VSS writer.

MANAGING CHECKPOINTS

Checkpoints represent the state of a VM at a particular point in time. You can create checkpoints when the VM is running or when the VM is shut down. When you create a checkpoint of a running VM, the running VM's memory state is also stored in the checkpoint. Restoring a checkpoint taken of a running VM returns the running

VM to a restored state. Creating a checkpoint creates either an .avhd or .avhdx file (depending on whether the VM is using virtual hard disks in the .vhd or .vhdx format).

Windows Server 2019 and Windows Server 2016 support two types of checkpoints:

- **Standard checkpoints.** These function just as checkpoints have functioned in previous versions of Hyper-V. These checkpoints capture the state, date, and hardware configuration of a virtual machine. They are designed for development and test scenarios.
- **Production checkpoints.** Available only in Windows Server 2019 and Windows Server 2016, production checkpoints use backup technology inside the guest as opposed to the saved-state technology used in standard checkpoints. Production checkpoints are fully supported by Microsoft and can be used with production workloads, which is something that was not supported with the standard version of checkpoints available in previous versions of Hyper-V.

You can switch between standard and production checkpoints on a per-virtual machine basis by editing the Properties of the virtual machine; in the Management section of the Properties dialog box, choose between Production and Standard checkpoints.

You can create checkpoints from Windows PowerShell with the *Checkpoint-VM* cmdlet. The other checkpoint-related Windows PowerShell cmdlets in Windows Server actually use the VMSnapshot noun, though on Windows 10, they confusingly have aliases for the

VMCheckPoint noun.

The Windows Server checkpoint related cmdlets are as follows:

- ***Restore-VMSnapshot.*** Restores an existing VM checkpoint.
- ***Export-VMSnapshot.*** Allows you to export the state of a VM as it exists when a particular checkpoint was taken. For example, if you took checkpoints at 2 PM and 3 PM, you could choose to export the checkpoint taken at 2 PM and then import the VM in the state that it was in at 2 PM on another Hyper-V host.
- ***Get-VMSnapshot.*** Lists the current checkpoints.
- ***Rename-VMSnapshot.*** Allows you to rename an existing VM checkpoint.
- ***Remove-VMSnapshot.*** Deletes a VM checkpoint. If the VM checkpoint is part of the chain but not the final link, changes are merged with the successive checkpoint, so that the checkpoint remains a representation of the VM at the point in time when the snapshot was taken. For example, if checkpoints were taken at 1 PM, 2 PM, and 3 PM, and you delete the 2 PM checkpoint, the .avhd/.avhdx files associated with the 2 PM snapshot would be merged with the .avhd/.avhdx files associated with the 3 PM snapshot, so that the 3 PM snapshot retains its integrity.

Checkpoints do not replace backups. Checkpoints are almost always stored on the same volume as the original VM hard disks, so a failure of that volume will result in all VM storage files—both original disks and checkpoint disks—being lost. If a disk in a checkpoint chain becomes corrupted, then that checkpoint and all subsequent checkpoints will be lost. Disks earlier in the checkpoint chain will remain unaffected. Hyper-V

supports a maximum of 50 checkpoints per VM.

VIRTUAL FIBRE CHANNEL ADAPTERS

Virtual Fibre Channel allows you to make direct connections from VMs running on Hyper-V to Fibre Channel storage. If the following requirements are met, Virtual Fibre Channel is supported on Windows Server 2019 and Windows Server 2016:

- The computer functioning as the Hyper-V virtualization host must have a Fibre Channel host bus adapter (HBA) that has a driver that supports virtual Fibre Channel.
- SAN must be NPIV (N_Port ID) enabled.
- The VM must be running a supported version of the guest operating system.
- Virtual Fibre Channel LUNs cannot be used to boot Hyper-V VMs.

VMs running on Hyper-V support up to four virtual Fibre Channel adapters, each of which can be associated with a separate Storage Area Network (SAN).

Before you can use a virtual Fibre Channel adapter, you will need to create at least one virtual SAN on the Hyper-V virtualization host. A virtual SAN is a group of physical Fibre Channel ports that connect to the same SAN.

VM live migration and VM failover clusters are supported; however, virtual Fibre Channel does not

support VM checkpoints, host-based backup, or live migration of SAN data.

STORAGE QOS

Storage Quality of Service (QoS) allows you to limit the maximum number of IOPS (input/output operations per second) for virtual hard disks. IOPS are measured in 8-KB increments. If you specify a maximum IOPS value, the virtual hard disk will be unable to exceed this value. You use Storage QoS to ensure that no single workload on a Hyper-V virtualization host consumes a disproportionate amount of storage resources.

It's also possible to specify a minimum IOPS value for each virtual hard disk. You would do this if you wanted to be notified that a specific virtual hard disk's IOPS has fallen below a threshold value. When the number of IOPS falls below the specified minimum, an event is written to the event log. You configure Storage QoS on a per-virtual hard disk basis.

HYPER-V STORAGE OPTIMIZATION

Several technologies built into Windows Server 2019 and Windows Server 2016 allow you to optimize the performance and data storage requirements for files associated with VMs.

Deduplication

In Windows Server 2019, both ReFS and NTFS volumes support deduplication. Deduplication is a process by which duplicate instances of data are removed from a volume and replaced with pointers to the original instance. Deduplication is especially effective when used with volumes that host virtual hard disk files because many of these files contain duplicate copies of data, such as the VM's operating system and program files.

Once installed, you can enable deduplication through the Volumes node of the File And Storage Services section of the Server Manager Console. When enabling deduplication, you specify whether you want to use a general file server data deduplication scheme or a virtual desktop infrastructure scheme. For volumes that host VM files, the VDI scheme is appropriate. You can't enable deduplication on the operating system volume; deduplication may only be enabled on data volumes. For this reason, remember to store VM configuration files and hard disks on a volume that is separate from the operating system volume.

Storage tiering

Storage tiering is a technology that allows you to mix fast storage, such as solid-state disk (SSD), with traditional spinning magnetic disks to optimize both storage performance and capacity. Storage tiering works on the premise that a minority of the data stored on a

volume is responsible for the majority of read and write operations. Storage tiering can be enabled through the storage spaces functionality, and rather than creating a large volume that consists entirely of SSDs, you create a volume comprised of both solid-state and spinning magnetic disks. In this configuration, frequently accessed data is moved to the parts of the volume hosted on the SSDs, and less frequently accessed data is moved to the parts of the volume hosted on the slower spinning magnetic disks. This configuration allows many of the performance benefits of an SSD-only volume to be realized without the cost of using SSD-only storage.

When used in conjunction with deduplication, frequently accessed deduplicated data is moved to the faster storage, providing reduced storage requirements, while improving performance over what would be possible if the volume hosting VM files were solely comprised of spinning magnetic disks. You also have the option of pinning specific files to the faster storage, which overrides the algorithms that move data according to accumulated utilization statistics. You configure storage tiering using Windows PowerShell.

HYPER-V VIRTUAL SWITCHES

Hyper-V virtual switches, called Hyper-V virtual networks in previous versions of Hyper-V, represent

network connections to which the Hyper-V virtual network adapters can connect. You can configure three types of Hyper-V virtual switches: external switches, internal switches, and private switches.

External switches

An external switch connects to a physical or wireless network adapter. Only one virtual switch can be mapped to a specific physical or wireless network adapter or NIC team. For example, if a virtualization host had four physical network adapters configured as two separate NIC teams, you could configure two external virtual switches. If a virtualization host had three physical network adapters that did not participate in any NIC teams, you could configure three external virtual switches. VMs connected to the same external switch can communicate with each other as well as with external hosts connected to the network to which the network adapter mapped to the external switch is connected. For example, if an external switch is connected to a network adapter that is connected to a network that can route traffic to the Internet, a VM connected to that external virtual switch will also be able to connect to hosts on the Internet. When you create an external switch, a virtual network adapter that maps to this switch is created on the virtualization host unless you clear the option that allows the management

operating system to share the network adapter. If you clear this option, the virtualization host will not be able to communicate through the network adapter associated with the external switch.

Internal switches

An internal switch allows communication between the VM and the virtualization host. All VMs connected to the same internal switch can communicate with each other and the virtualization host. For example, you could successfully initiate an RDP connection from the virtualization host to an appropriately configured VM or use the *Test-NetConnection* Windows PowerShell cmdlet from a Windows PowerShell prompt on the virtualization host to get a response from a VM connected to an internal network connection. VMs connected to an internal switch are unable to use that virtual switch to communicate with hosts on a separate virtualization host that are connected to an internal switch with the same name.

Private switches

VMs connected to the same private switch on a VM host can communicate with one another, but they cannot communicate directly with the virtualization host. Private switches only allow communication between VMs on the same virtualization host. For example, VMs

alpha and beta are connected to private switch p_switch_a on virtualization host h_v_one. VM gamma is connected to private switch p_switch_a on virtualization host h_v_two. VMs alpha and beta will be able to communicate with each other, but they won't be able to communicate with h_v_one or VM gamma.

VIRTUAL MACHINE NETWORK ADAPTERS

Generation 1 VMs support two types of network adapters: synthetic network adapters and legacy network adapters. A synthetic network adapter uses drivers that are provided when you install integration services in the VM operating system. If a VM operating system doesn't have these drivers or if integration services are not available for this operating system, then the network adapter will not function. Synthetic network adapters are unavailable until a VM operating system that supports them is running. This means that you can't perform a PXE boot from a synthetic network adapter if you have configured a generation 1 VM.

Legacy network adapters emulate a physical network adapter, similar to a multiport DEC/Intel 21140 10/100TX 100 MB card. Many operating systems, including those that do not support virtual machine integration services, support this network adapter. This

means that if you want to run an operating system in a VM that doesn't have virtual machine integration services support—such as a version of Linux or BSD that isn't officially supported for Hyper-V—you'll need to use a legacy network adapter because it is likely to be recognized by the guest VM operating system.

Legacy network adapters on generation 1 VMs also function before the VM guest operating system is loaded. This means that if you want to PXE boot a generation 1 VM—for example, if you wanted to use WDS to deploy an operating system to the VM—you'd need to configure the VM with a legacy network adapter.

Generation 2 VMs don't separate synthetic and legacy network adapters and have only a single network adapter type. Generation 2 VMs support PXE booting from this single network adapter type. Generation 2 VMs also support "hot add" network adapters, allowing you to add or remove network adapters to or from a virtual machine while it is running. It is important to remember that only recent Windows client and server operating systems and only certain Linux operating systems are supported as generation 2 VMs.

OPTIMIZING NETWORK PERFORMANCE

You can optimize network performance for VMs hosted

on Hyper-V in a number of ways. For example, you can configure the virtualization host with separate network adapters connected to separate subnets. You do this to separate network traffic related to the management of the Hyper-V virtualization host from network traffic associated with hosted VMs. You can also use NIC teaming on the Hyper-V virtualization host to provide increased and fault-tolerant network connectivity. You'll learn more about NIC teaming later in this chapter.

Bandwidth management

An additional method of optimizing network performance is to configure bandwidth management at the virtual network adapter level. Bandwidth management allows you to specify a minimum and the maximum traffic throughput for a virtual network adapter. The minimum bandwidth allocation is an amount that Hyper-V will reserve for the network adapter. For example, if you set the minimum bandwidth allocation to 10 Mbps for each VM, Hyper-V would ensure that when other VMs needed more, they would be able to increase their bandwidth utilization until they reached a limit defined by the combined minimum bandwidth allocation of all VMs hosted on the server. Maximum bandwidth allocations specify an upper limit for bandwidth utilization. By default, no minimum or maximum limits are set on virtual network

adapters.

You configure bandwidth management by selevting the Enable Bandwidth Management option on a virtual network adapter and specifying a minimum and maximum bandwidth allocation in megabits per second (Mbps).

SR-IOV

SR-IOV (Single Root I/O Virtualization) increases network throughput by bypassing a virtual switch and sending network traffic straight to the VM. When you configure SR-IOV, the physical network adapter is mapped directly to the VM. As such, SR-IOV requires that the VM's operating system include a driver for the physical network adapter. You can only use SR-IOV if the physical network adapter and the network adapter drivers used with the virtualization host support the functionality. You can only configure SR-IOV for a virtual switch during switch creation. Once you have an SR-IOV–enabled virtual switch, you can then enable SR-IOV on the virtual network adapter that connects to that switch.

Dynamic virtual machine queue

Dynamic virtual machine queue is an additional technology that you can use to optimize network performance. When a VM is connected through a virtual

switch to a network adapter that supports virtual machine queue and virtual machine queue is enabled on the virtual network adapter's properties, the physical network adapter can use Direct Memory Access (DMA) to forward traffic directly to the VM. With virtual machine queue, network traffic is processed by the CPU assigned to the VM rather than by the physical network adapter used by the Hyper-V virtualization host. Dynamic virtual machine queue automatically adjusts the number of CPU cores used to process network traffic. Dynamic virtual machine queue is automatically enabled on a virtual switch when you enable virtual machine queue on the virtual network adapter.

Virtual machine NIC teaming

NIC teaming allows you to aggregate bandwidth across multiple network adapters while also providing a redundant network connection if one of the adapters in the team fails. NIC teaming allows you to consolidate up to 32 network adapters and to use them as a single network interface. You can configure NIC teams using adapters that are from different manufacturers and that run at different speeds (though it's generally a good idea to use the same adapter make and model in production environments).

You can configure NIC teaming at the virtualization host level if the virtualization host has multiple network

adapters. The drawback is that you can't configure NIC teaming at the host level if the network adapters are configured to use SR-IOV. If you want to use SR-IOV and NIC teaming, create the NIC team instead in the VM. You can configure NIC teaming within VMs by adding adapters to a new team using the Server Manager console or the *New-NetLbfoTeam* PowerShell cmdlet.

When configuring NIC teaming in a VM, ensure that each virtual network adapter that will participate in the team has MAC address spoofing enabled

VIRTUAL MACHINE MAC ADDRESSES

By default, VMs running on Hyper-V hosts use dynamic MAC addresses. Each time a VM is powered on, it will be assigned a MAC address from a MAC address pool. You can configure the properties of the MAC address pool through the MAC Address Range settings available through Virtual Switch Manager.

When you deploy operating systems on physical hardware, you can use two methods to ensure that the computer is always assigned the same IP address configuration. The first method is to assign a static IP address from within the virtualized operating system. The second is to configure a DHCP reservation that always assigns the same IP address configuration to the

MAC address associated with the physical computer's network adapter.

This won't work with Hyper-V VMs in their default configuration because the MAC address may change if you power the VM off and then on. Rather than configure a static IP address using the VM's operating system, you can instead configure a static MAC address on a per-virtual network adapter basis. This will ensure that a VM's virtual network adapter retains the same MAC address whether the VM is restarted or even if the VM is migrated to another virtualization host.

To configure a static MAC address on a per-network adapter basis, edit the network adapter's advanced features. When entering a static MAC address, you will need to selevt a MAC address manually. You shouldn't use one from the existing MAC address pool because there is no way for the current virtualization hosts or other virtualization hosts on the same subnet to check whether a MAC address that is to be assigned dynamically has already been assigned statically.

NETWORK ISOLATION

Hyper-V supports VLAN (Virtual Local Area Network) tagging at both the network adapter and virtual switch level. VLAN tags allow the isolation of traffic for hosts connected to the same network by creating separate

broadcast domains. Enterprise hardware switches also support VLANs as a way of partitioning network traffic. To use VLANs with Hyper-V, the virtualization hosts' network adapter must support VLANs. A VLAN ID has 12 bits, which means you can configure 4,095 VLAN IDs.

You configure VLAN tags at the virtual network adapter level by selevting Enable Virtual LAN Identification in the Virtual Network Adapter Properties dialog box. VLAN tags applied at the virtual switch level override VLAN tags applied at the virtual network adapter level. To configure VLAN tags at the virtual switch level, selevt the Enable Virtual LAN Identification For Management Operating System option and specify the VLAN identifier.

HYPER-V REPLICA

Hyper-V replica provides a replica of a VM running on one Hyper-V host that can be stored and updated on another Hyper-V host. For example, you could configure a VM hosted on a Hyper-V failover cluster in Melbourne to be replicated through Hyper-V replica to a Hyper-V failover cluster in Sydney. Hyper-V replication allows for replication across site boundaries and does not require access to shared storage in the way that failover clustering does.

Hyper-V replication is asynchronous. While the replica

copy is consistent, it is a lagged copy with changes sent only as frequently as once every 30 seconds. Hyper-V replication supports multiple recovery points, with a recovery snapshot taken every hour. (This incurs a resource penalty, so it is off by default.) This means that when activating the replica, you can choose to activate the most up-to-date copy or a lagged copy. You would choose to activate a lagged copy if some form of corruption or change made the up-to-date copy problematic.

When you perform a planned failover from the primary host to the replica, you need to switch off the primary host. This ensures that the replica is in an up-to-date and consistent state. This is a drawback compared to failover or live migration where the VM will remain available during the process. A series of checks are completed before performing planned failover to ensure that the VM is off, that reverse replication is allowed back to the original primary Hyper-V host, and that the state of the VM on the current replica is consistent with the state of the VM on the current primary. Performing a planned failover will start the replicated VM on the original replica, which will now become the new primary server.

Hyper-V replica also supports unplanned failover. You perform an unplanned failover if the original Hyper-V host has failed or the site that hosts the primary replica

has become unavailable. When performing unplanned failover, you can choose either the most recent recovery point or a previous recovery point. Performing unplanned failover will start the VM on the original replica, which will now become the new primary server.

Hyper-V extended replication allows you to create a second replica of the existing replica server. For example, you could configure Hyper-V replication between a Hyper-V virtualization host in Melbourne and Sydney, with Sydney hosting the replica. You could then configure an extended replica in Brisbane using the Sydney replica.

Configuring Hyper-V replica servers

To configure Hyper-V replication, you need to configure the Replication Configuration settings. The first step is to selevt Enable This Computer As A Replica Server. Next, selevt the authentication method you will use. If the computers are parts of the same Active Directory environment, you can use Kerberos. When you use Kerberos, Hyper-V replication data isn't encrypted when transmitted across the network. If you are concerned about encrypting network data, you could configure IPsec. If you are concerned about encrypting replication traffic, another option is to use certificate-based authentication. This is useful if you are transmitting data across the public Internet without using an

encrypted VPN tunnel. When using certificate-based authentication, you'll need to import and selevt a public certificate issued to the partner server.

The final step when configuring Hyper-V replica is to selevt the servers from which the Hyper-V virtualization host will accept incoming replicated VM data. One option is to have the Hyper-V virtualization host accept replicated VMs from any authenticated Hyper-V virtualization host, using a single default location to store replica data. The other option is to configure VM replica storage on a per-server basis. For example, if you wanted to store VM replicas from one server on one volume and VM replicas from another server on a different volume, you'd configure VM replica storage on a per-server basis.

Once replication is configured on the source and destination servers, you'll also need to enable the pre-defined firewall rules to allow the incoming replication traffic. There are two rules: one for replication using Kerberos (HTTP) on port 80 and the other for using certificate-based authentication on port 443.

Configuring VM replicas

Once you have configured the source and destination replica servers, you need to configure replication on a per-VM basis. You do this by running the Enable Replication wizard, which you can trigger by clicking

Enable Replication when the VM is selevted in Hyper-V Manager. To configure VM replicas, you must perform the following steps:

- **Selevt Replica Server.** Selevt the replica server name. If you are replicating to a Hyper-V failover cluster, you'll need to specify the name of the Hyper-V replica broker. You'll learn more about Hyper-V replica broker later in this chapter.
- **Choose Connection Parameters.** Specify the connection parameters. The options will depend on the configuration of the replica servers. On this page, depending on the existing configuration, you can choose the authentication type and whether replication data will be compressed when transmitted over the network.
- **Selevt Replication VHDs.** When configuring replication, you have the option of not replicating some of a VM's virtual hard disks. In most scenarios, you should replicate all a VM's hard disk drives. One reason not to replicate a VM's virtual hard disk would be if the virtual hard disk only stores frequently changing temporary data that wouldn't be required when recovering the VM.
- **Replication Frequency.** Use this to specify the frequency with which changes are sent to the replica server. You can choose between intervals of 30 seconds, 5 minutes, and 15 minutes.
- **Additional Recovery Points.** You can choose to create additional hourly recovery points. Doing this gives you the option of starting the replica from a previous point in time rather than the most recent. The advantage is that this allows you to roll back to a previous version of the VM if data corruption occurs and has replicated to the most recent recovery point. The replica server can store a maximum of 24 recovery points.
- **Initial Replication.** The last step in configuring Hyper-V replica is choosing how to seed the initial replica. Replication works by sending changed blocks of data, so the initial replica, which sends

the entire VM, will be the largest transfer. You can perform an offline transfer with external media, use an existing VM on the replica server as the initial copy (the VM for which you are configuring a replica must have been exported and then imported on the replica server) or transfer all VM data across the network. You can perform replication immediately or at a specific time in the future, such as 2 AM when network utilization is likely to be lower.

Replica failover

You perform a planned replica failover when you want to run the VM on the replica server rather than on the primary host. Planned failover involves shutting down the VM, which ensures that the replica will be up to date. Contrast this with Hyper-V live migration, which you perform while the VM is running. When performing a planned failover, you can configure the VM on the replica server to automatically start once the process completes; you can also configure reverse replication, so that the current replica server becomes the new primary server, and the current primary becomes the new replica server.

If the primary server becomes unavailable, you can trigger an unplanned failover. You would then perform the unplanned failover on the replica server (as the primary is not available). When performing an unplanned failover, you can selevt any of the up to 24 previously stored recovery points.

Hyper-V replica broker

You need to configure and deploy Hyper-V replica broker if your Hyper-V replica configuration includes a Hyper-V failover cluster as a source or destination. You don't need to configure and deploy Hyper-V replica broker if both the source and destination servers are not participating in a Hyper-V failover cluster. You install the Hyper-V Replica Broker role using Failover Cluster Manager after you've enabled the Hyper-V role on cluster nodes.

HYPER-V FAILOVER CLUSTERS

One of the most common uses for failover clustering is to host Hyper-V virtual machines. The Hyper-V and failover clustering roles are even supported in Nano Server, which allows organizations to deploy failover clustering–capable virtualization hosts that have a minimal operating system footprint.

Hyper-V host cluster storage

When deployed on Hyper-V host clusters, the configuration and virtual hard disk files for highly available VMs are hosted on shared storage. This shared storage can be one of the following:

- **Serial Attached SCSI (SAS).** Suitable for two-node failover clusters where the cluster nodes are near each other.
- **iSCSI storage.** Suitable for failover clusters with two or more nodes. Windows Server includes iSCSI Target Software, allowing it

to host iSCSI targets that can be used as shared storage by Windows failover clusters.

- **Fibre Channel.** Fibre Channel/Fibre Channel over Ethernet storage requires special network hardware. While generally providing better performance than iSCSI, Fibre Channel components tend to be more expensive.
- **SMB 3.0 file shares configured as continuously available storage.** This special type of file share is highly available, with multiple cluster nodes able to maintain access to the file share. This configuration requires multiple clusters. One cluster hosts the highly available storage used by the VMs, and the other cluster hosts the highly available VMs.
- **Cluster Shared Volumes (CSVs).** CSVs can also be used for VM storage in Hyper-V failover clusters. As with continuously available file shares, multiple nodes in the cluster have access to the files stored on CSVs, ensuring that failover occurs with minimal disruption. As with SMB 3.0 file shares, multiple clusters are required, with one cluster hosting the CSVs and the other cluster hosting the VMs.

When considering storage for a Hyper-V failover cluster, remember the following:

- Ensure volumes used for disk witnesses are formatted as either NTFS or ReFS.
- Avoid allowing nodes from separate failover clusters to access the same shared storage by using LUN masking or zoning.
- Where possible, use storage spaces to host volumes presented as shared storage.

Cluster quorum

Hyper-V failover clusters remain functional until they do not have enough active votes to retain quorum. Votes

can consist of nodes that participate in the cluster as well as disk or file share witnesses. The calculation on whether the cluster maintains quorum is dependent on the cluster quorum mode. When you deploy a Windows Server failover cluster, one of the following modes will automatically be selevted, depending on the current cluster configuration:

- Node Majority
- Node and Disk Majority
- Node and File Share Majority
- No Majority: Disk Only

You can change the cluster mode manually; or, with Dynamic Quorum in Windows Server 2019, the cluster mode will change automatically when you add or remove nodes, a witness disk, or a witness share.

Node Majority

The Node Majority cluster quorum mode is chosen automatically during setup if a cluster has an odd number of nodes. When this cluster quorum mode is used, a file share or disk witness is not used. A failover cluster will retain quorum as long as the number of available nodes is more than the number of failed nodes that retain cluster membership. For example, if you deploy a nine-node failover cluster, the cluster will retain quorum as long as five cluster nodes are able to communicate with each other.

Node and Disk Majority

The Node and Disk Majority model is chosen automatically during setup if the cluster has an even number of nodes and shared storage is available to function as a disk witness. In this configuration, cluster nodes and the disk witness each have a vote when calculating quorum. As with the Node Majority model, the cluster will retain quorum as long as the number of votes that remain in communication exceeds the number of votes that cannot be contacted. For example, if you deployed a six-node cluster and a witness disk, there would be a total of seven votes. As long as four of those votes remained in communication with each other, the failover cluster would retain quorum.

Node and File Share Majority

The Node and File Share Majority model is used when a file share is configured as a witness. Each node and the file share have a vote when it comes to determining if quorum is retained. As with other models, a majority of the votes must be present for the cluster to retain quorum. Node and File Share Majority is suitable for organizations that are deploying multi-site clusters; for example, placing half the cluster nodes in one site, half the cluster nodes in another site, and the file share witness in a third site. If one site fails, the other site can retain communication with the site that hosts the file

share witness, in which case quorum is retained.

No Majority: Disk Only

The No Majority: Disk Only model must be configured manually and must only be used in testing environments because the only vote that counts toward quorum is that of the disk witness on shared storage. The cluster will retain quorum as long as the witness is available, even if every node but one fails. Similarly, the cluster will be in a failed state if all the nodes are available but the shared storage hosting the disk witness goes offline.

Cluster node weight

Rather than every node in the cluster having an equal vote when determining quorum, you can configure which cluster nodes can vote to determine quorum by running the Configure Cluster Quorum Wizard. Configuring node weight is useful if you are deploying a multi-site cluster and you want to control which site retains quorum if communication between the sites is lost. You can determine which nodes in a cluster are currently assigned votes by selevting Nodes in the Failover Cluster Manager.

Dynamic quorum

Dynamic quorum allows cluster quorum to be recalculated automatically each time a node is removed

from or added to a cluster. By default, dynamic quorum is enabled on Windows Server 2019 clusters. Dynamic quorum works in the following manner:

- The vote of the witness is automatically adjusted based on the number of voting nodes in the cluster. If the cluster has an even number of nodes, the witness has a vote. If a cluster has an even number of nodes and a node is added or removed, the witness loses its vote.
- In the event of a 50 percent node split, dynamic quorum can adjust the vote of a node. This is useful in avoiding "split brain" syndrome during site splits with multi-site failover clusters.

An advantage of dynamic quorum is that as long as nodes are evicted in a graceful manner, the cluster will reconfigure quorum appropriately. This means that you could change a nine-node cluster so that it was a five-node cluster by evicting nodes, and the new quorum model would automatically be recalculated assuming that the cluster had only five nodes. With dynamic quorum, it is a good idea to specify a witness even if the initial cluster configuration has an odd number of nodes; doing so means a witness vote will automatically be included in the event that an administrator adds or removes a node from the cluster.

Cluster networking

In lab and development environments, it's reasonable to have failover cluster nodes that are configured with a single network adapter. In production environments

with mission-critical workloads, you should configure cluster nodes with multiple network adapters, institute adapter teaming, and leverage separate networks. Separate networks should include:

- A network dedicated for connecting cluster nodes to shared storage
- A network dedicated for internal cluster communication
- The network that clients use to access services deployed on the cluster

When configuring IPv4 or IPv6 addressing for failover cluster nodes, ensure that addresses are assigned either statically or dynamically to cluster node network adapters. Avoid using a mixture of statically and dynamically assigned addresses, as this will cause an error with the Cluster Validation wizard. Also ensure that cluster network adapters are configured with a default gateway. While the cluster validation wizard will not provide an error if a default gateway is not present for the network adapters of each potential cluster node, you will be unable to create a failover cluster unless a default gateway is present.

Force Quorum Resiliency

Imagine you have a five-node, multi-site cluster in Melbourne and Sydney, with three nodes in Sydney. Also, imagine that Internet connectivity to the Sydney site is lost. Within the Sydney site itself, the cluster will remain running because with three nodes, it has

retained quorum. But if external connectivity to the Sydney site is not available, you may instead need to forcibly start the cluster in the Melbourne site (which will be in a failed state because only two nodes are present) using the */fq* (forced quorum) switch to provide services to clients.

In the past, when connectivity was restored, this would have led to a "split brain" or partitioned cluster, as both sides of the cluster would be configured to be authoritative. To resolve this with failover clusters running Windows Server 2012 or earlier, you would need to manually restart the nodes that were not part of the forced quorum set using the */pq* (prevent quorum) switch. Windows Server 2019 provides a feature known as Force Quorum Resiliency that automatically restarts the nodes that were not part of the forced quorum set, so that the cluster does not remain in a partitioned state.

Cluster Shared Volumes

Cluster Shared Volumes (CSVs) are a high-availability storage technology that allows multiple cluster nodes in a failover cluster to have read-write access to the same LUN. This has the following advantages for Hyper-V failover clusters:

- VMs stored on the same LUN can be run on different cluster nodes. This reduces the number of LUNs required to host VMs because the

VMs stored on a CSV aren't tied to one specific Hyper-V failover cluster node; instead, the VMs can be spread across multiple Hyper-V failover cluster nodes.

- Switch-over between nodes is almost instantaneous in the event of failover because the new host node doesn't have to go through the process of seizing the LUN from the failed node.

CSVs are hosted from servers running the Windows Server 2012 or later operating system and allow multiple nodes in a cluster to access the same NTFS- or ReFS-formatted file system. CSVs support BitLocker, with each node performing decryption of encrypted content using a cluster computer account. CSVs also integrate with SMB (Server Message Block) Multichannel and SMB Direct, allowing traffic to be sent through multiple networks and to leverage network cards that include Remote Direct Memory Access (RDMA) technology. CSVs can also automatically scan and repair volumes without requiring storage to be taken offline. You can convert cluster storage to a CSV using the Disks node of the Failover Cluster Manager.

Active Directory detached clusters

Active Directory detached clusters, also called "clusters without network names," are a feature of Windows Server 2012 R2 and later operating systems. Detached clusters have their names stored in DNS but do not require the creation of a computer account within Active Directory. The benefit of detached clusters is that it is

possible to create them without requiring that a computer object be created in Active Directory to represent the cluster; also, the account used to create the cluster need not have permissions to create computer objects in Active Directory. Although the account used to create a detached cluster does not need permission to create computer objects in Active Directory, the nodes that will participate in the detached cluster must still be domain joined.

Inside OUT

Detached Clusters

While many things are possible, not all things are a good idea. When discussing AD detached clusters with other IT professionals, the general consensus seems to be that deploying them is often more trouble than having the computer object be created in Active Directory. Ultimately, both AD detached clusters and traditional Windows Server failover clusters still require domain membership, so you might as well go all in and deploy the cluster in the traditional manner.

Preferred owner and failover settings

Cluster role preferences settings allow you to configure a preferred owner for a cluster role. When you do this, the role will be hosted on the node listed as the preferred owner. You can specify multiple preferred owners and configure the order in which a role will attempt to return to a specific cluster node.

Failover settings allow you to configure how many times a service will attempt to restart or failover in a specific period. By default, a cluster service can failover twice in

a six-hour period before the failover cluster will leave the cluster role in a failed state. The failback setting allows you to configure the amount of time a clustered role that has failed over to a node that is not its preferred owner will wait before falling back to the preferred owner.

HYPER-V GUEST CLUSTERS

A guest cluster is a failover cluster that consists of two or more VMs as cluster nodes. You can run a guest cluster on a Hyper-V failover cluster, or you can run guest clusters with nodes on separate Hyper-V failover clusters. While deploying a guest cluster on Hyper-V failover clusters may seem as though it is taking redundancy to an extreme, there are good reasons to deploy guest clusters and Hyper-V failover clusters together, including:

- **Failover clusters monitor the health of clustered roles to ensure that they are functioning.** This means that a guest failover cluster can detect when the failure of a clustered role occurs and can take steps to recover the role. For example, you deploy a SQL Server failover cluster as a guest cluster on a Hyper-V failover cluster. One of the SQL Server instances that participates in the guest cluster suffers a failure. In this scenario, failover occurs within the guest cluster, and another instance of SQL Server hosted on the other guest cluster node continues to service client requests.
- **Deploying guest and Hyper-V failover clusters together allows you to move applications to other guest cluster nodes while you are performing servicing tasks.** For

example, you may need to apply software updates that require a restart to the operating system that hosts a SQL Server instance. If this SQL Server instance is participating in a Hyper-V guest cluster, you could move the clustered role to another node, apply software updates to the original node, perform the restart, and then move the clustered SQL Server role back to the original node.

- **Deploying guest and Hyper-V failover clusters together allows you to live migrate guest cluster VMs from one host cluster to another host cluster while ensuring clients retain connectivity to clustered applications.** For example, a two-node guest cluster hosting SQL Server is hosted on one Hyper-V failover cluster in your organization's datacenter. You want to move the guest cluster from its current host Hyper-V failover cluster to a new Hyper-V failover cluster. By migrating one guest cluster node at a time from the original Hyper-V failover cluster to the new Hyper-V failover cluster, you'll be able to continue to service client requests without interruption, failing over SQL Server to the guest node on the new Hyper-V failover cluster after the first node completes its migration and before migrating the second node across.

Hyper-V guest cluster storage

Just as you can configure a Hyper-V failover cluster where multiple Hyper-V hosts function as failover cluster nodes, you can configure failover clusters within VMs, where each failover cluster node is a VM. Even though failover cluster nodes must be members of the same Active Directory domain, there is no requirement that they be hosted on the same cluster. For example, you could configure a multi-site failover cluster where the cluster nodes are hosted as highly available VMs, each hosted on its own Hyper-V failover clusters in each

site.

When considering how to deploy a VM guest cluster, you will need to choose how you will provision the shared storage that is accessible to each cluster node. The options for configuring shared storage for VM guest clusters include:

- iSCSI
- Virtual Fibre Channel
- Cluster Shared Volumes
- Continuously Available File Shares
- Shared virtual hard disks

The conditions for using iSCSI, Virtual Fibre Channel, Cluster Shared Volumes, and Continuously Available File Shares with VM guest clusters are essentially the same for VMs as they are when configuring traditional physically hosted failover cluster nodes.

Shared virtual hard disk

Shared virtual hard disks are a special type of shared storage only available to VM guest clusters. With shared virtual hard disks, each guest cluster node can be configured to access the same shared virtual hard disk. Each VM cluster node's operating system will recognize the shared virtual hard disk as shared storage when building the VM guest failover cluster.

Shared virtual hard disks have the following

requirements:

- Can be used with generation 1 and generation 2 VMs.
- Can only be used with guest operating systems running Windows Server 2012 or later. If the guest operating systems are running Windows Server 2012, they must be updated to use the Windows Server 2012 R2 integration services components.
- Can only be used if virtualization hosts are running the Windows Server 2012 R2 or later version of Hyper-V.
- Must be configured to use the .vhdx virtual hard disk format.
- Must be connected to a virtual SCSI controller.
- When deployed on a failover cluster, the shared virtual hard disk itself should be located on shared storage, such as a Continuously Available File Share or Cluster Shared Volume. This is not necessary when configuring a guest failover cluster on a single Hyper-V server that is not part of a Hyper-V failover cluster.
- VMs can only use shared virtual hard disks to store data. You can't boot a VM from a shared virtual hard disk.

The configuration of shared virtual hard disks differs from the traditional configuration of VM guest failover clusters because you configure the connection to shared storage by editing the VM properties rather than connecting to the shared storage from within the VM. Windows Server 2019 and Windows Server 2016 support shared virtual hard disks being resized and used with Hyper-V replica.

Hyper-V VHD Sets

VHD Sets are a newer version of shared virtual hard

disks. Hyper-V VHD Sets use a new virtual hard disk format that uses the .vhds extension. VHD Sets support online resizing of shared virtual disks, Hyper-V Replicas, and application-consistent Hyper-V checkpoints.

You can create a VHD Set file from Hyper-V Manager or by using the *New-VHD* cmdlet with the file type set to .vhds when specifying the virtual hard disk name. You can use the *Convert-VHD* cmdlet to convert an existing shared virtual hard disk file to a VHD Set file as long as you have taken the VMs that use the shared virtual hard disk file offline and have removed the shared virtual hard disk from the VM using the *Remove-VHHardDiskDrive* cmdlet.

More Info

Hyper-V VHD Sets

You can learn more about Hyper-V VHD Sets at https://docs.microsoft.com/en-us/windows-server/virtualization/hyper-v/manage/create-vhdset-file.

LIVE MIGRATION

Live migration is the process of moving an operational VM from one physical virtualization host to another with no interruption to VM clients or users. Live migration is supported between cluster nodes that share storage between separate Hyper-V virtualization hosts that are not participating in a failover cluster using a SMB 3.0 file share as storage; live migration is even

supported between separate Hyper-V hosts that are not participating in a failover cluster using a process called "shared-nothing live migration."

Live migration has the following prerequisites:

- There must be two or more servers running Hyper-V that use processors from the same manufacturer (for example, all Hyper-V virtualization hosts configured with Intel processors or all Hyper-V virtualization hosts configured with AMD processors).
- Hyper-V virtualization hosts need to be members of the same domain, or they must be members of domains that have a trust relationship with each other.
- VMs must be configured to use virtual hard disks or virtual Fibre Channel disks; pass-through disks are not allowed.

It is possible to perform live migration with VMs configured with pass-through disks under the following conditions:

- VMs are hosted on a Windows Server Hyper-V failover cluster.
- Live migration will be within nodes that participate in the same Hyper-V failover cluster.
- VM configuration files are stored on a Cluster Shared Volume.
- The physical disk that is used as a pass-through disk is configured as a storage disk resource that is controlled by the failover cluster. This disk must be configured as a dependent resource for the highly available VM.

If performing a live migration using shared storage, the following conditions must be met:

- The SMB 3.0 share needs to be configured so that the source and the destination virtualization host's computer accounts have read

and write permissions.

- All VM files (virtual hard disks, configuration files, and snapshot files) must be located on the SMB 3.0 share. You can use storage migration to move VM files to an SMB 3.0 share while the VM is running prior to performing a live migration using this method.

You must configure the source and destination Hyper-V virtualization hosts to support live migrations by enabling live migrations in the Hyper-V settings. When you do this, you specify the maximum number of simultaneous live migrations and the networks that you will use for live migration. Microsoft recommends using an isolated network for live migration traffic, though this is not a requirement.

The next step in configuring live migration is choosing which authentication protocol and live migration performance options to use. You selevt these in the Advanced Features area of the Live Migrations settings. The default authentication protocol is CredSSP (Credential Security Support Provider). CredSSP requires local sign-in to both source and destination Hyper-V virtualization hosts to perform live migration. Kerberos allows you to trigger live migration remotely. To use Kerberos, you must configure the computer accounts for each Hyper-V virtualization host with constrained delegation for the cifs and Microsoft Virtual System Migration Service services, granting permissions to the virtualization hosts that will participate in the live

migration partnership. The performance options allow you to speed up live migration. Compression increases processor utilization. SMB will use SMB Direct if both network adapters used for the live migration process support Remote Direct Memory Access (RDMA) and RDMA capabilities are enabled.

STORAGE MIGRATION

With storage migration, you can move a VM's virtual hard disk files, checkpoint files, smart paging files, and configuration files from one location to another. You can perform storage migration while the VM is running or while the VM is powered off. You can move data to any location that is accessible to the Hyper-V host. This allows you to move data from one volume to another, from one folder to another, or even to an SMB 3.0 file share on another computer. When performing storage migration, choose the Move The VM's Storage option.

For example, you could use storage migration to move VM files from one Cluster Share Volume to another on a Hyper-V failover cluster without interrupting the VM's operation. You have the option of moving all data to a single location, moving VM data to separate locations, or moving only the VM's virtual hard disk. To move the VM's data to different locations, selevt the items you want to move and the destination locations.

EXPORTING, IMPORTING, AND COPYING VMS

A VM export creates a duplicate of a VM that you can import on the same or different Hyper-V virtualization host. When performing an export, you can choose to export the VM, which includes all its VM checkpoints, or you can choose to export just a single VM checkpoint. Windows Server 2019, Windows Server 2016, and Windows Server 2012 R2 support exporting a running VM. With Hyper-V in Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008, it is necessary to shut down the VM before performing an export.

Exporting a VM with all its checkpoints will create multiple differencing disks. When you import a VM that was exported with checkpoints, these checkpoints will also be imported. If you import a VM that was running at the time of export, the VM is placed in a saved state. You can resume from this saved state, rather than having to restart the VM.

When importing a VM, you can choose from the following options:

- **Register The Virtual Machine In Place (Use The Existing ID).** Use this option when you want to import the VM while keeping the VM files in their current locations. Because this method uses the existing VM ID, you can only use it if the original VM on which the export was created is not present on the host to which

you wish to import the VM.

- **Restore The Virtual Machine (Use The Existing Unique ID).** Use this option when you want to import the VM while moving the files to a new location; for example, you would choose this option if you are importing a VM that was exported to a network share. Because this method also uses the existing VM ID, you can only use it if the original VM on which the export was created is not present on the host to which you wish to import the VM.
- **Copy The Virtual Machine (Create A New Unique ID).** Use this method if you want to create a separate clone of the exported VM. The exported files will be copied to a new location, leaving the original exported files unaltered. A new VM ID is created, meaning that the cloned VM can run concurrently on the same virtualization host as the original progenitor VM. When importing a cloned VM onto the same virtualization host as the original progenitor VM, ensure that you rename the newly imported VM; otherwise, you may confuse the VMs.

VM NETWORK HEALTH DETECTION

VM Network Health Detection is a feature for VMs that are deployed on Hyper-V host clusters. With VM Network Health Detection, you configure a VM's network adapter settings and mark certain networks as being protected. You do this in the Advanced Features section of the Network Adapter Properties dialog box.

If a VM is running on a cluster node where the network marked as protected becomes unavailable, the cluster will automatically live migrate the VM to a node where

the protected network is available. For example, you have a four-node Hyper-V failover cluster. Each node has multiple network adapters and a virtual switch named Alpha maps as an external virtual switch to a physical network adapter on each node. A VM, configured as highly available and hosted on the first cluster node, is connected to virtual switch Alpha. The network adapter on this VM is configured with the protected network option. After the VM has been switched on and has been running for some time, a fault occurs causing the physical network adapter mapped to virtual switch Alpha on the first cluster node to fail. When this happens, the VM will automatically be live migrated to another cluster node where virtual switch Alpha is working.

VM DRAIN ON SHUTDOWN

VM drain on shutdown is a feature that will automatically live migrate all running VMs off a node if you shut down that node without putting it into maintenance mode. If you are following best practice, you'll put nodes into maintenance mode and live migrating running workloads away from nodes that you will restart or intend to shut down anyway. The main benefit of VM drain on shutdown is that, in the event that you are having a bad day and forget to put a cluster node into maintenance mode before shutting it down or

restarting it, any running VMs will be live migrated without requiring your direct intervention.

DOMAIN CONTROLLER CLONING

Windows Server 2019 and Windows Server 2016 support creating copies of domain controllers that are running as VMs as long as certain conditions are met. Cloned domain controllers have the following prerequisites:

- The virtualization host supports VM-GenerationID, a 128-bit random integer that identifies each VM checkpoint. VM-GenerationID is supported by the version of Hyper-V available with Windows Server 2012 and later, as well as some third-party hypervisors.
- The domain controller must be running Windows Server 2012 or later as its operating system.
- The server that hosts PDC emulator Flexible Single Master Operations Role (FSMO) must be contactable. This server must be running Windows Server 2012 or later as its operating system. Remember that you should be running the latest version of Windows Server for domain controllers, so of course, you will be running Windows Server 2019 on the servers hosting your forest and domain FSMO roles!
- The computer account of the domain controller that will serve as the template for cloning must be added to the Cloneable Domain Controllers security group.

Once you have met these conditions, you'll need to create an XML configuration file named DCCloneConfig.xml using the *New-ADDCCloneConfigFile* Windows PowerShell cmdlet.

Once created, you'll need to edit this file and specify settings such as computer name, network settings and Active Directory site information. You should also check the template DC using the *Get-ADDCCloningExcludedApplicationsList* cmdlet to determine whether any services that will cause problems with the cloning are present on the template DC, such as the DHCP server service.

SHIELDED VIRTUAL MACHINES

Shielded virtual machines are a special type of virtual machine that has a virtual Trusted Platform Module (TPM) chip that is encrypted using BitLocker and can only run on specific approved hosts that support what is known as a guarded fabric. Shielded VMs allow sensitive data and workloads to be run on virtualization hosts without the concern that an attacker or the administrator of the virtualization host might export the virtual machine to gain access to the sensitive data stored there. Only certain operating systems can be used for shielded guest virtual machines. You'll learn more about shielded VMs and guarded virtualization fabrics in Chapter 19, "Hardening Windows Server and Active Directory."

MANAGING HYPER-V USING POWERSHELL

Technet24

The Hyper-V PowerShell module contains a large number of cmdlets that you can use to manage all aspects of Hyper-V and virtual machines on Windows Server 2019. These cmdlets are described in Table 6-1.

Table 6-1 Remote Access PowerShell cmdlets

Noun	Verbs	Function
VFD	New	Create a new virtual floppy disk.
VHD	Convert, Merge, Resize, Optimize, Dismount, New, Mount, Set, Get, Test	Manage virtual hard disks.
VHDSet	Optimize, Get	Manage virtual hard disk set files, an update to the VHDX files from previous versions of Hyper-V.
VHDSnapshot	Get, Remove	Used to manage snapshots with VHD Sets.
VM	Wait, Move, Get, Export, Debug, Measure, Import, Save, Resume, Restart, Set, Suspend, Stop, Start, Repair, Checkpoint, Remove, New, Compare, Rename	Use to manage virtual machines.
VMAssignable	Add, Get, Remove	Manage VM assignable devices.

Device		
VMBios	Set, Get	Manage VM BIOS settings.
VMComPort	Set, Get	Manage VM COM Port settings.
VMConnectAccess	Grant, Revoke, Get	Manage which users can connect to specific VMs.
VMConsoleSupport	Enable, Disable	Manage HID device support for virtual machines.
VMDvdDrive	Add, Set, Get, Remove	Manage VM DVD drive settings.
VMEventing	Disable, Enable	Enable or disable VM eventing.
VMFailover	Complete, Stop, Start	Manage VM failover.
VMFibreChannelHba	Set, Get, Remove, Add	Manage VM fibre channel HBA settings.
VMFile	Copy	Copy a file from the virtualization host to a location within the virtual machine.

Technet24

VMFirmware	Set, Get	Manage VM virtual firmware settings.
VMGpuPartitionAdapter	Set, Add, Remove, Get	Manage VM GPU settings.
VMGroup	New, Get, Remove, Rename	Allows you to manage VM Groups, which you use to manage groups of machines that might make up a multi-tier application.
VMGroupMember	Remove, Add	Add and remove VMs from a specific group.
VMHardDiskDrive	Set, Add, Remove, Get	Manage VM hard disk drive settings.
VMHost	Set, Get	Configure VM host settings.
VMHostAssignableDevice	Add, Mount, Get, Remove, Dismount	Manage VM host assignable devices.
VMHos	Get, Set	Configure VM host cluster

tCluster		settings.
VMHostNumaNode	Get	View VM host NUMA node information.
VMHostNumaNodeStatus	Get	View VM host NUMA node status.
VMHostSupportedVersion	Get	View VM host supported information data.
VMIdeController	Get	Manage VM IDE controller settings.
VMInitialReplication	Import, Start, Stop	Configure initial VM replication settings.
VMIntegrationService	Enable, Disable, Get	Manage VM integration services.
VMKeyProtector	Get, Set	Manage VM key protector settings.

Technet24

VMKeyStorageDrive	Set, Remove, Get, Add	Manage VM key storage drive settings.
VMMemory	Set, Get	Manage VM memory.
VMMigration	Enable, Disable	Manage VM Migration.
VMMigrationNetwork	Add, Remove, Get, Set	Configure VM migration network settings.
VMNetworkAdapter	Disconnect, Connect, Remove, Test, Add, Set, Rename, Get	Configure VM network adapters.
VMNetworkAdapterAcl	Remove, Add, Get	Manage VM network adapter ACLs.
VMNetworkAdapterExtendedAcl	Remove, Add, Get	Manage VM Network adapter extended ACL settings.
VMNetworkA	Get, Set	Manage VM network adapter failover settings.

dapterFailoverConfiguration		
VMNetworkAdapterIsolation	Set, Get	Manage VM network adapter isolation settings.
VMNetworkAdapterRdma	Set, Get	Manage virtual network adapter RDMA settings.
VMNetworkAdapterRoutingDomainMapping	Set, Get, Remove, Add	Configure virtual network adapter routing domains.
VMNetworkAdapterTeamMapping	Remove, Get, Set	Configure virtual network adapter team mapping.
VMNetworkAdapterVlan	Get, Set	Manage VM network adapter VLAN settings.

Technet24

VMPartitionableGpu	Get, Set	Manage VM GPU settings.
VMProcessor	Get, Set	Configure VM Processor settings.
VMRemoteFx3dVideoAdapter	Remove, Set, Get, Add	Manage RemoteFX3d video adapters.
VMRemoteFXPhysicalVideoAdapter	Disable, Get, Enable	Manage RemoteFX physical adapters.
VMReplication	Get, Stop, Suspend, Remove, Measure, Enable, Resume, Set	Manage VM replication.
VMReplicationAuthorizationEntry	Set, Remove, New, Get	Manage VM replication authorization entries.
VMReplicationConnection	Test	Check VM replication.

VMReplicationServer	Get, Set	Configure VM replication server settings.
VMReplicationStatistics	Reset	Reset VM replication statistics.
VMResourceMetering	Reset, Disable, Enable	Configure VM resource metering.
VMResourcePool	Set, Rename, Get, Measure, New, Remove	Manage VM resource pool settings.
VMSan	Remove, Disconnect, Connect, New, Set, Rename, Get	Manage VM SAN settings.
VMSavedState	Remove	Remove saved state data.
VMScsiController	Remove, Add, Get	Manage VM SCSI controller.
VMSecurity	Get, Set	Manage VM security information.
VMSecurityPo	Set	Configure VM security policies.

licy		
VMSnapshot	Rename, Restore, Export, Get, Remove	Manage VM checkpoints.
VMStorage	Move	Move VM storage.
VMStoragePath	Get, Remove, Add	Configure VM storage path.
VMSwitch	Set, Get, Rename, Remove, Add, New	Manage virtual switches.
VMSwitchExtension	Enable, Disable, Get	Manage extensions on a virtual switch.
VMSwitchExtensionPortData	Get	View port extensions on a virtual switch.
VMSwitchExtensionPortFeature	Set, Add, Remove, Get	Manage switch port features.
VMSwitchExtensionS	Get	View switch port data.

witchData		
VMSwitchExtensionSwitchFeature	Add, Get, Set, Remove	Manage switch features.
VMSwitchTeam	Set, Get	Manage switch teams.
VMSwitchTeamMember	Add, Remove	Manage switch team membership.
VMSystemSwitchExtension	Get	View switch extensions installed on VM host.
VMSystemSwitchExtensionPortFeature	Get	View port-level features supported by switch extensions on VM host.
VMSystemSwitchExte	Get	View switch-level features supported by virtual switch extensions on the

nsionSwitchFeature		VM host.
VMTPM	Disable, Enable	Manage virtual TPM settings.
VMTrace	Stop, Start	Manage trace settings for VM debugging.
VMVersion	Update	Update the version of the VM.
VMVideo	Get, Set	Configure video settings for virtual machines.

Chapter 7

Storage

Storage spaces and storage pools

Storage Replica

SMB 3.1.1

iSCSI

iSNS server

Scale-Out File Servers

Server for NFS

Deduplication

Storage Quality of Service

ReFS

Storage-related PowerShell cmdlets

Windows Server 2019 provides several software-defined storage solutions that allow you to have complex storage configurations managed by the Windows Server

Technet24

operating system, rather than having to manage storage solutions using storage hardware–specific vendor management tools. Many of the storage service solutions available in Windows Server are used by Microsoft in their Azure datacenters where storage provisioning, performance, resiliency, and redundancy are all managed through software rather than through specific hardware devices.

STORAGE SPACES AND STORAGE POOLS

Storage spaces and storage pools were first introduced to Windows Server with the release of Windows Server 2012. A *storage pool* is a collection of storage devices that you can use to aggregate storage. You expand the capacity of a storage pool by adding storage devices to the pool. A storage space is a virtual disk that you create from the free space that is available in a storage pool. Depending on how you configure it, a storage space can be resilient to failure and have improved performance through storage tiering.

You manage storage spaces and storage pools through Windows Admin Center, PowerShell, the RSAT Tools, or Server Manager.

Inside OUT

Persistent memory support

Windows Server 2019 includes support for persistent memory, also termed "storage class memory." Examples include Intel Optane DC PM and NVDIMM-N modules.

Storage pools

A storage pool is a collection of storage devices—usually disks but can also include items such as virtual hard disks—from which you can create one or more storage spaces. A storage space is a special type of virtual disk that has the following features:

- **Resilient storage.** Configure to use disk mirroring or parity in the structure of the underlying storage (if available). Resilience is discussed later in this chapter.
- **Tiering.** Configure to leverage a combination of SSD and HDD disks to achieve maximum performance. Tiering is discussed later in this chapter.
- **Continuous availability.** Storage spaces integrate with failover clustering, and you can cluster pools across separate nodes within a single failover cluster.
- **Write-back cache.** If a storage space includes SSDs, a write-back cache can be configured in the pool to buffer small random writes. These random writes are then later offloaded to SSDs or HDDs that make up the virtual disk in the pool.
- **Multitenancy.** You can configure Access Control Lists on each storage pool. This allows you to configure isolation in multitenant scenarios.

To add storage to a storage pool, perform the following general steps:

1. In Server Manager, navigate to the Storage Pools node under the Volumes node of the File And Storage Services workspace. In this

workspace, you will see available storage that hasn't been allocated in the Primordial pool, as shown in Figure 7-1. The Primordial pool is the collection of disks that have not been provisioned.

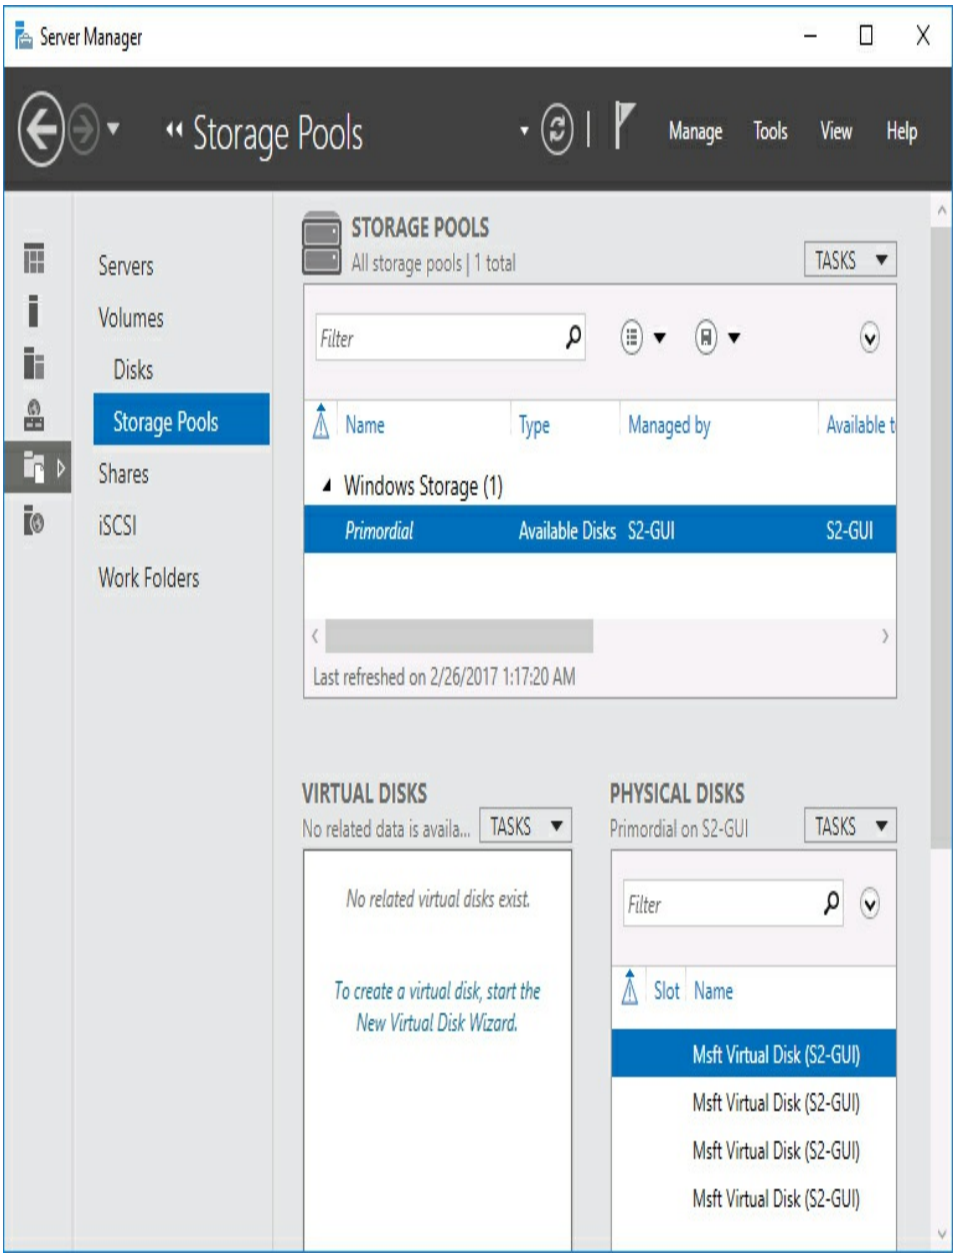


Figure 7-1 Primordial storage

2. On the Tasks menu, click New Storage Pool. You can also use the Tasks menu to rescan storage. You might need to use that option if you can't

see the storage in the Primordial pool that you want to add to a potential storage pool.

3. On the Storage Pool Name page of the New Storage Pool Wizard, specify a name for the storage pool and select which group of available disks you want to allocate to the storage pool. By default, the disks in the Primordial pool will be selected.
4. On the Physical Disks page, specify which available disks you want to add to the new pool. Figure 7-2 shows three disks being allocated using the Automatic allocation type and one disk being allocated as a Hot Spare type. Hot Spares will be swapped in if the existing disk in the storage pool fails.

Technet24

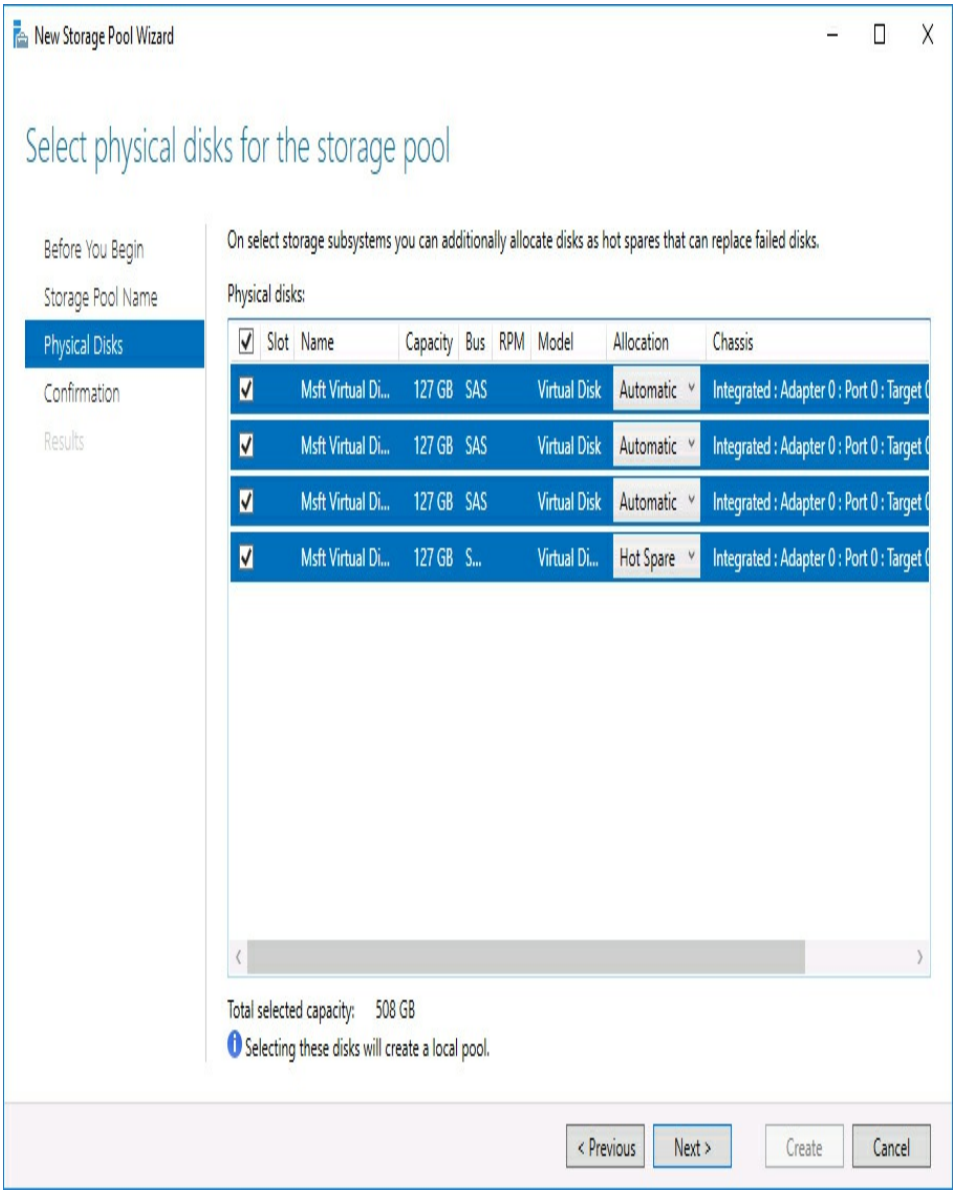


Figure 7-2 Adding disks

5. On the confirmation page, review your settings and click Create. When the pool is created, click Close.

You can add and remove disks from an existing storage pool. To add disks to an existing storage pool, perform the following steps:

1. In the Storage Pools node, verify that the disks that you want to add to the pool are available in the Primordial pool.
2. Right-click the storage pool where you wish to add the additional disks and click Add Physical Disk.
3. In the Add Physical Disk dialog box, select the disks from the Primordial pool that you wish to add to the existing storage pool and specify which Allocation type they will use. You can choose between Automatic, Hot Spare, and Manual. Figure 7-3 shows two disks being added using the Automatic allocation type.

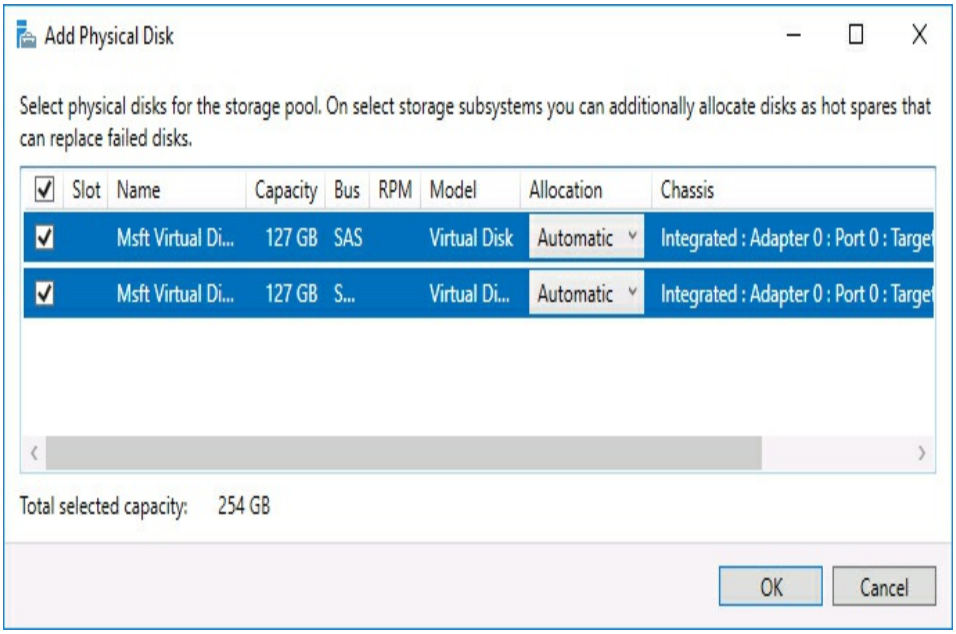


Figure 7-3 Adding disks to the existing pool

To remove a disk from an existing storage pool, perform the following steps:

1. In the Storage Pools node, select the storage pool from which you want to remove the disk.
2. In the list of physical disks, select the disk that you want to remove. You can right-click the disk within the pool and select the Toggle Drive Light option, shown in Figure 7-4, as a way of ensuring that you are removing the correct disk.

Technet24

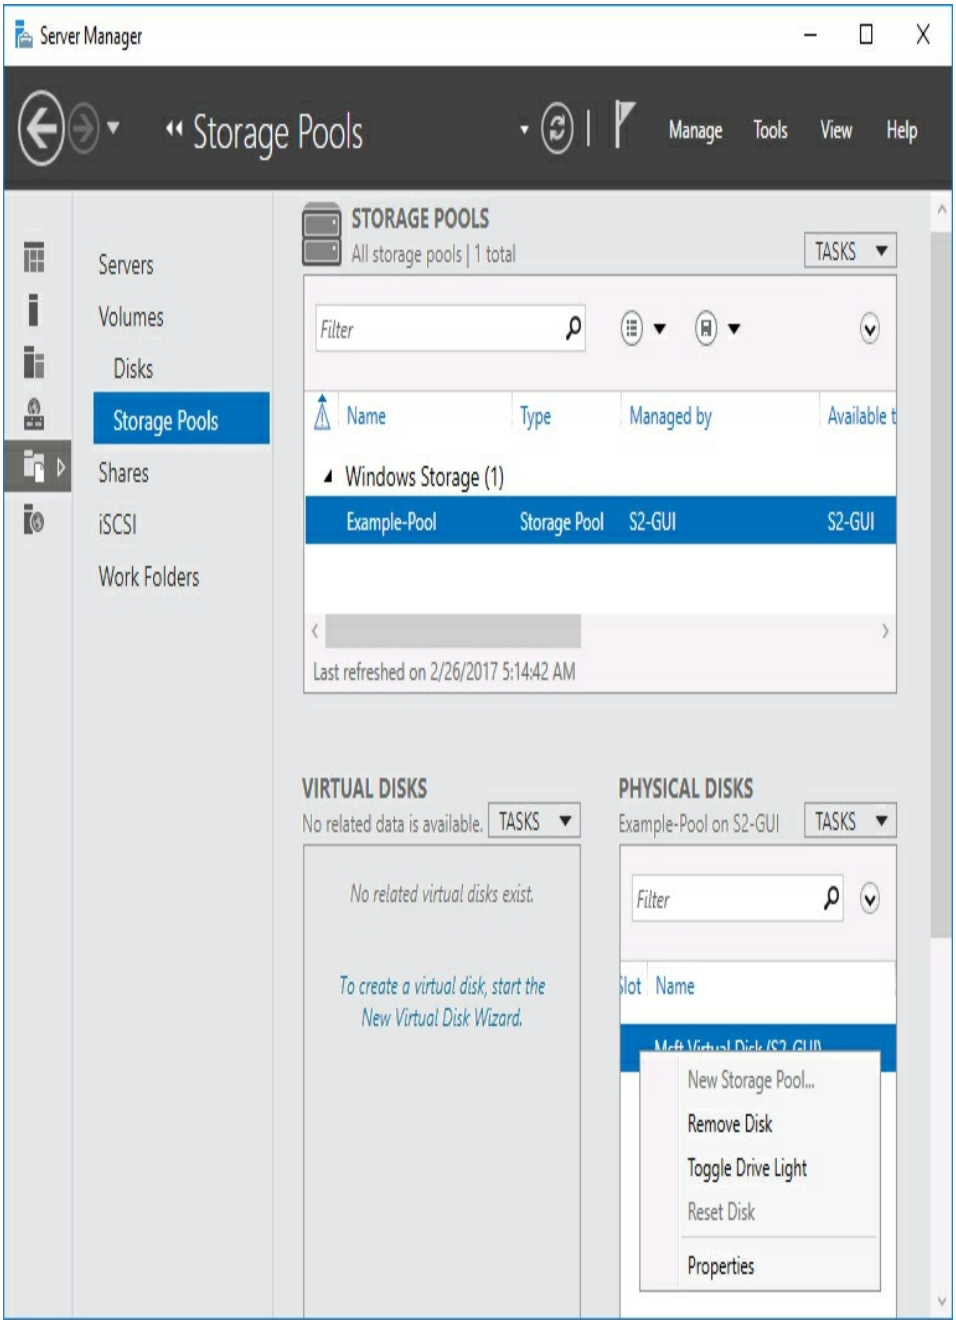


Figure 7-4 Toggle Drive Light

3. Once you are ready to remove the appropriate disk from the pool, right-click the disk, and select the Remove Disk option.
4. You will be presented with a warning message informing you that the operating system will attempt to rebuild virtual disks that store data on

the disk that you are removing from the pool. This can only happen if there are enough disks remaining in the pool to store the data that is currently stored on the physical disk that has been removed from the pool.

More Info

Storage spaces

To learn more about storage spaces, consult the following link: https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/overview.

Storage space resiliency

When creating a virtual disk on a storage pool that has enough disks, you can choose between several Storage Layouts, as shown in Figure 7-5.

Technet24

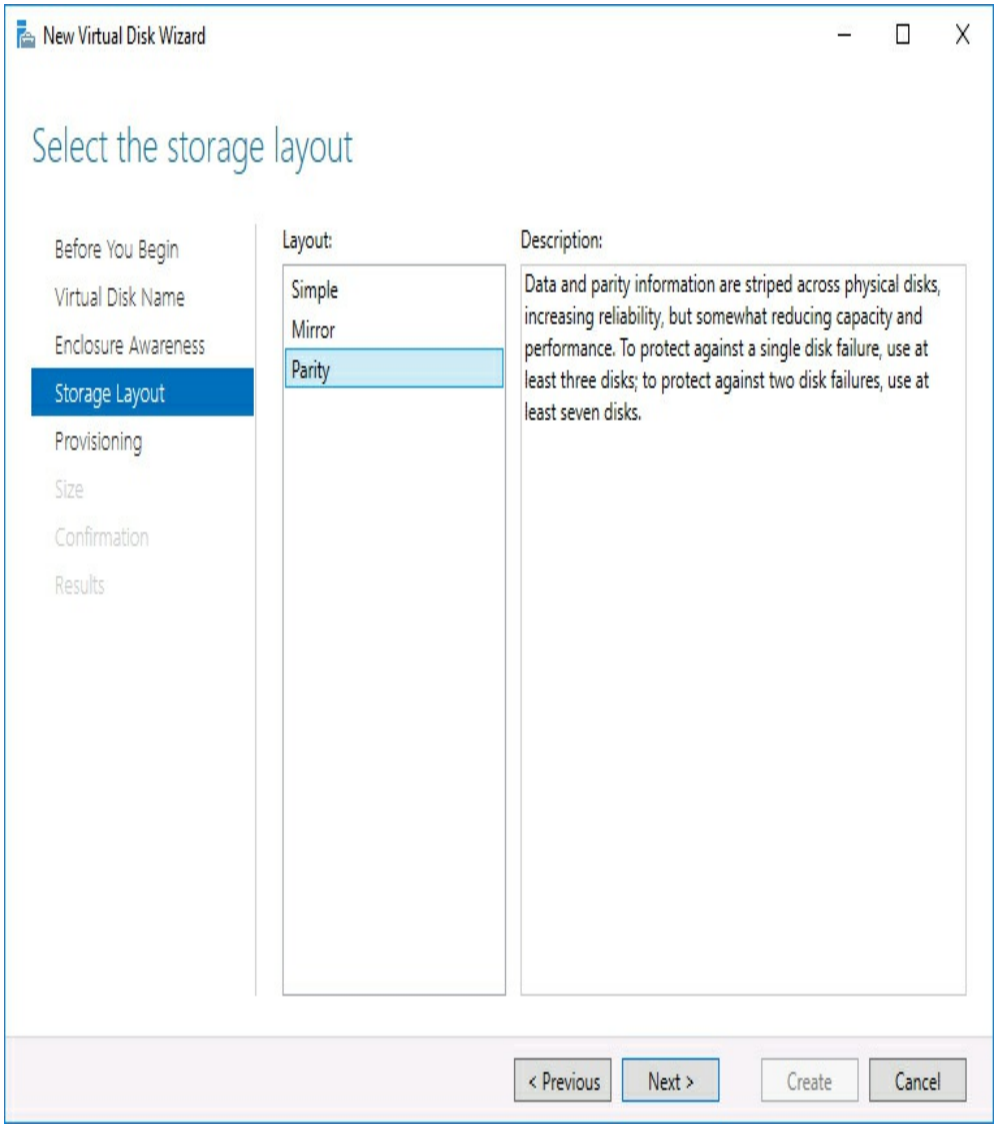


Figure 7-5 Choosing the Storage Layout

These options provide the following benefits:

- **Mirror.** Multiple copies of data are written across separate disks in the pool, which protects the virtual disk from failure of the physical disks that constitute the storage pool. Mirroring can be used with storage tiering. Depending on the number of disks in the pool, Windows Server 2019 storage spaces provide two-way or three-way mirroring. Two-way mirroring writes two copies of data, and three-

way mirroring writes three copies of data. Three-way mirroring provides better redundancy, but it also consumes more disk space.

- **Parity.** Similar to RAID 5, parity data is stored on disks in the array that are separate from where data is stored. Parity provides greater capacity than using the Mirror option, but it has the drawback of slower write performance. Windows Server 2019 provides two types of parity: Single Parity and Dual Parity. Single Parity provides protection against one failure at a time. Dual Parity provides protection against two failures at a time. You need to have a minimum of three disks for Single Parity and a minimum of seven disks for Dual Parity. You get the option to select between Single Parity and Dual Parity when configuring storage layout when there are more than seven disks.
- **Simple.** This option provides no resiliency for the storage, which means that if one of the disks in the storage pool fails, the data on any virtual hard disks built from that pool will also be lost.

If you configure disks in the storage pool that use the Hot Spare option, storage spaces will be able to automatically repair virtual disks that use the Mirror or Parity resiliency options. It's also possible for automatic repairs to occur if spare unallocated capacity exists within the pool.

Storage space tiering

Storage space tiering allows you to create a special type of virtual disk from a pool of storage that is a combination of SSD and traditional HDD disks. Storage tiering provides the virtual disk with performance similar to that of an array built out of SSD disks but without the cost of building a large capacity array

composed of SSD disks. It accomplishes this by moving frequently accessed files to faster physical disks within the pool, thus moving less-frequently accessed files to slower storage media.

You can only configure storage tiers when creating a virtual disk if there is a mixture of physical disks with the HDD and the SSD disk type in the pool upon which you want to create the disk. Once the disk is created, you cannot undo storage tiering from a virtual disk. You configure storage tiering for a virtual disk by selecting the Create Storage Tiers On This Virtual Disk option during virtual disk creation, as shown in Figure 7-6.

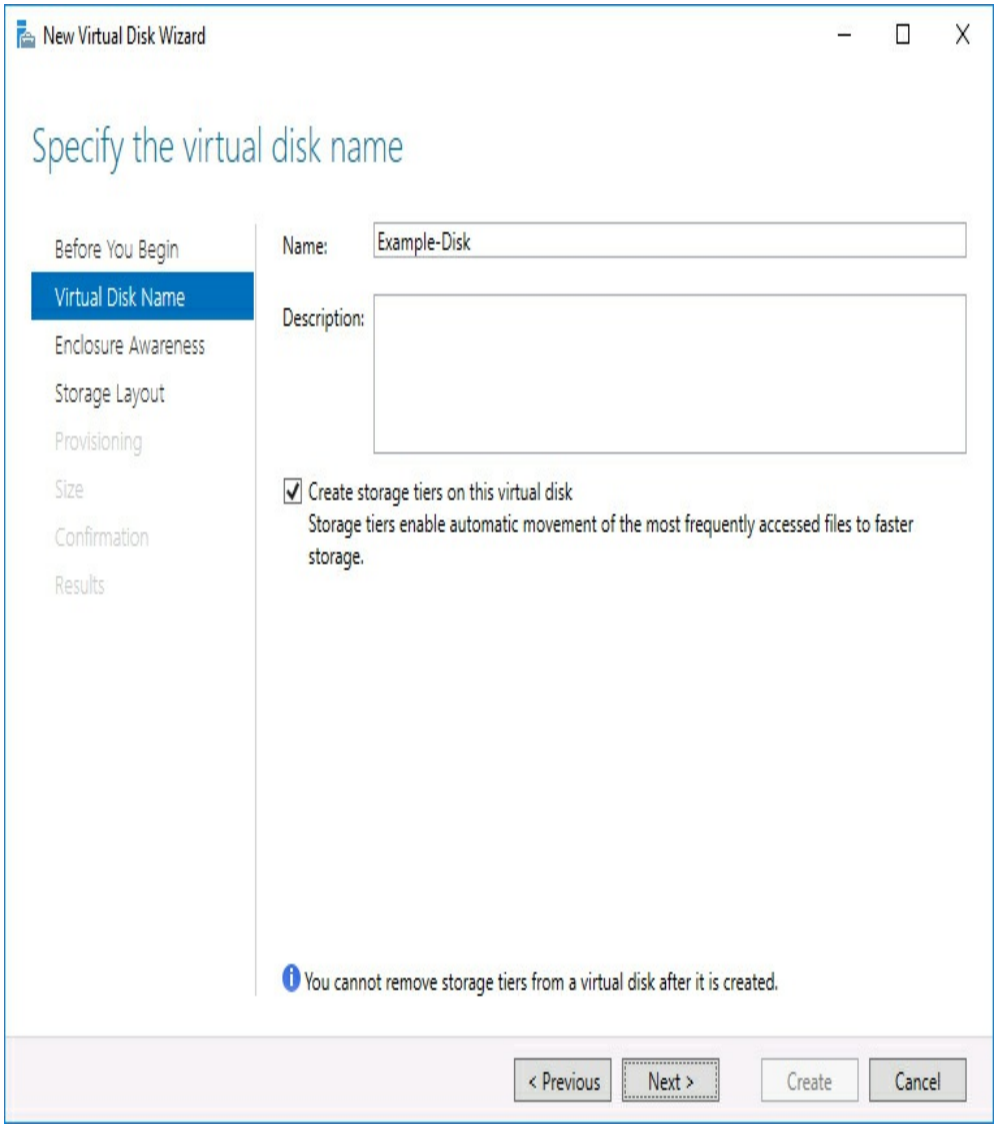


Figure 7-6 Enable storage tiering

One challenge when configuring storage tiering is ensuring that you have media marked as SSD and HDD in the pool. While media will usually be recognized correctly, in some cases you must specify that a disk is of the SSD type, which allows storage tiering to be configured.

You can specify the disk media type using the following PowerShell procedure:

1. First determine the storage using the *Get-StoragePool* cmdlet.
2. To view whether physical disks are configured as SSD or HDD, use the *Get-StoragePool* cmdlet and then pipe that cmdlet to the *Get-PhysicalDisk* cmdlet, as shown in Figure 7-7.

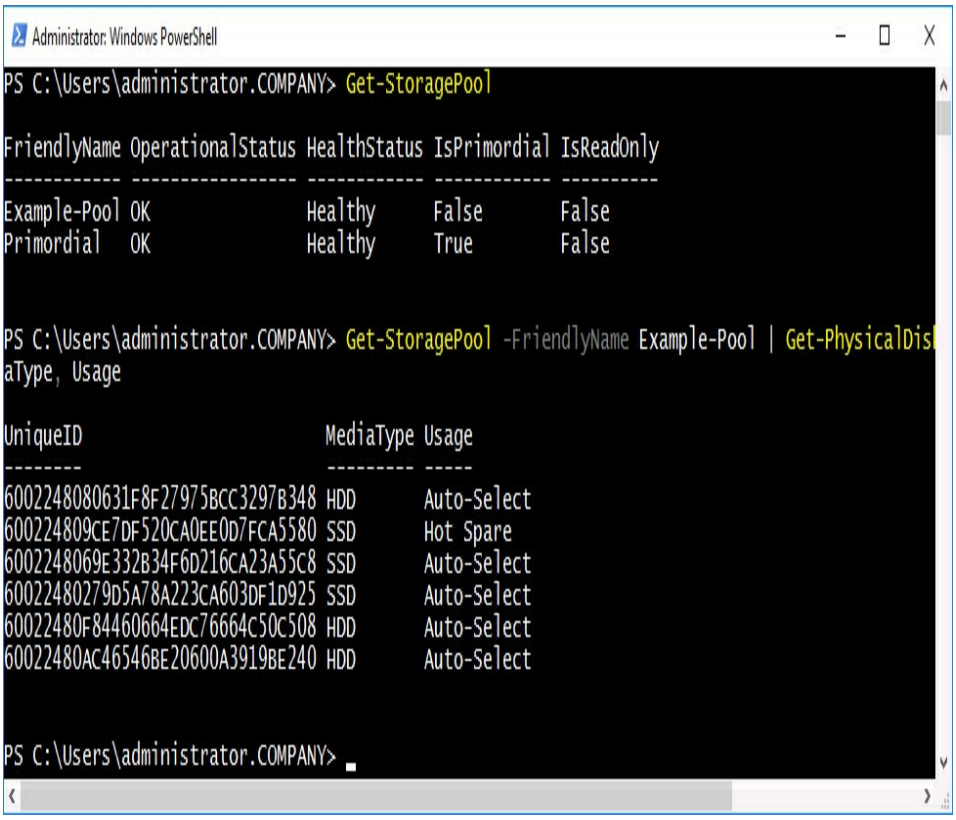


Figure 7-7 View disk media types

3. For example, to view the identity and media type of physical disks in the storage pool named *Example-Pool*, issue the command:

Click here to view code image

```
Get-StoragePool -FriendlyName ExamplePool |
Get-PhysicalDisk | Select UniqueID, MediaType,
Usage
```

4. Once you have determined the UniqueIDs of the disks that you want to configure as the SSD type, you can configure a disk to have the SSD type by using the *Set-PhysicalDisk* cmdlet with the *UniqueID* parameter and the *MediaType* parameter set to SSD. Similarly, you can change the type back to HDD by setting the *MediaType* parameter to HDD.

Thin provisioning and trim

Thin provisioning allows you to create virtual disks where you specify a total size for the disk, but only the space that is actually used will be allocated. For example, with thin provisioning, you might create a virtual hard disk that can grow to 500 GB in size but is only currently 10 GB in size because only 10 GB of data is currently stored on the volumes hosted on the disk. Figure 7-8 shows where you select the Provisioning Type in the New Virtual Disk Wizard.

Technet24

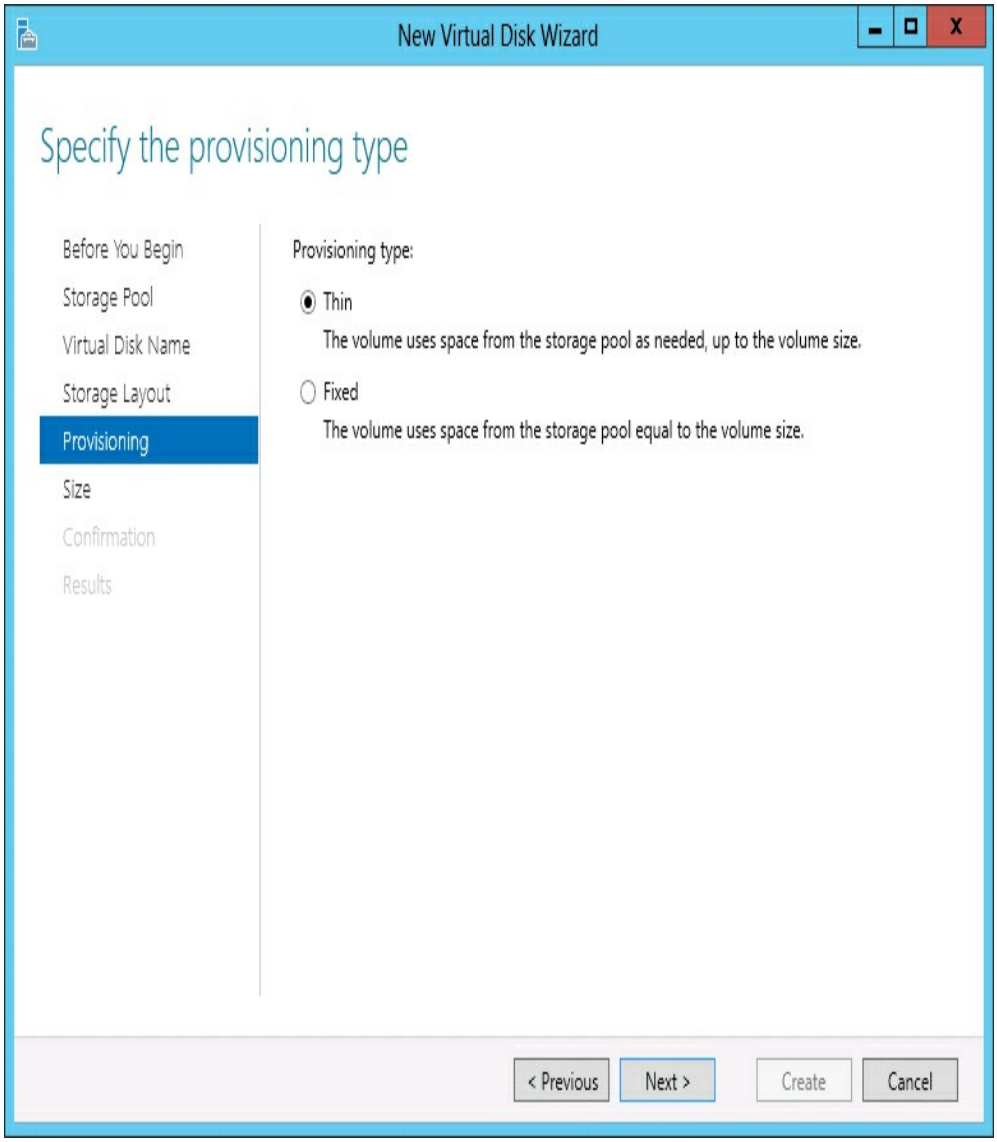


Figure 7-8 Thin provisioning

You can view the amount of space that has been allocated to a thin-provisioned virtual disk, and you can see the total capacity in the Virtual Disks area when the Storage Pools node is selected in the Server Manager console, as shown in Figure 7-9. When you create a

virtual disk, the maximum disk size available is determined by the amount of free space on the physical disks that make up the storage pool, rather than the maximum capacity of the existing thin-provisioned disks. For example, if you have a storage pool with two 10-TB physical disks, you can create more than two thin-provisioned disks that have a maximum size of 10 TB. You can create thin-provisioned disks 10 TB in size as long as the actual allocated space on the storage pool doesn't exceed 10 out of the 20 available. It is possible to create thin-provisioned disks in such a way that the total thin-provisioned disk capacity exceeds the storage capacity of the underlying storage pool. If the space is overallocated, you'll need to monitor how much of the underlying storage pool capacity is consumed and add disks to the storage pool because that capacity is exhausted.

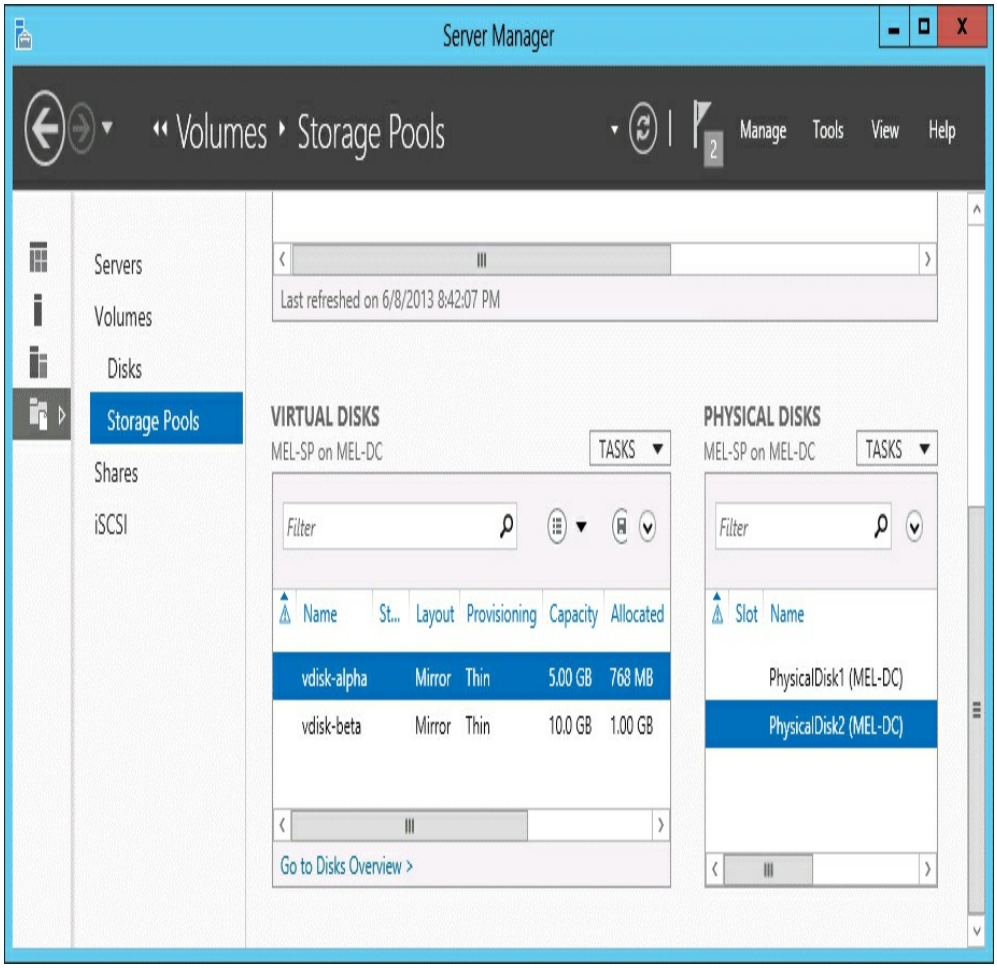


Figure 7-9 Capacity and allocation

Trim is an automatic process that reclaims space when data is deleted from thin provisioned disks. For example, if you have a 10-TB thin-provisioned virtual disk that stores 8 TB of data, 8 TB will be allocated from the storage pool that hosts that virtual disk. If you delete 2 TB of data from that thin-provisioned virtual disk, trim ensures that the storage pool that hosts that virtual disk can reclaim that unused space. The 10-TB

thin-provisioned virtual disk will appear to be 10 TB in size, but after the trim process is complete, it will only consume 6 TB of space on the underlying storage pool. Trim is enabled by default.

More Info

Thin provisioning and trim storage

To learn more about thin provisioning, consult the following link: https://docs.microsoft.com/en-us/windows-hardware/drivers/storage/thin-provisioning.

Creating virtual disks

A storage space virtual disk is composed from capacity available in the storage pool. As long as there is available capacity, you can create multiple virtual disks from a single storage pool. To create a new virtual disk using an existing storage pool, perform the following steps:

1. In Server Manager, navigate to the Storage Pools node under the Volumes node of the File And Storage Services workspace. In the Tasks menu of the Virtual Disks area, select New Virtual Disk.
2. On the list of available storage pools, shown in Figure 7-10, select an available storage pool upon which to create a new virtual disk and click OK.

Technet24

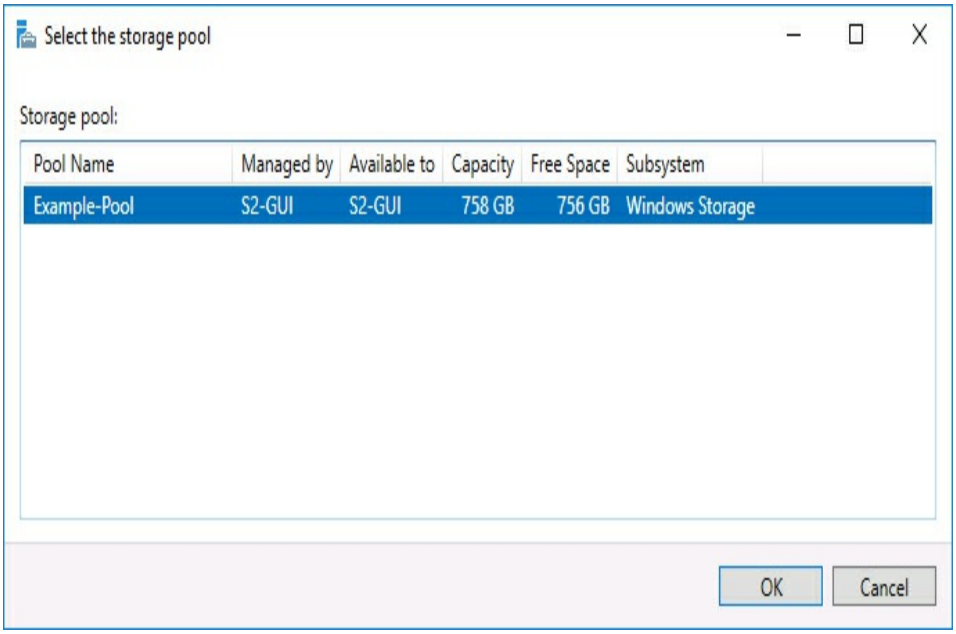


Figure 7-10 Select a storage pool

3. On the Enclosure Awareness page, select whether you want to enable Enclosure Awareness. Doing so will mean that copies of data are stored across separate storage enclosures. For enclosure awareness to work, you'll need multiple enclosures to host the disks that make up the storage pool.
4. On the Storage Layout page, you will be presented with the available layout options. Depending on the type and number of disks, you can choose between Simple, Mirror, and Parity.
5. On the Provisioning Type page, you can choose between thin and fixed provisioning.
6. On the Specify The Size of the Virtual Disk page, shown in Figure 7-11, you can choose a size for the disk. If you choose the Thin provisioning option, which you saw in Figure 7-8, the maximum amount of space in the storage pool is not listed. If you chose the Fixed option (back in Figure 7-8), the total amount of storage currently available in the pool is shown.

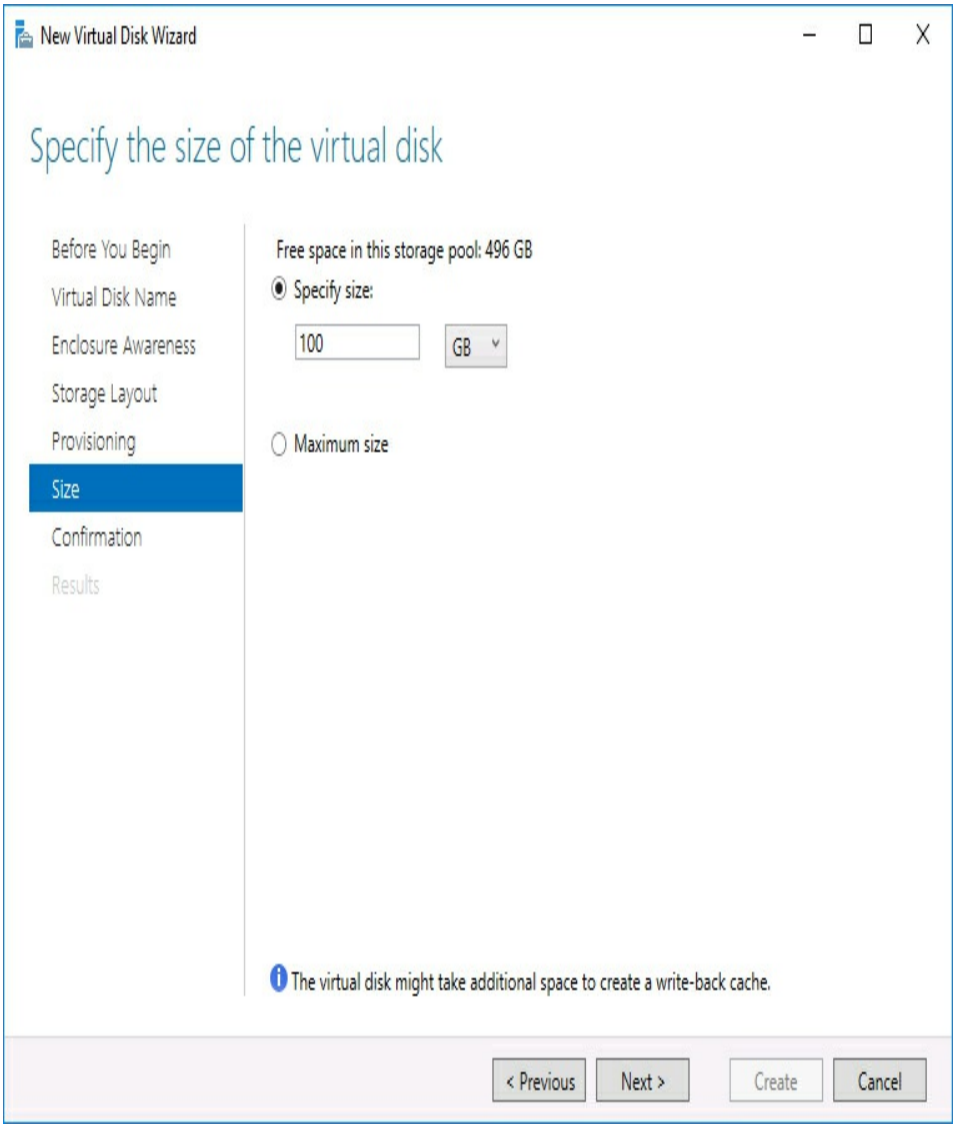


Figure 7-11 Specify the Virtual Disk size

7. On the Confirmation page, verify that the settings match what you intend and click Create.

Storage Spaces Direct

Storage Spaces Direct allows you to use Windows Server with locally attached storage to create highly available software-defined storage. Storage Spaces Direct (which

uses the S2D abbreviation because the SSD abbreviation is already used for solid-state disks) provides a form of distributed, software-defined, shared-nothing storage that has similar characteristics to RAID in terms of performance and redundancy. S2D allows you to create volumes from a storage pool of physical drives that are attached to multiple nodes that participate in a Windows Server failover cluster. Storage Spaces Direct functions as a replacement for expensive large-scale hardware storage arrays. S2D is only available with the Datacenter Edition of Windows Server.

> **Inside OUT**
>
> **Supported hardware**
>
> Although you can probably figure out a way to get an S2D cluster up and running on hardware that isn't officially supported, it's a terrible idea if you're planning to handle production workloads and data. Check with your vendor to determine whether its hardware includes support for S2D.

Storage Spaces Direct has the following properties:

- You can scale out by adding additional nodes to the cluster.
- When you add a node to a cluster configured for Storage Spaces Direct, all eligible drives on the cluster node will be added to the Storage Spaces Direct pool.
- You can have between 2 and 16 nodes in a Storage Spaces Direct failover cluster.
- It requires each node to have at least two solid state drives and at least four additional drives.
- A cluster can have more than 400 drives and can support more

than 4 petabytes of storage.

- Storage spaces direct works with locally attached SATA, SAS, persistent memory, or NVMe drives.
- Cache is automatically built from SSD media. All writes up to 256 KB and all reads up to 64 KB will be cached. Writes are then de-staged to HDD storage in optimal order.
- Storage Spaces Direct volumes can be part mirror and part parity. To have a three-way mirror with dual parity, it is necessary to have four nodes in the Windows Server failover cluster that hosts Storage Spaces Direct.
- If a disk fails, the plug-and-play replacement will automatically be added to the storage spaces pool when connected to the original cluster node.
- A storage spaces direct cluster can be configured with rack and chassis awareness as a way of further ensuring fault tolerance.
- Storage Spaces Direct clusters are not supported where nodes span multiple sites.
- While NTFS is supported for use with S2D clusters, ReFS is recommended.

S2D supports two deployment options:

- **Hyper-Converged.** With the Hyper-Converged deployment option, both storage and compute resources are deployed on the same cluster. This has the benefit of not requiring you to configure file server access and permissions and is most commonly used in small to medium sized Hyper-V deployments.
- **Converged.** With the converged (also known as “disaggregated”) deployment option, storage and compute resources are deployed in separate clusters. Often used with Hyper-V Infrastructure as a Service (IaaS) deployments, a scale-out file server is deployed on S2D to provide network attached storage over SMB3 file shares. The compute resources for the IaaS virtual machines are located on a

separate cluster to the S2D cluster.

Storage Spaces Direct in Windows Server 2019 supports nested resiliency. Nested resiliency is a capability designed for two-server S2D clusters that allows storage to remain available in the event of multiple hardware failures. When nested resiliency is configured, volumes can remain online and accessible even if one server goes offline and a drive fails. Nested resiliency only works when a cluster has two nodes. Nested resiliency requires a minimum of four capacity drives per server node and two cache drives per server node.

More Info

Storage Spaces Direct

To learn more about storage spaces direct, consult the following link: https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/storage-spaces-direct-overview.

S2D resiliency types

S2D resiliency options are dependent on how many fault domains are present, the failure tolerance required, and the storage efficiency that can be achieved. A fault domain is a collection of hardware, such as a rack of servers, where a single failure can affect every component in that collection.

Table 7-1 lists the different resiliency types, failure tolerances, storage efficiencies, and minimum fault domains.

Table 7-1 S2D resiliency

Resiliency	Failure tolerance	Storage efficiency	Minimum fault domains
Two-way mirror	1	50.0%	2
Three-way mirror	2	33.3%	3
Dual parity	2	50.0%-80.0%	4
Mixed	2	33.3%-80.0%	4

More Info

S2D Resiliency

To learn more about S2D resiliency, consult the following link: https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/storage-spaces-fault-tolerance.

Deploying S2D clusters

Deploying an S2D cluster is essentially the same as deploying a normal Windows Server failover cluster with the extra step of enabling S2D once the cluster has been created. To deploy an S2D cluster, perform the following general steps:

1. Ensure that the drives that you will use for S2D are empty and contain no existing partitions or data.

Technet24

2. Install the Failover Clustering, Hyper-V, File Server, Data Center Bridging, and roles and features; also, you install the Hyper-V and clustering PowerShell cmdlets on each cluster node by using the following command:

   Click here to view code image

   ```
   Install-WindowsFeature -Name "Failover-
   Clustering", "Hyper-V", "Data-Center-Bridging",
   "RSAT-Clustering-PowerShell", "Hyper-V-
   PowerShell", "FS-FileServer"
   ```

3. Use the *Test-Cluster* cmdlet with the *-Include Storage Spaces Direct*, *Inventory, Network*, and *System Configuration* parameters to verify that the cluster configuration is valid.
4. Use the *New-Cluster* cmdlet with the *-NoStorage* parameter to create the cluster. For example, run this command to create a new cluster named *TestS2DCluster* with machines *S2DN1* and *S2DN2*:

   Click here to view code image

   ```
   New-Cluster -Name TestS2DCluster -Node
   S2DN1,S2DN2 -NoStorage
   ```

5. Configure a cluster witness. This can be a cloud witness, a file share witness, or a disk witness.
6. Enable Storage Spaces Direct using the *Enable-ClusterStorageSpacesDirect* cmdlet. For example, to enable S2D on a cluster named *TestS2DCluster*, run the following command:

   Click here to view code image

   ```
   Enable-ClusterStorageSpacesDirect -CimSession
   TestS2DCluster
   ```

7. The next step is to create volumes and optionally enable the CSV cache.

More Info

S2D deployment

To learn more about S2D deployment, consult the following link: https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/deploy-storage-spaces-direct.

Managing S2D volumes

You can use Windows Admin Center to manage S2D cluster volumes. To add a volume to an S2D cluster, perform the following steps:

1. Use Windows Admin Center to connect to the S2D cluster by name.
2. Select Volumes in the Tools pane.
3. On the Volume page, click Inventory, as shown in Figure 7-12, and then click Create Volume.

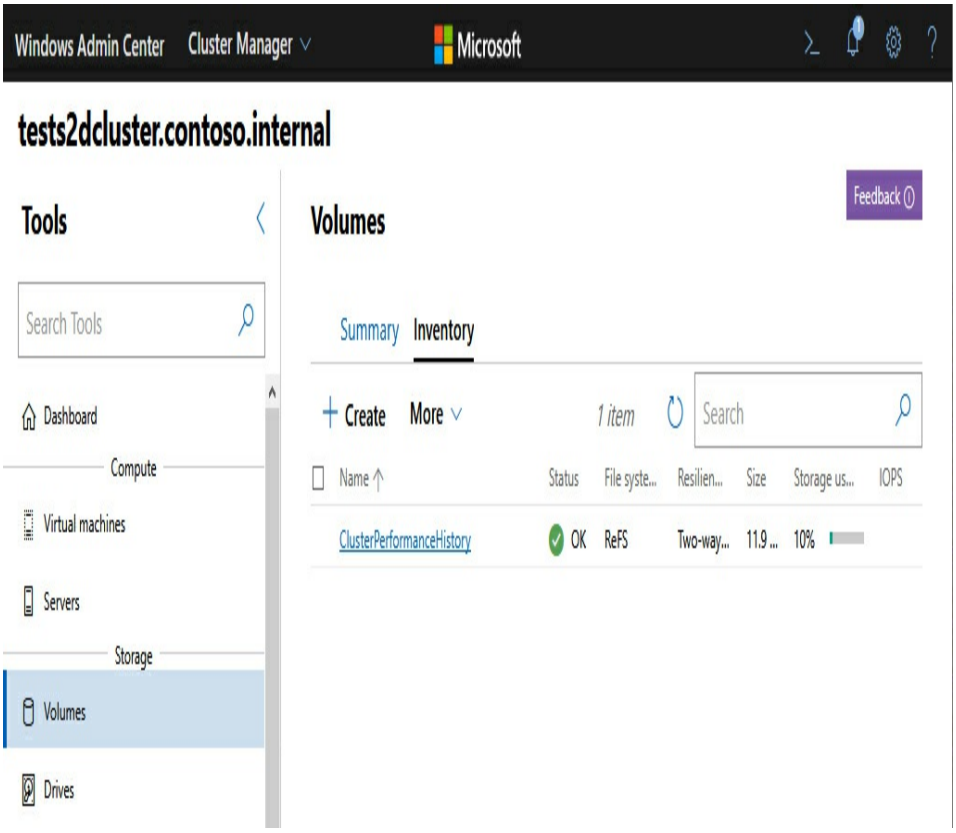


Figure 7-12 S2D cluster volume inventory

4. On the Create Volume page, shown in Figure 7-13, provide a Name for the volume, and then configure the volume Resiliency, volume size (Size On HDD), and whether options such as Deduplication will be enabled.

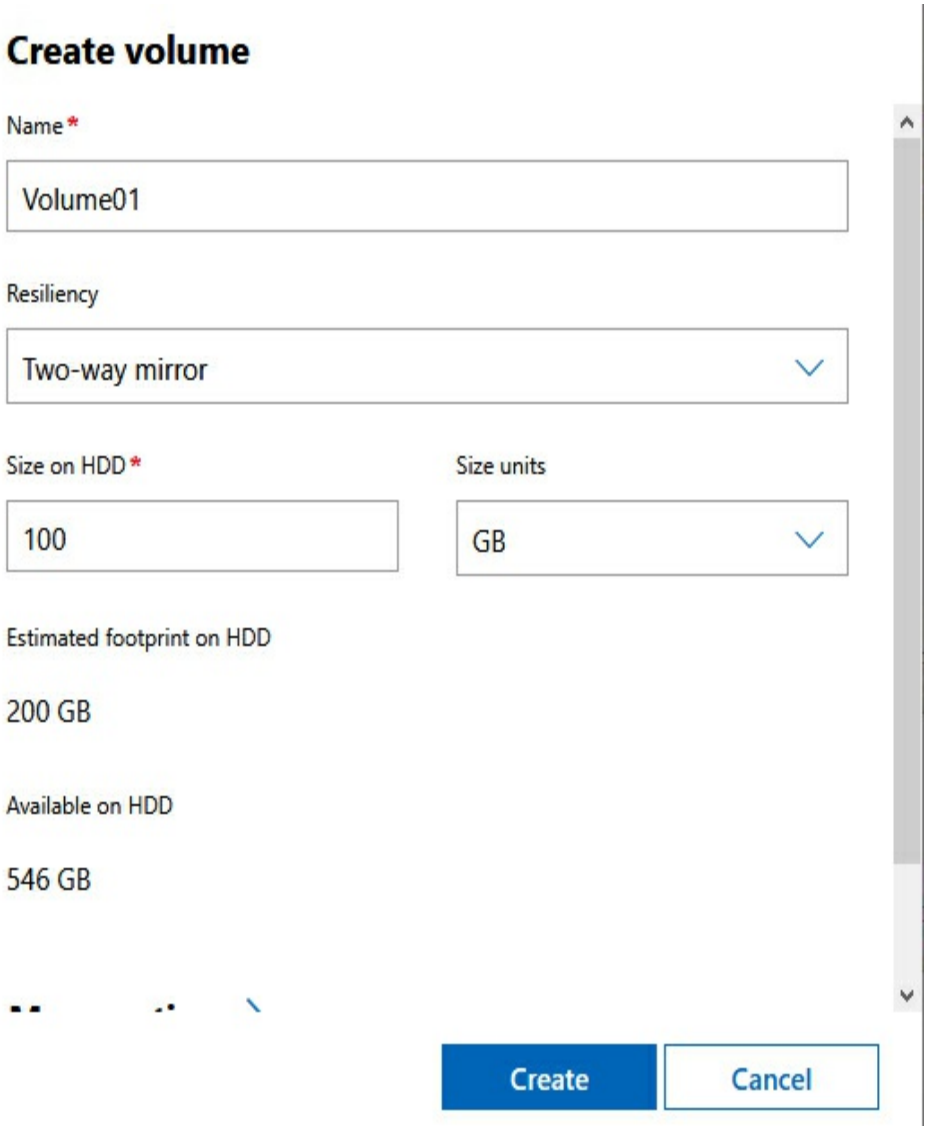


Figure 7-13 New S2D Volume

You might wish to make an S2D volume larger. You can do this by expanding the volume. To expand an S2D volume, perform the following steps:

1. Use Windows Admin Center to connect to the S2D cluster by name.
2. Select Volumes in the Tools pane and click the Volume.
3. On the Volume Properties page, shown in Figure 7-14, click Expand.

Technet24

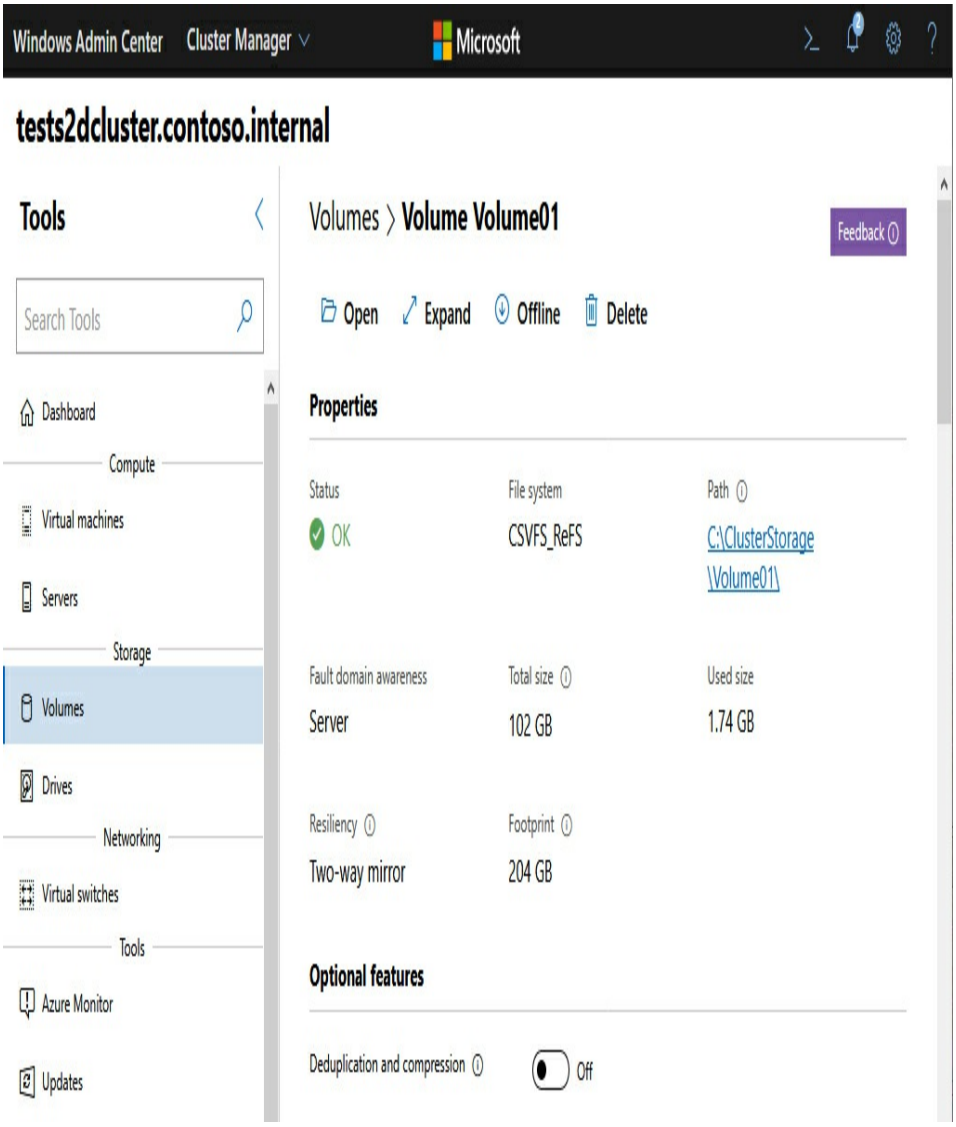


Figure 7-14 S2D Volume Properties

4. On the Expand Volume page, shown in Figure 7-15, specify the size to which you wish to expand the volume (up to the Available limit) and click Expand.

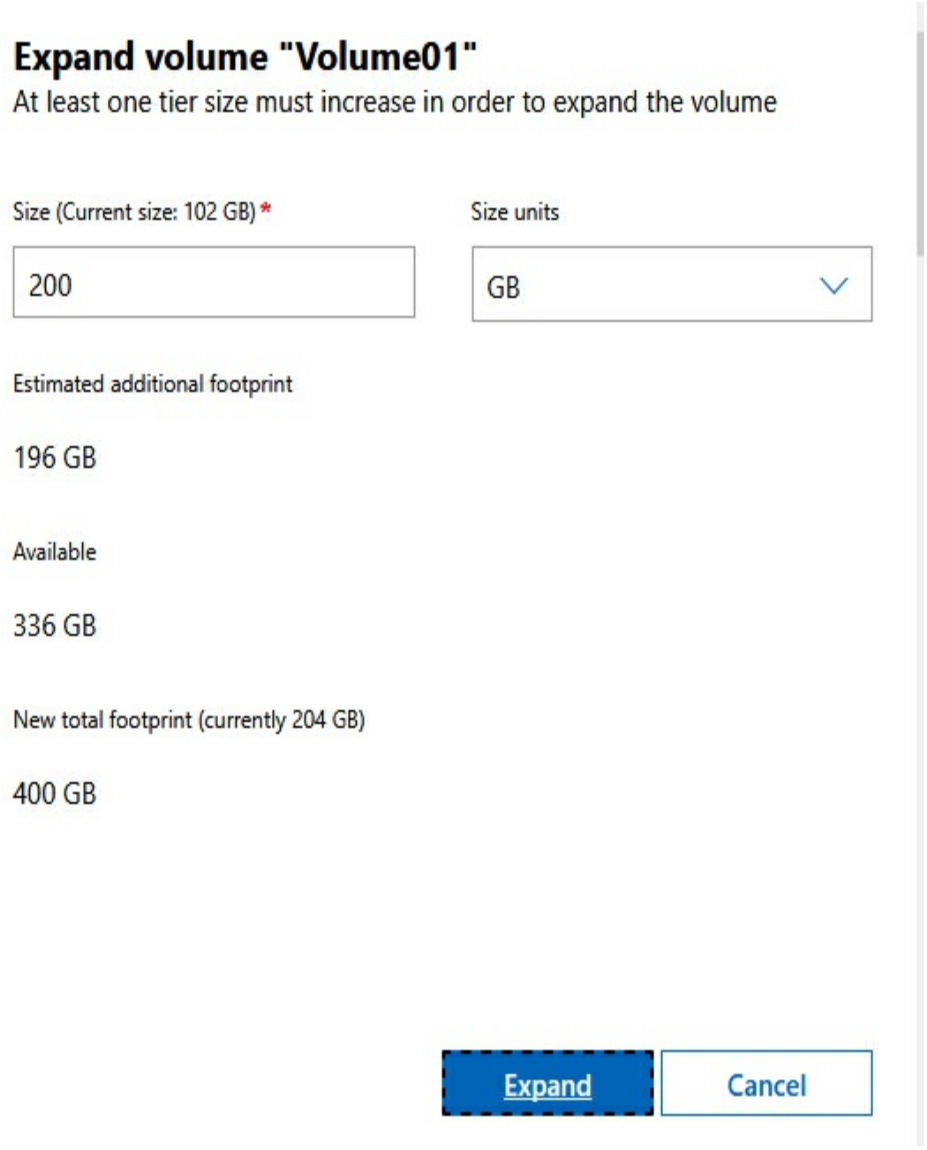


Figure 7-15 Expand the S2D volume

You can also use the Volume Properties page to take a volume offline or delete the volume.

More Info

Add S2D volumes

To learn more about adding S2D volumes, consult the following link:
https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/create-

volumes.

Adding S2D cluster nodes

Prior to adding a server to an existing S2D cluster, run the *test-cluster* cluster validation cmdlet, while including current and existing nodes. For example, run this command to add the cluster node S2DN3 to a cluster that included nodes *S2DN1* and *S2DN2*:

Click here to view code image

```
Test-Cluster -Name TestS2DCluster -node S2DN1,S2DN2,S2DN3 -
Include "Storage Spaces
Direct", "Inventory", "Network", "System Configuration"
```

Once the validation has completed, run the *Add-Clusternode* cmdlet on one of the existing cluster nodes and specify the new cluster node name.

More Info

Adding nodes to S2D clusters

To learn more about adding nodes to S2D clusters, consult the following link: https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/add-nodes.

S2D cluster node maintenance

Prior to performing maintenance on a cluster node, you should pause and drain the node. You can do this with the *Suspend-ClusterNode* cmdlet with the *-Drain* parameter. Once the node is drained, which will also put it in a paused state, you can either shut down the node

or perform other maintenance operations, such as restarting the node. When you have completed maintenance on the node, you can return it to operation by using the *Resume-Clusternode* cmdlet.

> More Info
>
> **S2D cluster node maintenance**
>
> To learn more about S2D cluster node maintenance, consult the following link: https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/maintain-servers.

STORAGE REPLICA

Storage Replica allows you to replicate volumes between servers, including clusters, for the purposes of disaster recovery. You can also use Storage Replica to configure asynchronous replication to provision failover clusters that span two geographically disparate sites, while all nodes remain synchronized.

You can configure Storage Replica to support the following types of replication:

- **Synchronous replication.** Use this when you want to mirror data and you have very low latency between the source and the destination. This allows you to create crash-consistent volumes. Synchronous replication ensures zero data loss at the file system level should a failure occur.
- **Asynchronous replication.** Asynchronous Storage Replica is suitable when you want to replicate storage across sites where you are experiencing higher latencies.

Storage Replica is available for single volumes under 2 TB in the Standard Edition of Windows Server 2019. These limits do not apply to the Datacenter Edition of the operating system. Participant servers must be members of the same Active Directory Domain Services Forest.

Storage Replica operates at the partition layer. This means that it replicates all VSS snapshots created by the Windows Server operating system or by backup software that leverages VSS snapshot functionality.

Storage Replica includes the following features:

- **Zero data loss, block-level replication.** When used with synchronous replication, there is no data loss. By leveraging block-level replication, even files that are locked will be replicated at the block level.
- **Guest and host.** Storage Replica works when Windows Server is a virtualized guest or when it functions as a host operating system. It is possible to replicate from third-party virtualization solutions to IaaS virtual machines hosted in the public cloud as long as Windows Server functions as the source and target operating system.
- **Supports manual failover with zero data loss.** You can perform manual failover when both the source and destination are online, or you can have failover occur automatically if the source storage fails.
- **Leverage SMB3.** This allows Storage Replica to use multichannel, SMB direct support on RoCE, iWARP, and InfiniBand RDMA network cards.
- **Encryption and authentication support.** Storage Replica

supports packet signing, AES-128-GCM full data encryption; it also provides support for Intel AES-NI encryption acceleration and Kerberos AES 256 authentication.

- **Initial seeding.** You can perform initial seeding by transferring data using a method other than Storage Replica between source and destination. This is especially useful when transferring large amounts of data between disparate sites where it may make more sense to use a courier to transport a high-capacity hard disk drive than it does to transmit data across a WAN link. The initial replication will then only copy blocks that have been changed since the replica data was exported from source to destination.
- **Consistency groups.** Consistency groups implement write ordering guarantees. This ensures that applications such as Microsoft SQL Server, which may write data to multiple replicated volumes, will have that data replicated such that it remains replicated sequentially in a consistent way.

Supported configurations

Storage Replica is supported in the following configurations:

- **Server-to-server.** In this configuration, Storage Replica supports both synchronous and asynchronous replication between two standalone services. Local drives, storage spaces with shared SAS storage, SAN, and iSCSI-attached LUNs can be replicated. You can manage this configuration either using Server Manager or PowerShell. Failover can only be performed manually.
- **Cluster-to-cluster.** In this configuration, replication occurs between two separate clusters. The first cluster might use Storage Spaces Direct, Storage Spaces with shared SAS storage, SAN, and iSCSI-attached LUNs. You manage this configuration using PowerShell and Azure Site Recovery. Failover must be performed manually.

Technet24

- **Stretch cluster.** A single cluster where nodes are located in geographically disparate sites. Some nodes share one set of asymmetric storage, and other nodes share another set of asymmetric storage. Storage is replicated either synchronously or asynchronously, depending on bandwidth considerations. This scenario supports Storage Spaces with shared SAS storage, SAN, and iSCSI-attached LUNs. You manage this configuration using PowerShell and the Failover Cluster Manager GUI tool. This scenario allows for automated failover.

The following configurations are not supported on Windows Server, though they may be supported at some point in the future:

- Storage Replica only supports one-to-one replication in Windows Server. You cannot configure Storage Replica to support one-to-many replication or transitive replication. Transitive replication is where there is a replica of the replica server.
- Storage Replica on Windows Server does not support bringing a replicated volume online for read-only access in Windows Server 2016. In Windows Server 2019, you can perform a test failover and temporarily mount a snapshot of the replicated storage on an unused NTFS- or ReFS-formatted volume.
- Deploying Scale-Out File Servers on stretch clusters participating in Storage Replica is not a supported configuration.
- Deploying Storage Spaces Direct in a stretch cluster with Storage Replica is not supported.

Configuring replication

To configure replication between two servers that are members of the same domain, perform the following steps:

1. Ensure that each server meets the following requirements:

1. Has two volumes: one volume hosts the data that you want to replicate; the other hosts the replication logs.
2. Ensure that both log and data disks are formatted as GPT and not MBR.
3. Data volume on source and destination server must be the same size and use the same sector sizes.
4. Log volume on source and destination volume should be the same size and use the same sector sizes. The log volume should be a minimum of 9 GB in size.
5. Ensure that the data volume does not contain the system volume, page file, or dump files.
6. Install the File Server role service and Storage Replica feature on each server as shown in Figure 7-16. Installing these features requires that you restart the computer.

Technet24

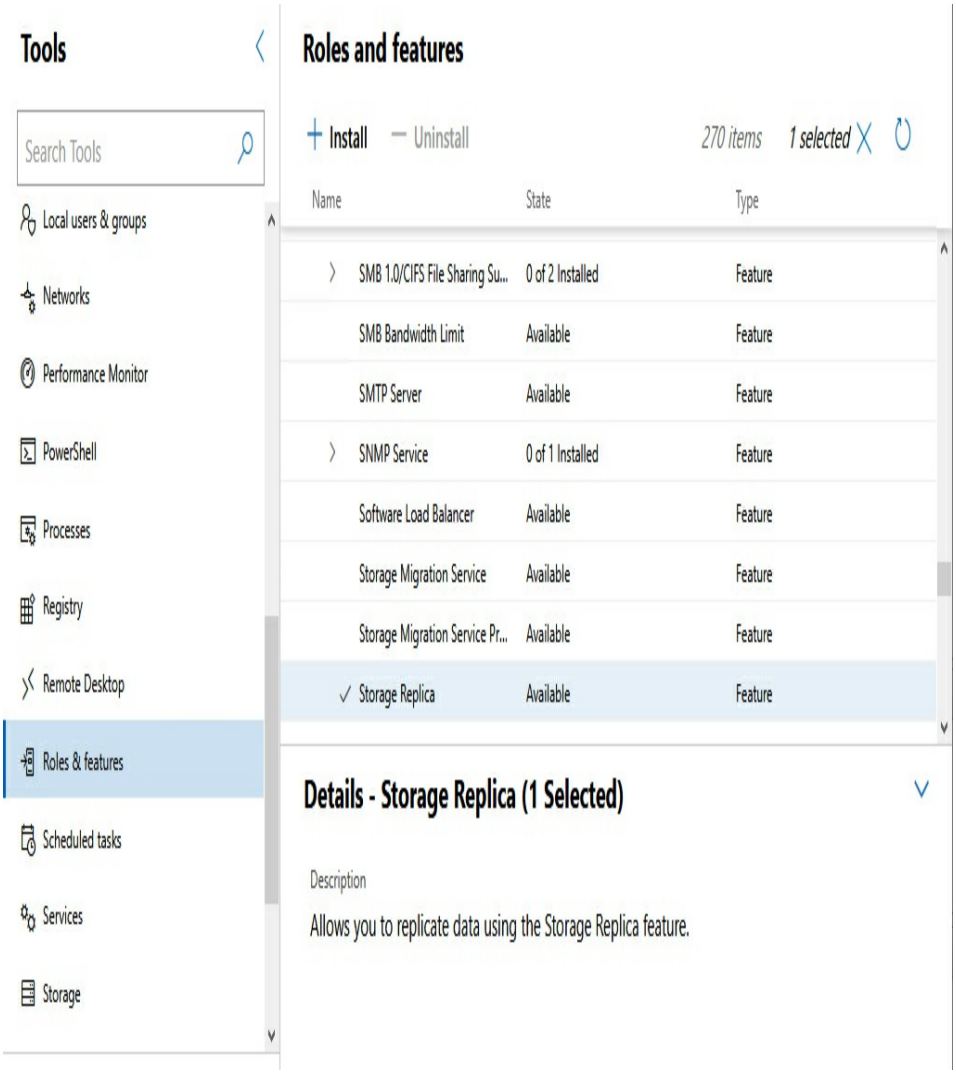


Figure 7-16 Install Storage Replica feature

2. Use the *Test-SRTopology* cmdlet to verify that all Storage Replica requirements have been met. To do this, perform the following steps:

    1. Create a temp directory on the source server that will store the Storage Replica Topology Report.
    2. Make a note of the drive letters of the source storage and log volumes and destination storage and log volumes.

The following command will test Storage Topology for a

source computer with the following details:

- It is named *S3-DCenter*:
- The source volume is E: and the log volume is F:
- The destination computer is named *S4-DCenter*.
- The volume is E: and the destination log volume is F:
- The duration is 15 minutes.
- The results are written to the *c:\temp* folder.

Click here to view code image

```
Test-SRTopology -SourceComputerName  S3-DCenter -SourceVolumeName
e:
-SourceLogVolumeName f: -DestinationComputerName S4-DCenter -
DestinationVolumeName e:
-DestinationLogVolumeName f: -DurationInMinutes 15 -ResultPath
c:\temp
```

When the test completes, view the *TestSrTopologyReport.html* file to verify that your configuration meets Storage Replica requirements. Once you have verified that the configuration does meet requirements, you can use the *New-SRPartnerShip* cmdlet to create a Storage Replica partnership.

For example, the following command will configure a Storage Replica group named *SR71* for a source computer named *S3-DCenter* with the source volume of E: and the log volume of F: to a destination computer named *S4-DCenter* with the destination volume of E: and destination log volume of F:

Click here to view code image

```
New-SRPartnerShip -SourceRGName SR71 -DestinationRGName SR72 -
SourceComputerName
S3-DCenter -SourceVolumeName e: -SourceLogVolumeName f: -
DestinationComputerName
S4-DCenter -DestinationVolumeName e: -DestinationLogVolumeName f:
```

You can check the status of Storage Replica replication by running the *Get-SRGroup*, *Get-SRPartnership*, and *(Get-SRGroup).replicas* cmdlets. Once Storage Replica is active, you won't be able to access the replica storage device on the destination computer unless you reverse replication or remove replication.

To switch direction, use the *Set-SRPartnerShip* cmdlet. For example, to reverse the replication direction of the Storage Replica partnership configured above, use the following command:

Click here to view code image

```
Set-SRPartnerShip -NewSourceComputerName S4-DCenter -SourceRGName
SR72
-DestinationComputerName S3-DCenter -DestinationRGName SR71
```

When you run this command, you will receive a warning that data loss may occur, and you will be asked whether you want to complete the operation.

To remove replication, run the following commands on the source node:

```
Get-SRPartnership | Remove-SRPartnership
Get-SRGroup | Remove-SRGroup
```

Then run the following command on the destination

node:

```
Get-SRGroup | Remove-SRGroup
```

By default, all replication when Storage Replica is configured is synchronous. You can switch between synchronous and asynchronous replication using the *Set-SRPartnership* cmdlet with the *ReplicationMode* parameter.

Storage Replica uses the following ports:

- **445.** Used for SMB, which is the replication transport protocol.
- **5985.** Used for WSManHTTP, which is the management protocol for WMI/CIM/PowerShell.
- **5445.** Used for SMB with iWARP. This port is only required if you are using iWARP RDMA networking.

More Info

Storage Replica

To learn more about Storage Replica, consult the following link: https://docs.microsoft.com/en-us/windows-server/storage/storage-replica/storage-replica-overview.

SMB 3.1.1

Windows Server 2016 and Windows Server 2019 use SMB (Server Message Block) 3.1.1. SMB 3.1.1 contains substantial security and performance improvements over previous versions of the SMB protocol. You can determine which version of SMB a server is using by

running the *Get-SmbConnection* cmdlet, as shown in Figure 7-17.

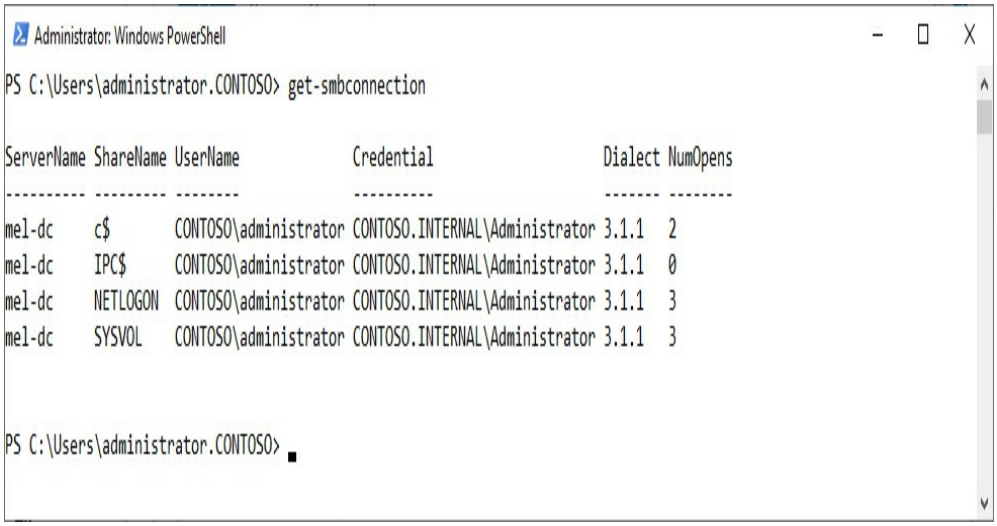


Figure 7-17 Check SMB version

The version of SMB that is used depends on the operating system of the client. Because newer versions of SMB perform much better and are much more secure than older versions, updating clients and server operating systems to the most recent version ensures that you are running the fastest version of SMB. Table 7-2 lists which version of SMB is used when clients or servers communicate with one another.

Table 7-2 Protocol negotiation

OS	Win 10 / WS 2016 / WS	Win 8.1 / WS	Win 8 / WS	Win 7 / WS 2008	Win Vista / WS	Older vers

	2019	2012 R2	2012	R2	2008	ions
Win 10 / WS 2016 / WS 2019	SMB 3.1.1	SMB 3.0.2	SMB 3.0	SMB 2.1	SMB 2.0.2	SMB 1.x
Win 8.1 / WS 2012 R2	SMB 3.0.2	SMB 3.0.2	SMB 3.0	SMB 2.1	SMB 2.0.2	SMB 1.x
Win 8 / WS 2012	SMB 3.0	SMB 3.0	SMB 3.0	SMB 2.1	SMB 2.0.2	SMB 1.x
Win 7 / WS 2008 R2	SMB 2.1	SMB 2.1	SMB 2.1	SMB 2.1	SMB 2.0.2	SMB 1.x
Win Vista / WS 2008	SMB 2.0.2	SMB 2.0.2	SMB 2.0.2	SMB 2.0.2	SMB 2.0.2	SMB 1.x
Older versions	SMB 1.x	SMB 1.x	SMB 1.x	SMB 1.x	SMB 1.x	SMB 1.x

SMB 1 is an old protocol that doesn't include key protections available in later versions of SMB, including pre-authentication integrity, secure dialect negotiation, encryption, insecure guest authentication blocking, and improved message signing. SMB1 is less efficient than

version SMB 3 or later, with SMB 3 supporting larger read and writes, peer caching of folder and file properties, durable handles, the client oplock leasing model, SMB direct, and directory leasing. SMB 1 is only necessary if your organization is running older versions of Windows that are no longer in mainstream support.

By default, SMB1 is disabled on Windows Server 2019. If your organization has servers running Windows Server 2016 or earlier, you'll need to configure the operating system so that SMB1 is not present. If you are worried that you do have applications on the network that use SMB 1, you can audit SMB 1 utilization by running the following command on any file server that has SMB 1 enabled and check the *SMBServer\Audit event log*:

Click here to view code image

```
Set-SmbServerConfiguration -AuditSmb1Access $true
```

Run the following command to remove SMB 1 from a computer running Windows Server 2016:

Click here to view code image

```
Remove-WindowsFeature -Name FS-SMB1
```

ISCSI

iSCSI allows access to storage devices across a TCP/IP network. The iSCSI initiator is a special software component that allows connections to iSCSI targets,

which are storage devices to which an iSCSI initiator connects over the network. When you install the iSCSI Target Server role service on a computer running the Windows Server operating system, you can configure virtual hard disks as iSCSI targets. When an iSCSI initiator connects to this virtual hard disk, it appears on the client as a local disk. While you can use iSCSI targets as local storage, you can also use an iSCSI-connected disk as shared storage in failover clusters.

By default, the iSCSI Initiator component is installed on all computers running supported versions of Windows Server. When you select the iSCSI Initiator for the first time in the Tools menu of the Server Manager console, it will prompt you to have the service start and to have the service be configured to start automatically in the future. Once the service is started, you can configure the iSCSI initiator on the iSCSI Initiator Properties dialog box, which is available when you select the iSCSI Initiator in the Tools menu of the Server Manager console. In most cases, the configuration involves entering the FQDN of the server that hosts the iSCSI target in the Target text box and clicking Quick Connect. The iSCSI Initiator Properties dialog box is shown in Figure 7-18.

Technet24

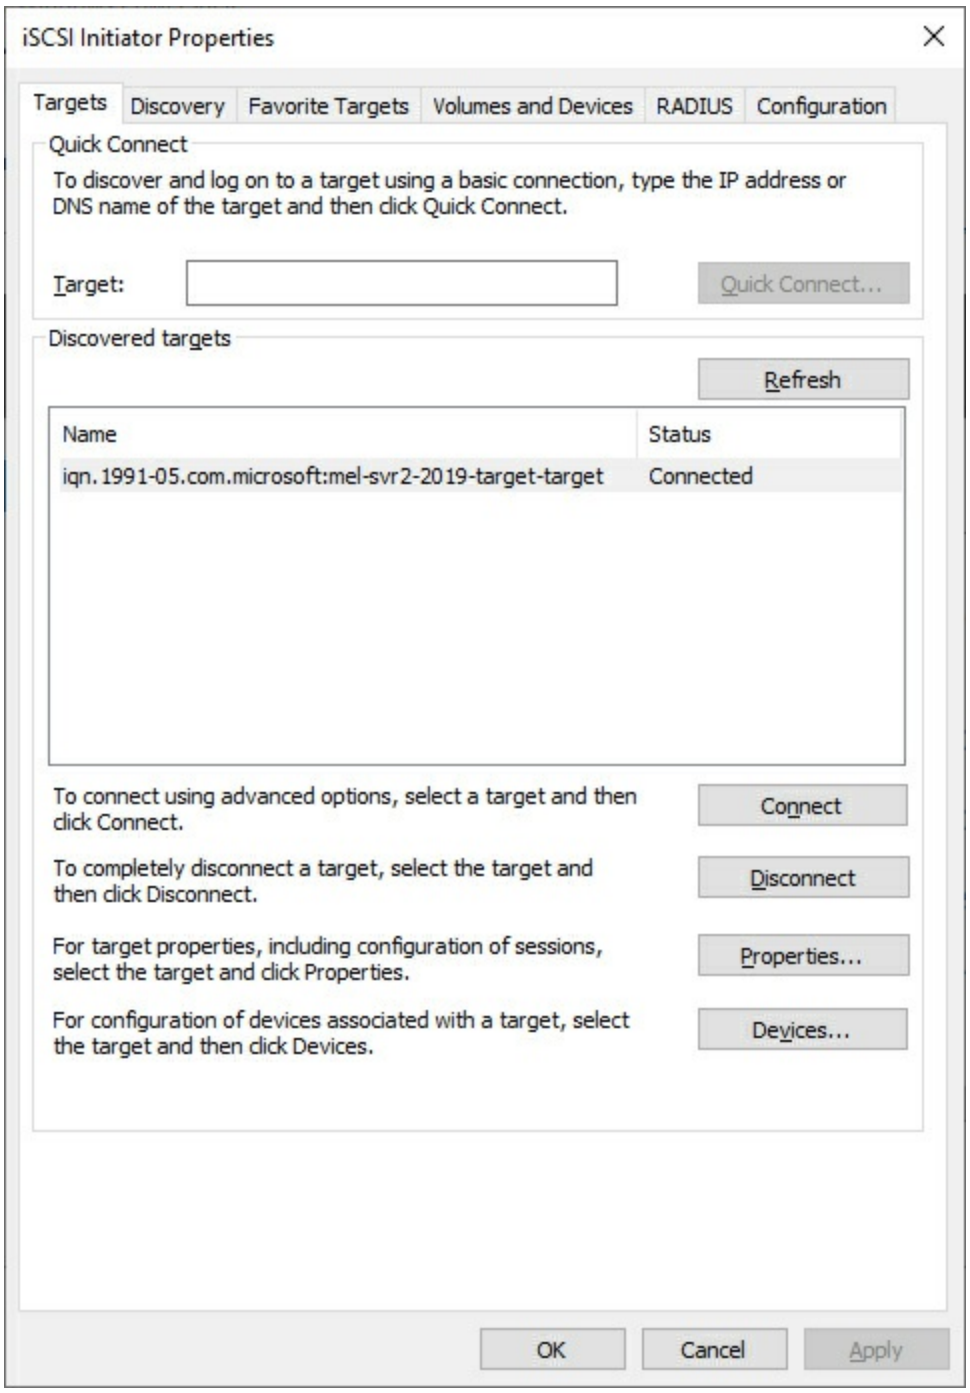


Figure 7-18 iSCSI Initiator Properties

You can only connect to an iSCSI target that has been

configured to accept connections from the iSCSI initiator to which you are attempting to connect. When creating an iSCSI target, you specify which iSCSI initiators can access the target, as shown in Figure 7-19. You can specify initiators on the basis of IQN, DNS Name, IP Address, or MAC address.

Technet24

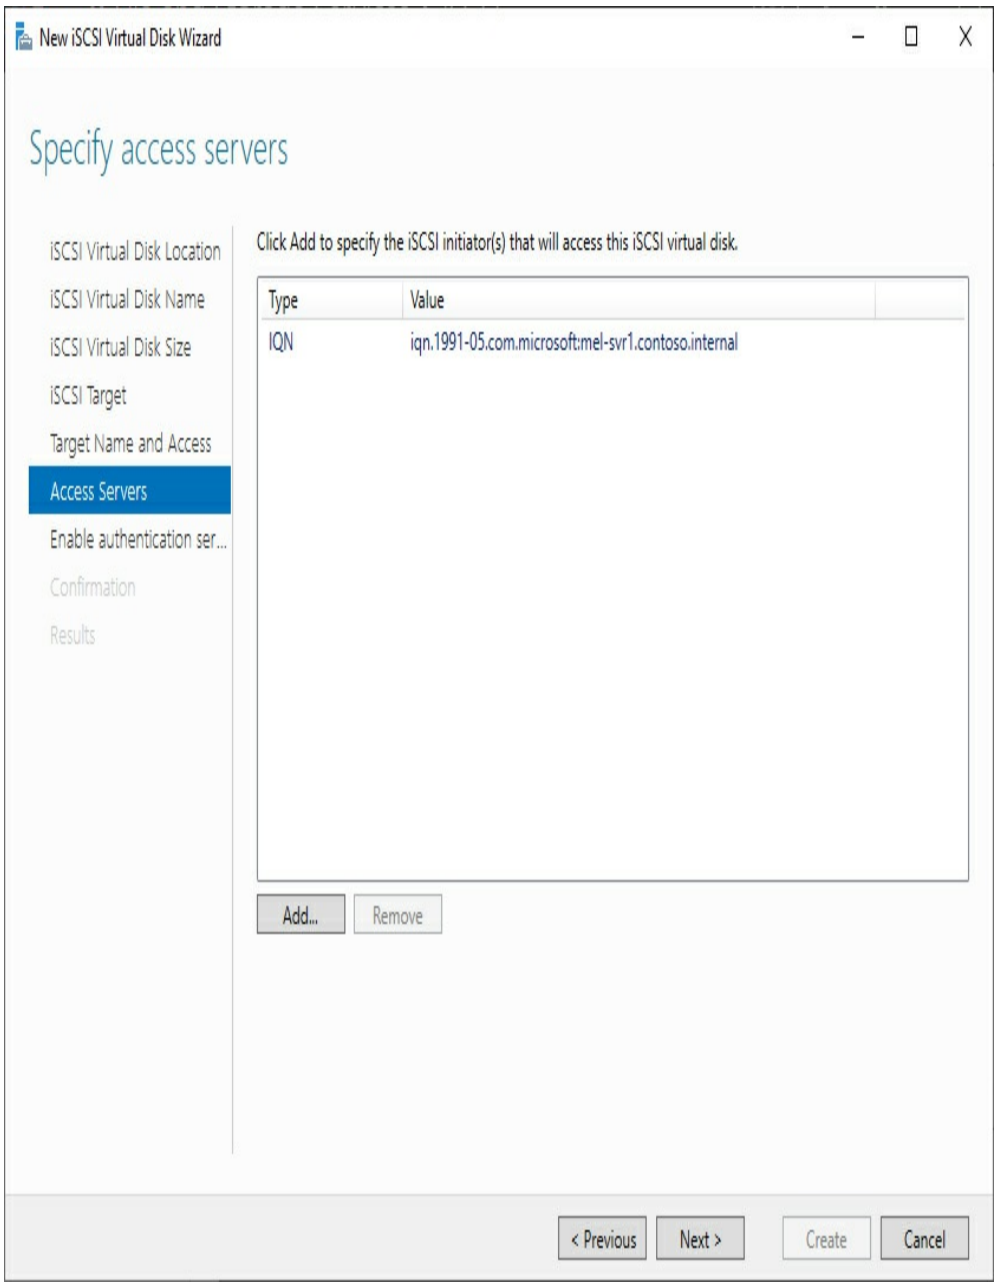


Figure 7-19 Configuring access servers

If the initiator is running on a domain-joined computer running a Windows 8 client operating system or later or a Windows Server 2012 server operating system or later,

you can query Active Directory to determine the initiator ID as shown in Figure 7-20. When configuring an initiator ID for a client running a Microsoft operating system, the format is *iqn.1991-05.com.microsoft:FQDN*. For example, if you are specifying the initiator ID of *cbr.contoso.internal*, you would type **iqn.1991-05.com.microsoft:cbr.contoso.internal**.

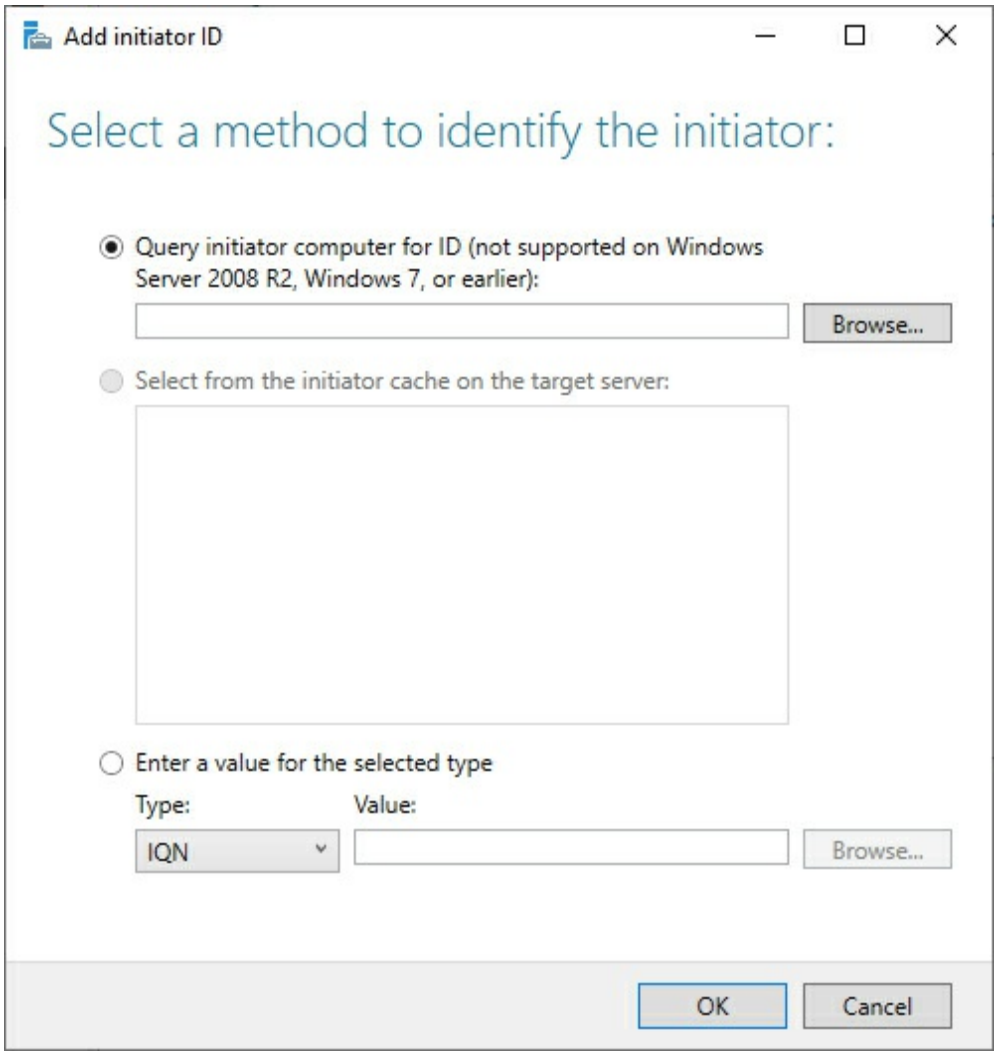


Figure 7-20 Identify the initiator

Technet24

You can create a new iSCSI virtual disk using the *New-IscsiVirtualDisk* cmdlet. For example, to create a new 10GB iSCSI virtual disk that has the path *E:\Disks\Disk1.vhd*, you can use this command:

Click here to view code image

```
New-IscsiVirtualDisk –Path “e:\disks\disk1.vhd” –Size 10GB
```

More Info

iSCSI Target

To learn more about iSCSI Target, consult the following link: https://docs.microsoft.com/en-us/windows-server/storage/iscsi/iscsi-target-server.

You can use the *New-IscsiServerTarget* cmdlet to create an iSCSI server target. For example, to configure an iSCSI target that allows computers *syd-a.contoso.com* and *syd-b.contoso.com* to access an iSCSI virtual disk, you can use this command:

Click here to view code image

```
New-IscsiServerTarget –Targetname “Syd-A-Syd-B-Target” –
InitiatorIDs
DNSName:Syd-a.contoso.com,DNSName:syd-b.contoso.com
```

ISNS SERVER

The iSNS (Internet Storage Name Service) server role allows you to centralize the discovery of iSCSI initiators and targets. Rather than having to enter the address of a server that hosts an iSCSI target when configuring the

iSCSI initiator, the iSNS server provides a list of available targets. You can use the iSNS server with both Microsoft and third-party iSCSI initiators and targets. Figure 7-21 shows the iSNS Server Properties dialog box and the addresses of two initiators that have registered with iSNS. Discovery domains allow you to partition initiators and targets into more manageable groups and are useful if your organization uses iSCSI extensively.

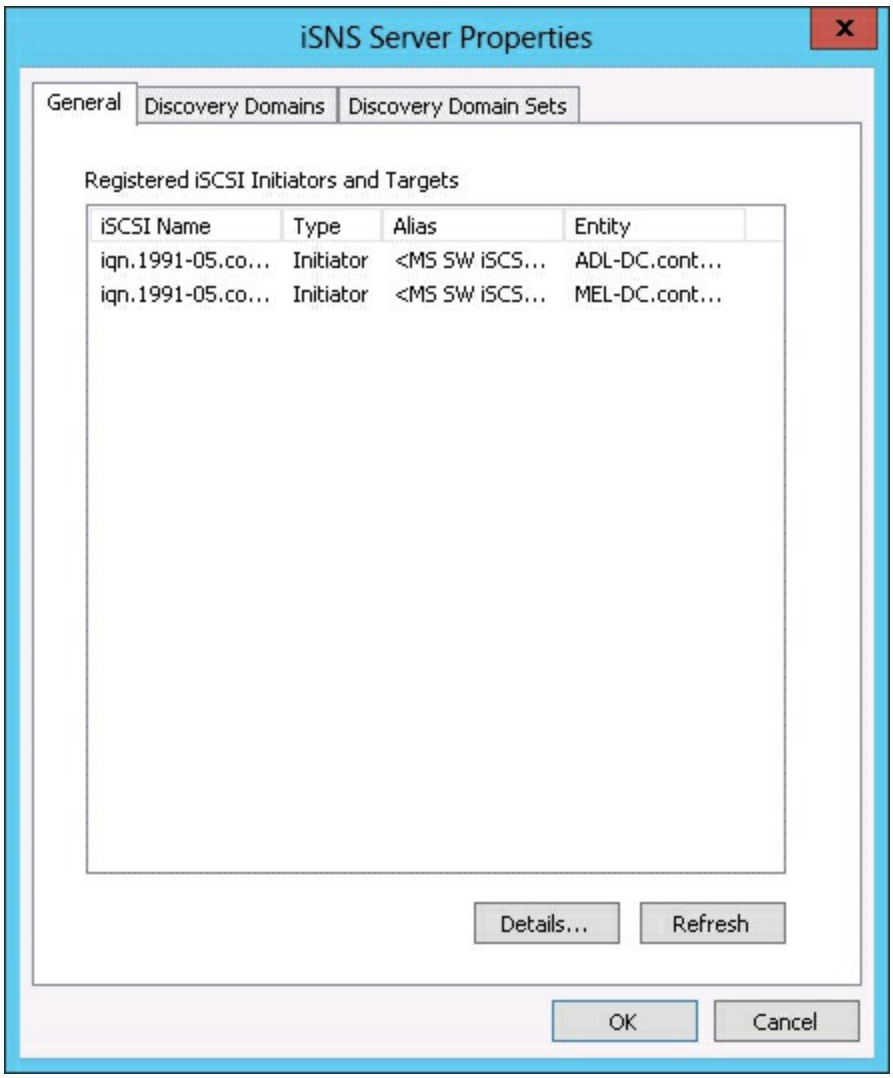


Figure 7-21 List of initiators

You can register an iSCSI Initiator running on Windows Server with an iSNS server through the Discovery tab of the iSCSI Initiators Properties dialog box, shown in Figure 7-22. Using an iSNS server to locate initiators is useful when you have iSCSI initiators running on Windows 7 or Windows Server 2008 R2 operating systems. These operating systems fell out of support in January 2020, but many organizations are still using them with extended support agreements.

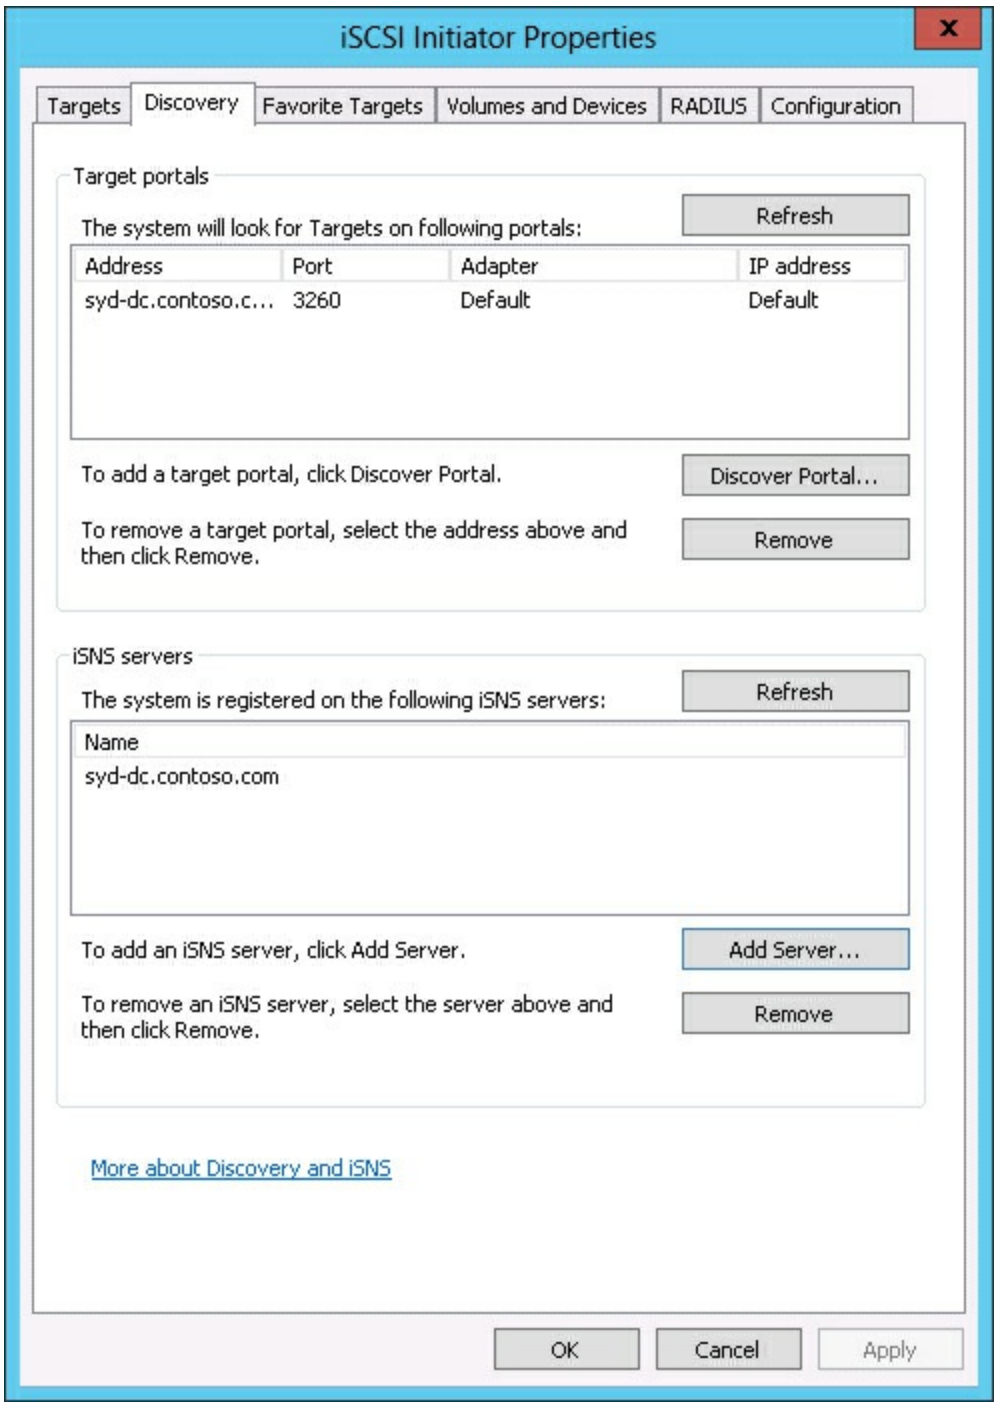


Figure 7-22 Configure iSNS server

With clients running Windows 8 and Windows Server

2012 and later operating systems, you can locate iSCSI IDs through Active Directory.

Windows Server does not support registering iSCSI targets through the GUI. You can register iSCSI targets associated with a virtual disk with an iSNS server using PowerShell. Use the following command to register all the iSCSI targets hosted on a computer running Windows Server with an iSNS server:

Click here to view code image

```
Set-WmiInstance –Namespace root\wmi –Class WT_iSNSServer –
Arguments
@{ServerName=”ISNSservername”}
```

More Info

iSNS Server

To learn more about the iSNS server role, consult the following link: https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/cc772568(v=ws.11).

SCALE-OUT FILE SERVERS

A Scale-Out File Server (SoFS) allows you to share a single folder from multiple nodes of the same cluster. You can use SoFS to deploy file shares that are continuously available for file-based server application storage. This storage is suitable for hosting Hyper-V virtual machine files or Microsoft SQL Server databases with a level of reliability, availability, manageability, and

performance that equates to what is possible with a storage area network.

Benefits of an SoFS deployment include:

- **Active-Active file share topology.** SoFS allows the same folder to be accessed from multiple nodes of the same cluster. An SoFS file share remains online should one or more cluster nodes fail or be taken down for planned maintenance.
- **Scalable bandwidth.** You can respond to a requirement for increased bandwidth by adding nodes.
- **Automatic rebalancing of clients.** SMB client connects are tracked on a per file share basis, with clients being redirected to the cluster node that has the best access to the storage device used by the file share.
- **CHKDSK with zero downtime.** The Cluster Shared Volume File System, used with SoFS, allows CHKDSK operations to occur without affecting applications that have open handles on the file system.

You should consider SoFS file shares for the following scenarios:

- Store Hyper-V configuration and live virtual disks
- Store live SQL Server database files
- Store shared IIS configuration data

SoFS has the following requirements:

- The storage configuration must be explicitly supported by failover clustering. This means that you must be able to successfully run the Cluster Validation Wizard prior to adding an SoFS.
- SoFS requires Cluster Shared Volumes.

Windows Server 2019 introduces a new SoFS role called

the Infrastructure File Server. An infrastructure SoFS uses a single namespace share for the Cluster Shared Volume drive. The benefit of the Infrastructure File Server role is that it allows a Hyper-V host to communicate using guaranteed continuous availability to the Infrastructure SoFS SMB server. A failover cluster can only support a single Infrastructure SoFS instance. To create an Infrastructure SoFS, run the following PowerShell command:

Click here to view code image

```
Add-ClusterScaleOutFileServerRole -Cluster ClusterName -
Infrastructure -Name
InfrastructureSoFSName
```

More Info

SoFS

To learn more about SoFS, consult the following link: https://docs.microsoft.com/en-us/windows-server/failover-clustering/sofs-overview.

SERVER FOR NFS

Server for NFS allows clients that use the Network File System (NFS) protocol to access data stored on computers running Windows Server. NFS is primarily used by UNIX and Linux clients, but it is also used by some third-party hypervisors. You can configure Server for NFS to leverage Windows-to-UNIX user mappings from Active Directory Domain Services, User Name

Mapping, or both. Identity Management for UNIX can also be deployed to update the Active Directory schema to support UNIX user account data.

Server for NFS supports continuous availability when deployed on a Windows Server 2019 failover cluster. Once installed, you configure Server for NFS by editing the Server For NFS Properties node of the Services For Network File System tool, which is available from the Tools menu of the Server Manager console when Server for NFS is installed. This dialog box is shown in Figure 7-23.

Technet24

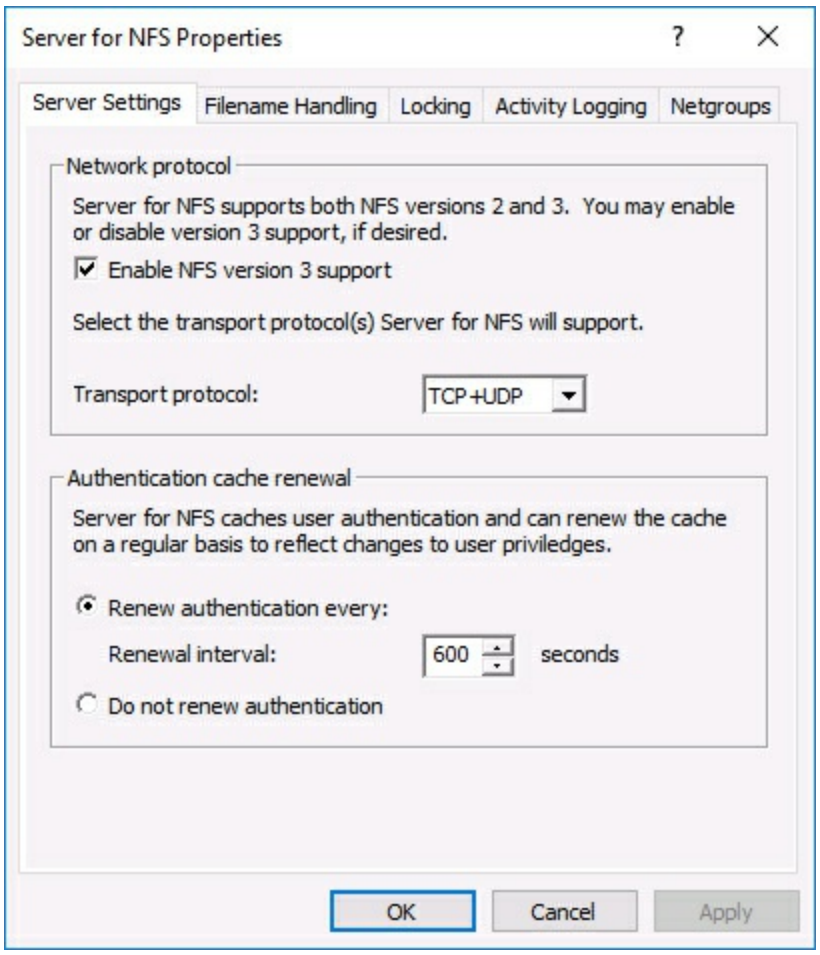


Figure 7-23 Server For NFS Properties dialog box

You can use this dialog box to configure the following:

- Support for NFS version 3.
- Transport protocol (TCP/UPD, TCP, or UDP).
- Authentication cache renewal.
- Filename translation for file characters supported by NFS. (NTFS does not support filename translation.)
- File locking wait period.
- Activity Logging.

- Netgroups for managing access to NFS shares.

Even though the Server for NFS Properties dialog box mentions NSF versions 2 and 3, Server for NFS in Windows Server 2016 and Windows Server 2019 also includes support for NFS version 4.1. There is also support in the NFS server implementation of Windows Server 2019 for the Pseudo file system, sessions and session trunking, and compound RPCs.

More Info

Server for NFS

To learn more about Server for NFS, consult the following link: https://docs.microsoft.com/en-us/windows-server/storage/nfs/nfs-overview.

DEDUPLICATION

Deduplication works by analyzing files, locating the unique chunks of data that make up those files, and storing only one copy of each unique data chunk on the volume. (A chunk is a collection of storage blocks.) Deduplication can reduce the amount of storage consumed on the volume because when analyzed, it turns out that a substantial number of data chunks stored on a volume are identical. Rather than store multiple copies of the same identical chunk, deduplication ensures that one copy of the chunk is stored with placeholders in other locations pointing at the single copy of the chunk, rather than storing the

chunk itself. Windows Server 2019 supports deduplication on both NTFS- and ReFS-formatted volumes.

Before you can enable deduplication, you need to install the Data Deduplication role service, as shown in Figure 7-24.

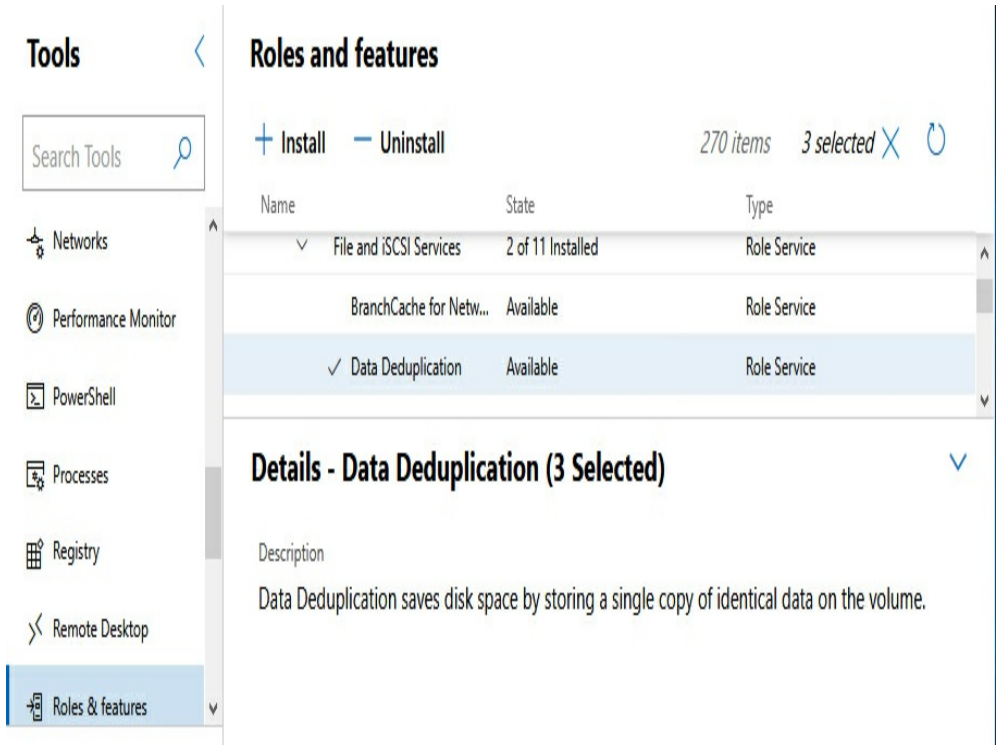


Figure 7-24 Enable deduplication

When you configure deduplication, you choose from one of the following usage types, as shown in Figure 7-25:

- **General Purpose File Server.** Appropriate for general-purpose file servers, optimization is performed on the background on any file that is older than three days. Files that are in use and partial

files are not optimized.

- **Virtual Desktop Infrastructure (VDI) server.** Appropriate for VDI servers, optimization is performed in the background on any file that is older than three days, but files that are in use and partial files will also be optimized.
- **Virtualized Backup Server.** This usage type is suitable for backup applications such as System Center Data Protection Manager or an Azure Backup Server. It performs priority optimization on files of any age. It will optimize in-use files, but it will not optimize partial files.

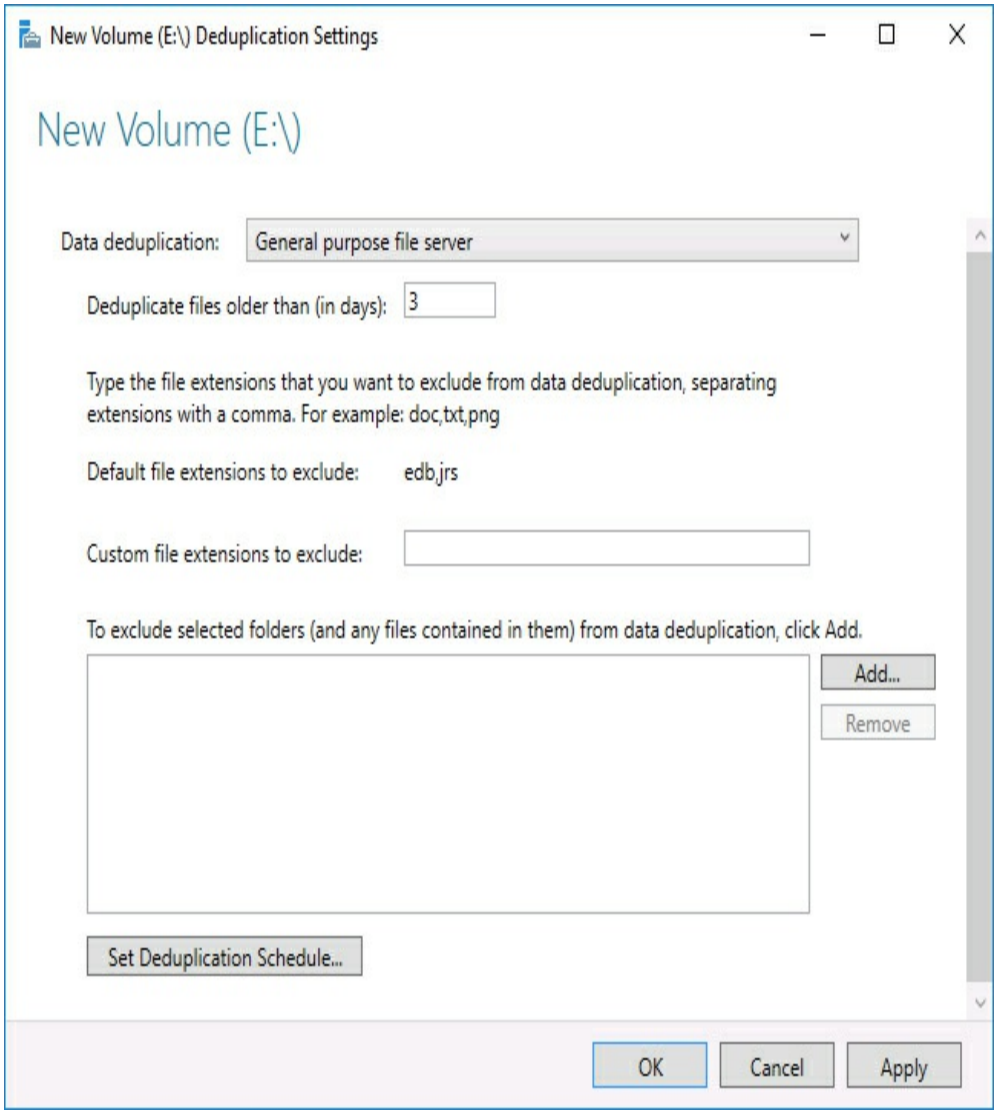


Figure 7-25 Deduplication settings

When configuring deduplication settings, you can configure files to be excluded based on file extension, or you can exclude entire folders from data deduplication. Deduplication involves running a series of jobs outlined in Table 7-3.

Table 7-3 Deduplication jobs

Name	Description	Schedule
Optimization	Deduplicates and optimizes the volume.	Once per hour
Garbage collection	Reclaims disk space by removing unnecessary chunks. You may want to run this job manually after deleting a substantial amount of data in an attempt to reclaim space by using the *Start-DedupeJob* cmdlet with the *Type* parameter set to *GarbageCollection*.	Every Saturday at 2:35 AM
Integrity scru	Identifies corruption in the chunk store and uses volume features where possible to repair and reconstruct corrupted data.	Every Saturday at 3:35

Technet24

binding		AM
Unoptimization	A special job that you run manually when you want to disable deduplication for a volume.	Run manually

More Info

Deduplication

To learn more about Deduplication, consult the following link: https://docs.microsoft.com/en-us/windows-server/storage/data-deduplication/overview.

STORAGE QUALITY OF SERVICE

Storage Quality of Service (QoS) allows you to centrally manage and monitor the performance of storage used for virtual machines that leverage the Scale-Out File Server and Hyper-V roles. The Storage QoS feature will automatically ensure that access to storage resources is distributed equitably between virtual machines that use the same file server cluster. It allows you to configure

minimum and maximum performance settings as policies in units of IOPS.

Storage QoS allows you to accomplish the following goals:

- **Reduce the impact of noisy neighbor VMs.** A noisy neighbor VM is a virtual machine that is consuming a disproportionate amount of storage resources. Storage QoS allows you to limit the extent to which such a VM can consume storage bandwidth.
- **Monitor storage performance.** When you deploy a virtual machine to a scale-out file server, you can review storage performance from a central location.
- **Allocate minimum and maximum available resources.** Through Storage QoS policies, you can specify the minimum and maximum resources available. This allows you to ensure that each VM has the minimum storage performance it requires to run reserved for it.

There are two types of Storage QoS policy:

- **Aggregated.** An aggregated policy applies maximum and minimum values for a combined set of virtual hard disk files and virtual machines. For example, by creating an aggregated policy with a minimum of 150 IOPS and a maximum of 250 IOPS and applying it to three virtual hard disk files, you can ensure that the three virtual hard disk files will have a minimum of 150 IOPS between them when the system is under load and will consume a maximum of 250 IOPs when the virtual machines associated with those hard disks are heavily using storage.
- **Dedicated.** A dedicated policy applies a minimum and maximum value to each individual virtual hard disk. For example, if you apply a dedicated policy to each of three virtual hard disks that specify a minimum of 150 IOPS and a maximum of 250 IOPS, each virtual hard disk will individually be able to use up to 250 IOPS, while

having a minimum of 150 IOPS reserved for use if the system is under pressure.

You create policies with the *New-StorageQosPolicy* cmdlet, which is specified by using the *PolicyType* parameter if the policy is Dedicated or Aggregated; also, you must specify the minimum and maximum IOPS. For example, run this command to create a new policy called *Alpha* of the Dedicated type that has a minimum of 150 IOPS and a maximum of 250 IOPS:

Click here to view code image

```
New-StorageQosPolicy -Name Alpha -PolicyType Dedicated -
MinimumIops 150 -MaximumIops 250
```

Once you've created a policy, you can apply it to a virtual hard disk using the *Set-VMHardDiskDrive* cmdlet.

More Info

Storage Quality of Service

To learn more about Storage QoS, consult the following link: https://docs.microsoft.com/en-us/windows-server/storage/storage-qos/storage-qos-overview.

REFS

The ReFS (Resilient File System) is appropriate for very large workloads where you need to maximize data availability and integrity and ensure that the file system is resilient to corruption. ReFS in Windows Server 2019 supports deduplication.

The ReFS file system is suitable for hosting specific types of workloads such as virtual machines and SQL Server data, because it includes the following features that improve upon NTFS:

- **Integrity.** ReFS uses checksums for both metadata and file data. This means that ReFS can detect data corruption.
- **Storage spaces integration.** When integrated with storage spaces that are configured with Mirror or Parity options, ReFS has the ability to automatically detect and repair corruption using a secondary or tertiary copy of data stored by Storage Spaces. The repair occurs without downtime.
- **Proactive error correction.** ReFS includes a data integrity scanner that scans the volume to identify latent corruption and proactively repair corrupt data.
- **Scalability.** ReFS is specifically designed to support data sets in the petabyte range.
- **Advanced VM operations.** ReFS includes functionality specifically to support virtual machine operations. Block cloning accelerates copy operations, which accelerate VM checkpoint merges. Sparse VDL allows ReFS to substantially reduce the amount of time required to create very large fixed-sized virtual hard disks.

It is important to note that ReFS is suitable only for hosting specific types of workloads. It isn't suitable for many workloads used in small and medium enterprises that aren't hosting large VMs or huge SQL Server databases. You won't choose to deploy ReFS for a file server because it doesn't support functionality such as File Server Resource Manager, file compression, file

encryption, extended attributes, and quotas.

You configure ReFS as an option when formatting or creating a volume, as shown in Figure 7-26.

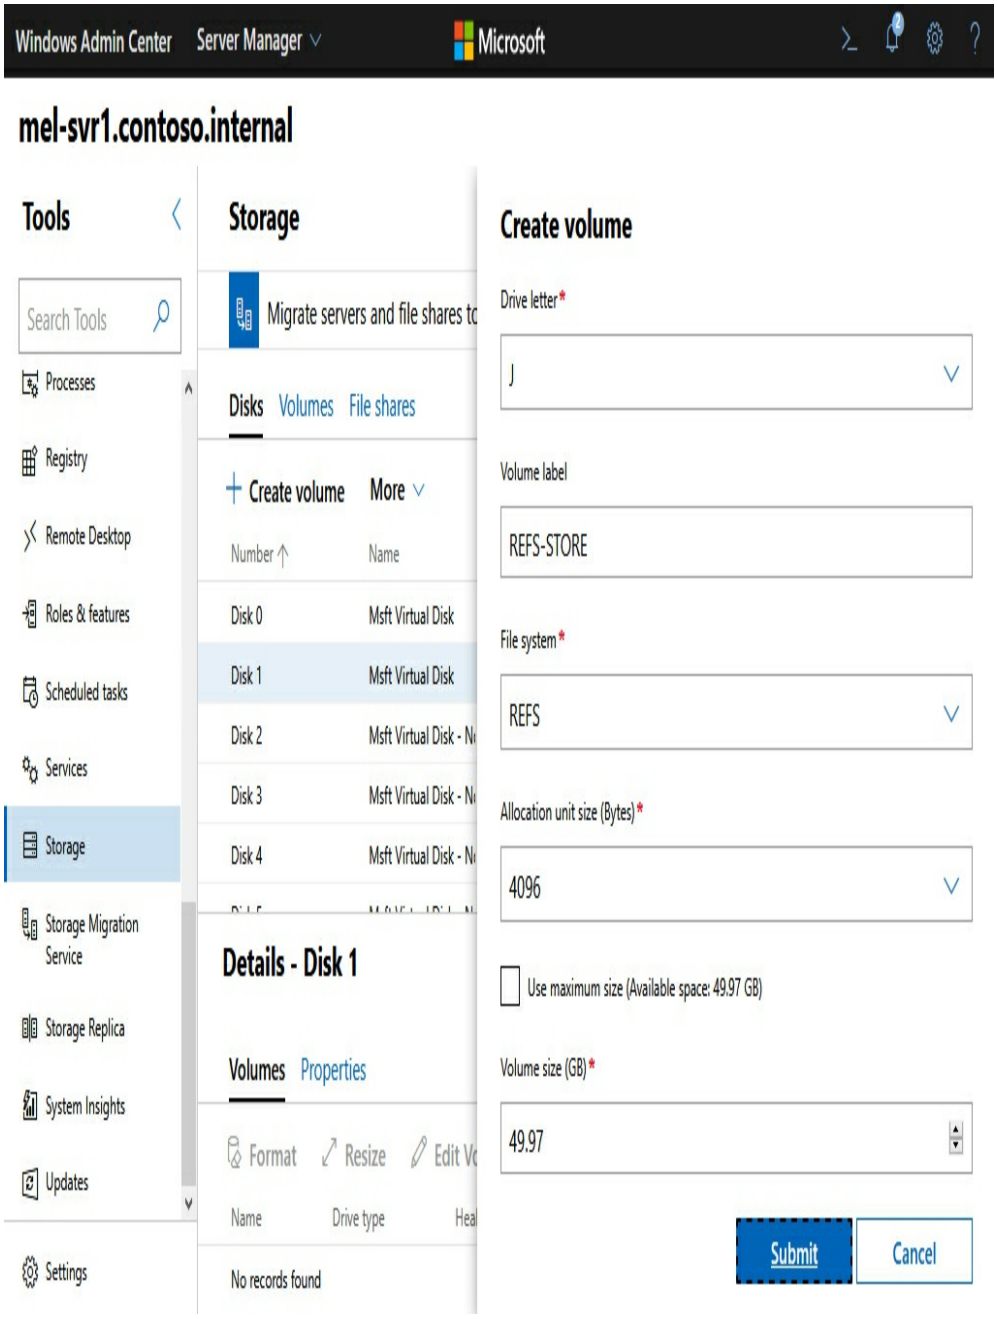


Figure 7-26 ReFS file system

When used with Storage Spaces Direct and Windows Server 2019, ReFS allows mirror-accelerated parity. Mirror-accelerated parity provides fault tolerance without impacting performance. To create a mirror-accelerated parity volume for use with Storage Spaces Direct, use the following PowerShell command:

Click here to view code image

```
New-Volume -FriendlyName “ExampleVolume” -filesystem CSVFS_ReFS -
StoragePoolFriendlyName
“ExamplePool” -StorageTierFriendlyNames Performance, Capacity -
StorateTierSizes 200GB,
800GB
```

More Info

Resilient File System

To learn more about ReFS, consult the following link: https://docs.microsoft.com/en-us/windows-server/storage/refs/refs-overview.

STORAGE-RELATED POWERSHELL CMDLETS

Several modules contain PowerShell cmdlets that are relevant to storage services. These are listed in the following tables.

Deduplication

Table 7-4 lists the PowerShell cmdlets in the

Technet24

deduplication module.

Table 7-4 Deduplication module cmdlets

Noun	Verbs	Function
Dedupfile	*Expand*	Expand a deduplicated file or directory
DedupFileMetadata	*Measure*	View file deduplication information
DedupJob	*Stop, Get, Start*	Manage deduplication jobs
DedupMetadata	*Get*	View deduplication metadata
DedupSchedule	*New, Remove, Get, Set*	Configure deduplication schedule
DedupStatus	*Update, Get*	View deduplication status
DedupVolume	*Disable, Enable, Get, Set*	Configure deduplication for a volume

iSCSI

Table 7-5 lists the PowerShell cmdlets in the iSCSI module.

Table 7-5 iSCSI module cmdlets

Noun	Verbs	Function
iSCSIChapSecret	*Set*	Configure iSCSI secret
iSCSIConnection	*Get*	View iSCSI connection
iSCSISession	*Unregister, Register, Get*	Configure iSCSI session
iSCSITarget	*Connect, Update, Disconnect, Get*	Manage iSCSI Target
iSCSITargetPortal	*Update, Get, New, Remove*	Configure iSCSI Target Portal

iSCSITarget

Table 7-6 lists the PowerShell cmdlets in the iSCSI Target module.

Table 7-6 iSCSI Target module cmdlets

Noun	Verbs	Function
iSCSIServerTarget	*Get, Remove, New, Set*	Manage iSCSI Server Target
iSCSITargetServerConfiguration	*Import, Export*	Backup and recover iSCSI target server configuration

Technet24

iSCSITarget ServerSettin g	*Set, Get*	View and manage iSCSI target server setting
iSCSIVirtual Disk	*Import, New, Remove, Restore, Convert, Checkpoint, Set, Get, Resize*	Manage iSCSI virtual disks
iSCSIVirtual DiskOperatio n	*Stop*	Stop iSCSI virtual disk operation
iSCSIVirtual DiskSnapsho t	*Set, Export, Dismount, Mount, Get, Remove*	Configure iSCSI virtual disk snapshots
iSCSIVirtual DiskTargetM apping	*Add, Remove, Expand*	Manage iSCSI virtual disk target mapping

NFS

Table 7-7 lists the PowerShell cmdlets in the NFS module.

Table 7-7 NFS module cmdlets

Noun	Verbs	Function
NFSClientConf iguration	*Get, Set*	View and configure NFS client settings

NFSClientGroup	*Remove, Rename, Set, Get, New*	Manage NFS client groups
NFSClientLock	*Get, Remove*	Manage NFS client locks
NFSMappedIdentity	*Set, Test, Remove, New, Get, Resolve*	Configure NFS mapped identities
NfsMappingStore	*Revoke, Get*	Configure NFS mapping store
NfsMountedClient	*Get, Revoke*	Manage clients connected to NFS server
NfsNetgroup	*Remove, Get, Set, New*	Configure NFS net groups
NfsNetgroupStore	*Set, Get*	Configure NFS net group store
NfsOpenFile	*Revoke, Get*	Manage NDF open files
NfsServerConfiguration	*Get, Set*	Manage NFS server configuration
NfsSession	*Get, Disconnect*	Manage NFS sessions
NfsShare	*Remove, Get, Set, New*	Configure NFS shares
NfsSharePermission	*Get, Revoke, Grant*	Manage NFS share permissions
NfsStatistics	*Reset, Get*	View and reset NFS

		statistics

Storage

Table 7-8 lists the PowerShell cmdlets in the X module.

Table 7-8 Storage module cmdlets

Noun	Verbs	Function
DedupProperties	*Get*	View deduplication properties
Disk	*Clear, Initialize, Get, Set, Update*	Configure disk
DiskImage	*Mount, Get, Dismount*	Mount, dismount, and view disk image
DiskStorageNodeView	*Get*	View disk storage node view
FileIntegrity	*Set, Repair, Get*	Manage file integrity
FileShare	*Get, Remove, Debug, Set, New*	Manage file shares
FileShareAccess	*Revoke, Grant, Block, Unblock*	Manage file share access
FileShareAccessControlEn	*Get*	View file share access control entries

try		
FileStorageTier	*Clear, Set, Get*	Configure and view file storage tiers
HostStorageCache	*Update*	Update host storage cache
InitiatorID	*Remove, Get*	View and remove initiator ID
InitiatorIdFromMaskingSet	*Remove*	Remove initiator ID from masking set
InitiatorIDToMaskingSet	*Add*	Add initiator ID to masking set
InitiatorPort	*Set, Get*	View and configure initiator port
MaskingSet	*Get, Remove, New, Rename*	Manage masking set
OffloadDataTransferSetting	*Get*	View offload data transfer setting
Partition	*Resize, New, Get, Remove, Set*	Manage partitions
PartitionAccessPath	*Remove, Add*	Manage partition access paths

Technet24

PartitionSupportSize	*Get*	View partition support size
PhysicalDisk	*Add, Set, Remove, Reset, Get*	Manage physical disks
PhysicalDiskIdentification	*Enable, Disable*	Enable or disable physical disk identification
PhysicalDiskStorageNodeView	*Get*	View physical disk storage node view data
PhysicalExtent	*Get*	View physical extent information
PhysicalExtentAssociation	*Get*	View physical extent association
ResiliencySetting	*Get, Set*	View and configure resiliency setting
StorageAdvancedProperty	*Get*	View storage advanced properties
StorageDiagnosticInfo	*Clear, Get*	View and reset storage diagnostic information
StorageDiagnosticLog	*Start, Stop*	Start and stop diagnostic logs

StorageEnclosure	*Get*	View storage enclosure information.
StorageEnclosureIdentification	*Disable, Enable*	Enable or disable storage enclosure identification
StorageEnclosureStorageNodeView	*Get*	View storage enclosure storage node view information
StorageEnclosureVendorData	*Get*	View storage enclosure vendor data
StorageFaultDomain	*Get*	View storage fault domain information
StorageFileServer	*Remove, Get, New, Set*	Configure storage file server
StorageFirmware	*Update*	Update storage firmware
StorageFirmwareInformation	*Get*	View storage firmware information
StorageHealthAction	*Get*	View storage health action

StorageHealthReport	*Get*	View storage health report
StorageHealthSetting	*Get, Set, Remove*	Configure and manage storage health setting
StorageHighAvailability	*Disable, Enable*	Enable or disable storage high availability
StorageJob	*Stop, Get*	Stop or view storage job
StorageMaintenanceMode	*Enable, Disable*	Enable or disable storage maintenance mode
StorageNode	*Get*	View storage node information
StoragePool	*Optimize, Set, Update, Remove, New, Get*	Manage storage pool
StorageProvider	*Set, Get*	View and configure storage provider
StorageProviderCache	*Update*	Update storage provider cache
StorageReliabilityCounter	*Get, Reset*	Manage storage reliability counter

StorageSetting	*Set, Get*	View and configure storage setting
StorageSubSystem	*Get, Unregister, Set, Register, Debug*	Manage storage subsystem
StorageSubsystemVirtualDisk	*New*	New storage subsystem virtual disk
StorageTier	*Get, Remove, New, Resize, Set*	Manage storage tiers
StorageTierSupportedSize	*Get*	View storage tier supported size
SupportedClusterSizes	*Get*	View supported cluster sizes
SupportedFileSystems	*Get*	View supported file system information
TargetPort	*Get*	View target port information
TargetPortal	*Get*	View target portal information
TargetPortFromMaskingSet	*Remove*	Remove target port from a masking set

Technet24

VirtualDisk	*Show, New, Get, Hide, Resize, Remove, Repair, Set, Disconnect, Connect*	Manage virtual disk
VirtualDiskClone	*New*	Create a new virtual disk clone
VirtualDiskFromMaskingSet	*Remove*	Remove a virtual disk from a masking set
VirtualDiskSnapshot	*New*	Create a virtual disk snapshot
VirtualDiskSupportedSize	*Get*	View virtual disk supported size
VirtualDiskToMaskingSet	*Add*	Add virtual disk to masking set
Volume	*Optimize, Set, Repair, Format, Get, Debug, New*	Manage volumes
VolumeCache	*Write*	Flushes the volume cache to storage
VolumeCorruptionCount	*Get*	View volume corruption data
VolumeScrubPolicy	*Get, Set*	View and configure volume scrub policy

Storage Replica

Table 7-9 lists the PowerShell cmdlets in the Storage Replica module.

Table 7-9 Storage Replica module cmdlets

Noun	Verbs	Function
SRAccess	*Get, Grant, Revoke*	Manage Storage Replica access
SRConfiguration	*Export*	Backup existing Storage Replica configuration
SRDelegation	*Grant, Revoke, Get*	Configure Storage Replica delegation
SRGroup	*Suspend, Sync, Remove, Set, Get, New*	Manage Storage Replica groups
SRMetadata	*Clear*	Clear Storage Replica metadata
SRNetworkConstraint	*Get, Set, Remove*	Manage Storage Replica network constraints
SRPartnership	*Set, Get, New, Remove*	Configure Storage Replica partnerships
SRTopology	*Test*	Verify Storage Replica topology

Technet24

Chapter 8

File servers

Shared folder permissions

File Server Resource Manager

Distributed File System

BranchCache

PowerShell commands

Traditionally, the most common use for Windows servers has been as file servers. No matter how advanced technology gets, people who work in an organization need a way to share files with each other that is less chaotic than emailing them or handing them over on USB drives. Even with cloud storage options such as SharePoint Online and OneDrive, many organizations still make use of the humble file server as a way of storing and sharing documents. Shared folders also provide a central location that can be backed up on a regular basis, something that gives them utility

beyond just being a central location to store documents.

In this chapter, you'll learn about how you can get more out of your Windows Server 2019 file server by taking advantage of the functionality of File Server Resource Manager, Distributed File System, and BranchCache. Azure File Sync, which allows you to take advantage of the cloud to assist with file tiering and replication, is covered in Chapter 17, "Azure IaaS and hybrid services."

Inside OUT

Enable deduplication

We talked about deduplication in detail in Chapter 7, "Storage." As mentioned there, deduplication is a really useful technology to use with file servers. This is because a surprising number of files stored on file servers are duplicates of one another. Some organizations that have implemented deduplication on file servers found that they reduced storage utilization by up to 80 percent. Remember, to enable deduplication, you need to have all the file shares on a volume that is separate from the OS volume. Both NTFS and ReFS support deduplication in Windows Server 2019.

SHARED FOLDER PERMISSIONS

The basic idea with shared folders is that you create a folder and assign permissions to a group, such as a department within your organization, and the people in that group use that space on the file server to store files to be shared with the group.

For example, you create a shared folder named *Managers* on a file server named *FS1*. Next, you set share permissions on the shared folder and file system

permissions on the files and folders within the shared folder. Permissions allow you to control who can access the shared folder and what users can do with that access. For example, permissions determine whether they are limited to just read-only access or whether they can create and edit new and existing files.

When you set share permissions and file system permissions for a shared folder, the most restrictive permissions apply. For example, if you configure the Share permission so that the Everyone group has Full Control and then configure the file permission so that the Domain Users group has Read Access, a user who is a member of the Domain Users group accessing the file over the network has Read Access.

Things get a little more complicated when a user is a member of multiple groups; in this case, the most cumulative permission applies. For example, in a file where the Domain Users group has Read Access, but the Managers group has Full Control, a user who is a member of both Domain Users and Managers who accesses the file over the network has the Full Control permission. This is great for a certification exam question, but it can be needlessly complex when you're trying to untangle permissions to resolve a service desk ticket.

Inside OUT

We need to talk about permissions

If you've ever studied for a Microsoft certification and come across file servers, you'll have learned about the difference between shared folder and file system permissions. The general problem with these permissions is that they tend to confuse people, which is why permissions are often incorrectly applied manually and why service desks are inundated with calls about the wrong people having access to a file while the right people can't open it. There are several solutions to the permissions problem. The first is to have permissions set automatically using Dynamic Access Control where access is determined by rules around the nature of the file and the user trying to access the file. You'll learn about Dynamic Access Control and how you can use it to solve the permissions problem in Chapter 14, "Dynamic Access Control and Active Directory Rights Management Services." The next permissions related technology is AD RMS, also covered in Chapter 14, which ensure that permissions stay with the file even when the file is moved away from the file share.

Using File Explorer

You can create shared folders in a variety of ways. The way that many administrators do it—often out of habit—is by using the built-in functionality of File Explorer. If you are using File Explorer to share folders, you have two general options when it comes to permissions:

- **Simple Share Permissions.** When you use the Simple Share Permissions option, you specify whether a user or group account has Read or Read/Write permissions to a shared folder. Figure 8-1 shows this for the Managers group. When you use Simple Share Permissions, both the share permissions and the file and folder-level permissions are set at the same time. It is important to note that any files and folders in the shared folder path have their permissions reset to match those configured through the File Sharing dialog box. This, however, doesn't happen with other forms of share permission configurations.

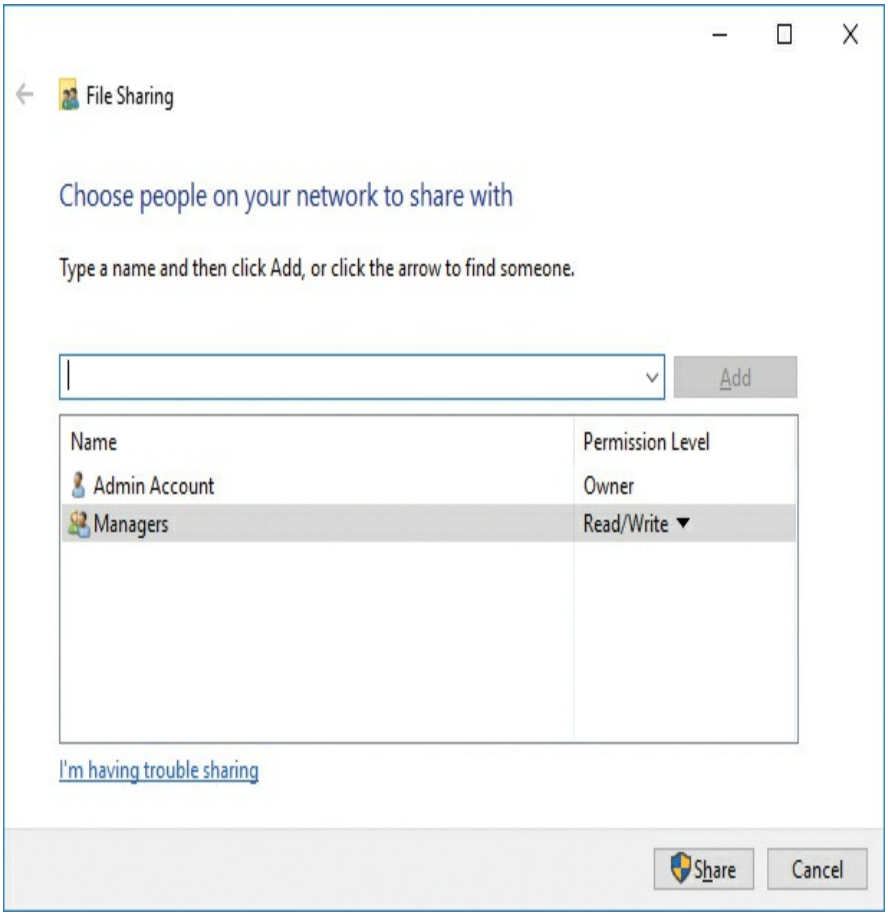


Figure 8-1 Simple share permissions

- **Advanced Share Permissions.** Advanced Share Permissions are what administrators who have been managing Windows file servers since the days of Windows NT 4 are likely to be more familiar with. With Advanced Permissions, you configure share permissions separately from file system permissions. You configure advanced share permissions through the Advanced Sharing button on the Sharing tab of a folder's Properties dialog box. Figure 8-2 shows how you configure Advanced Sharing Permissions. When you configure Advanced Sharing Permissions, permissions are only set on the share and are not reset on the files and folders within the share. If you are using Advanced Sharing Permissions, you set the file system permissions separately.

Technet24

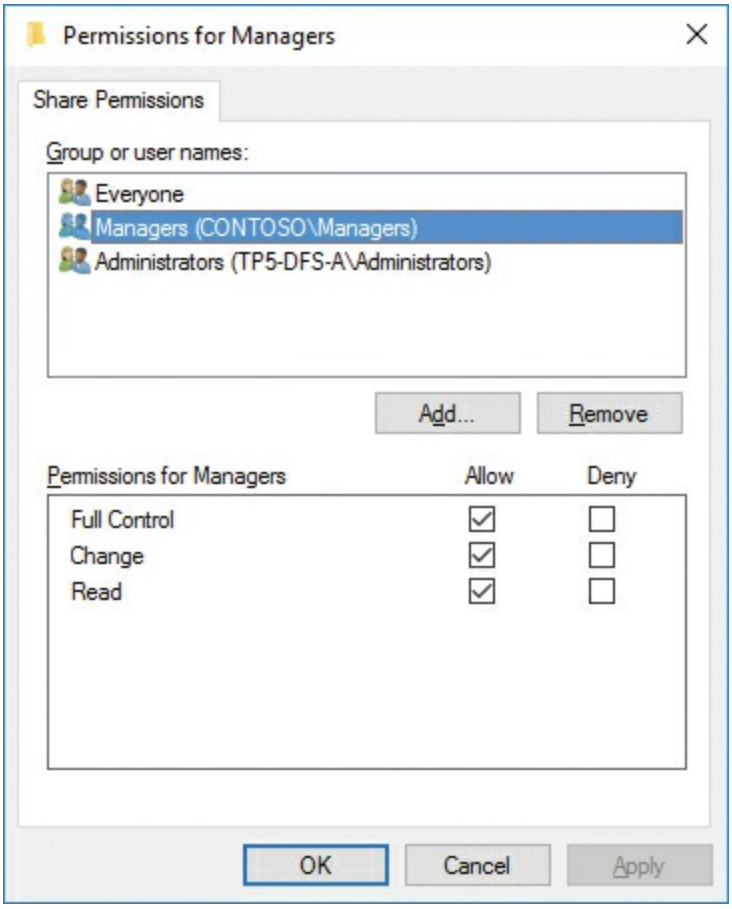


Figure 8-2 Share Permissions dialog box

Windows Admin Center

Windows Admin Center (WAC) also provides basic file share–configuration functionality, though this is not currently as sophisticated as what can be accomplished through Server Manager or File Explorer. As WAC evolves, it's likely that more file sharing functionality will be added to the tool. At the moment, file sharing is accessible through the Files node of WAC. To access the permissions of a share, you must select the host folder in the WAC interface and then use the More drop-down

menu to view share permissions. Share Permissions for a folder named *Shared-WAC* are displayed in Figure 8-3.

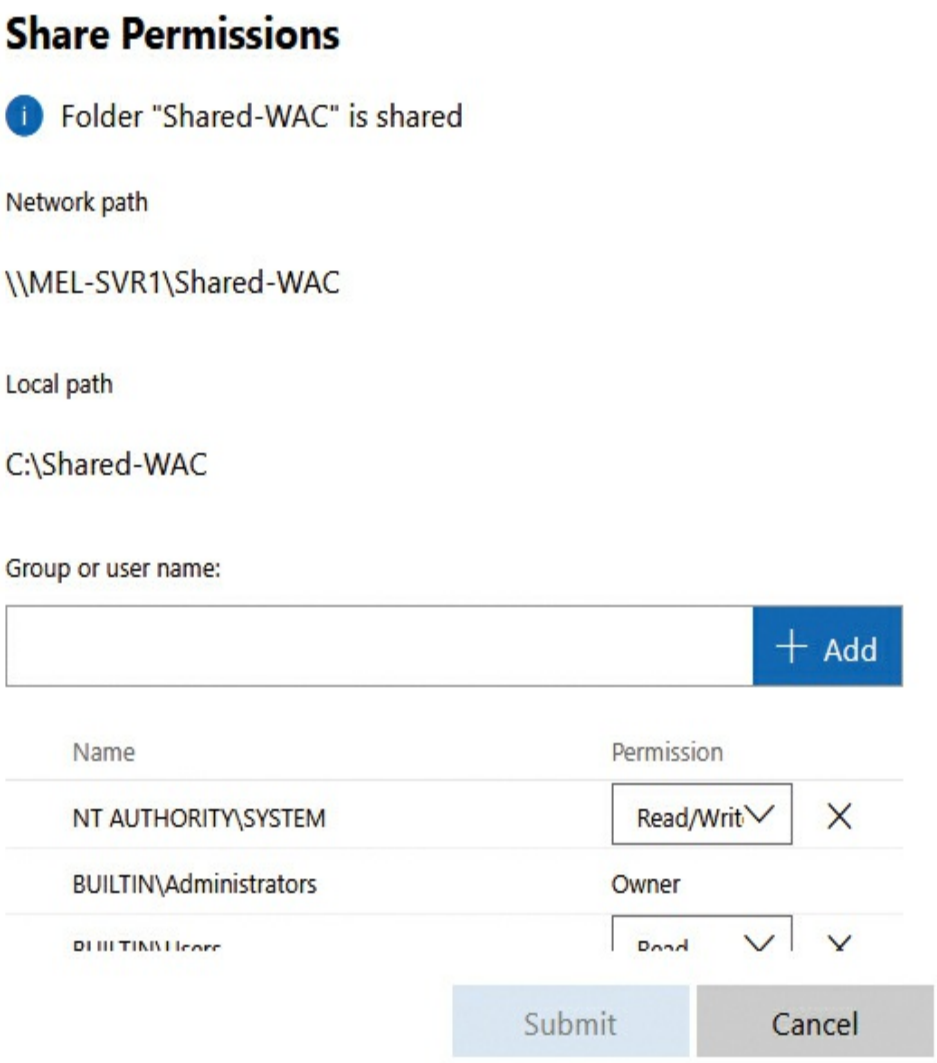


Figure 8-3 Windows Admin Center Share Permissions dialog box

Server Manager

You can manage shares centrally through the Shares

Technet24

area of the Server Manager console, as shown in Figure 8-4. The advantage of the Server Manager console is that you can use it to connect to manage shares on remote servers, including servers running the Server Core installation option.

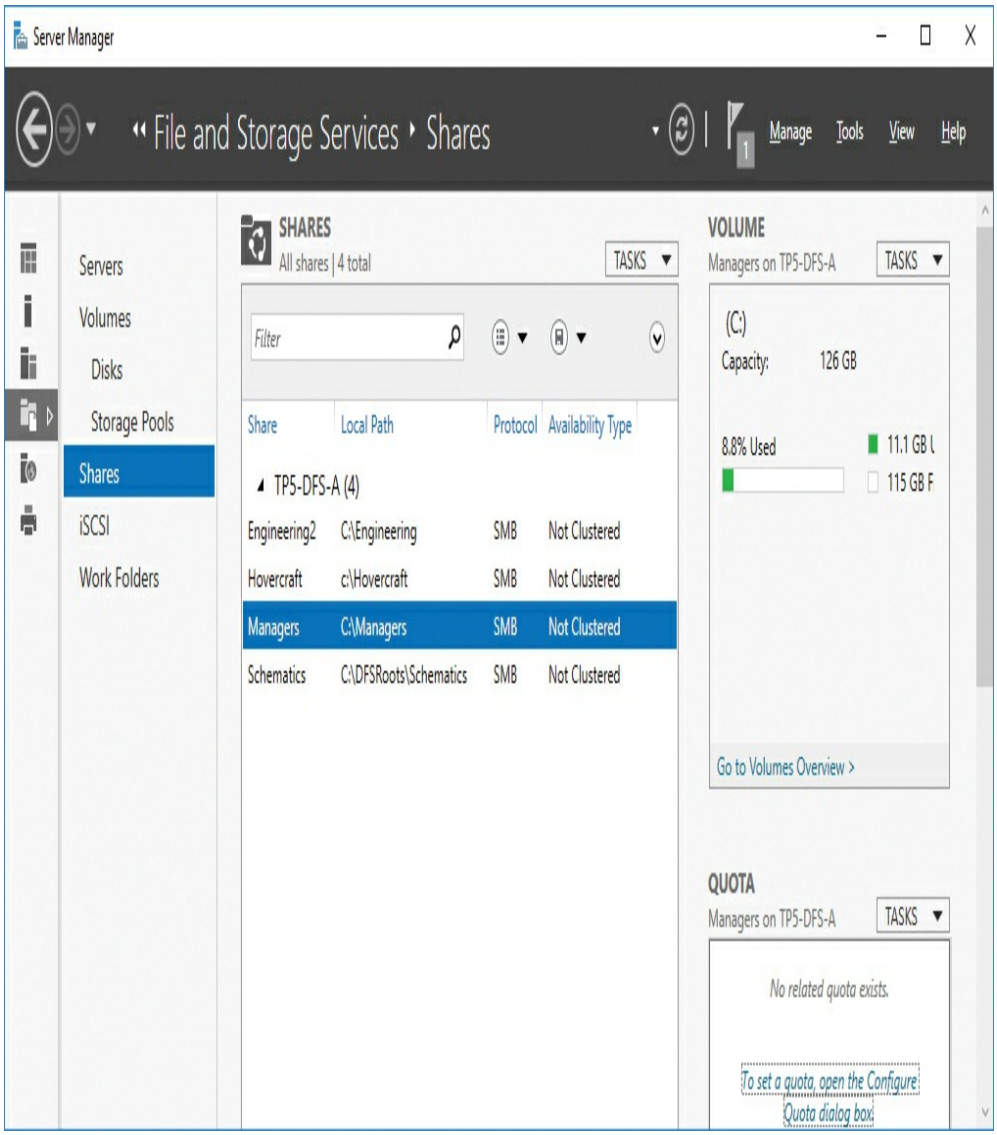


Figure 8-4 Server Manager for share management

When you edit the properties of a share through Server Manager, you can also edit the permissions, as shown in Figure 8-5. This functions in the same way as editing Advanced Share Permissions through File Explorer in that it won't reset the permissions on the file system itself; permissions are only reset on the share.

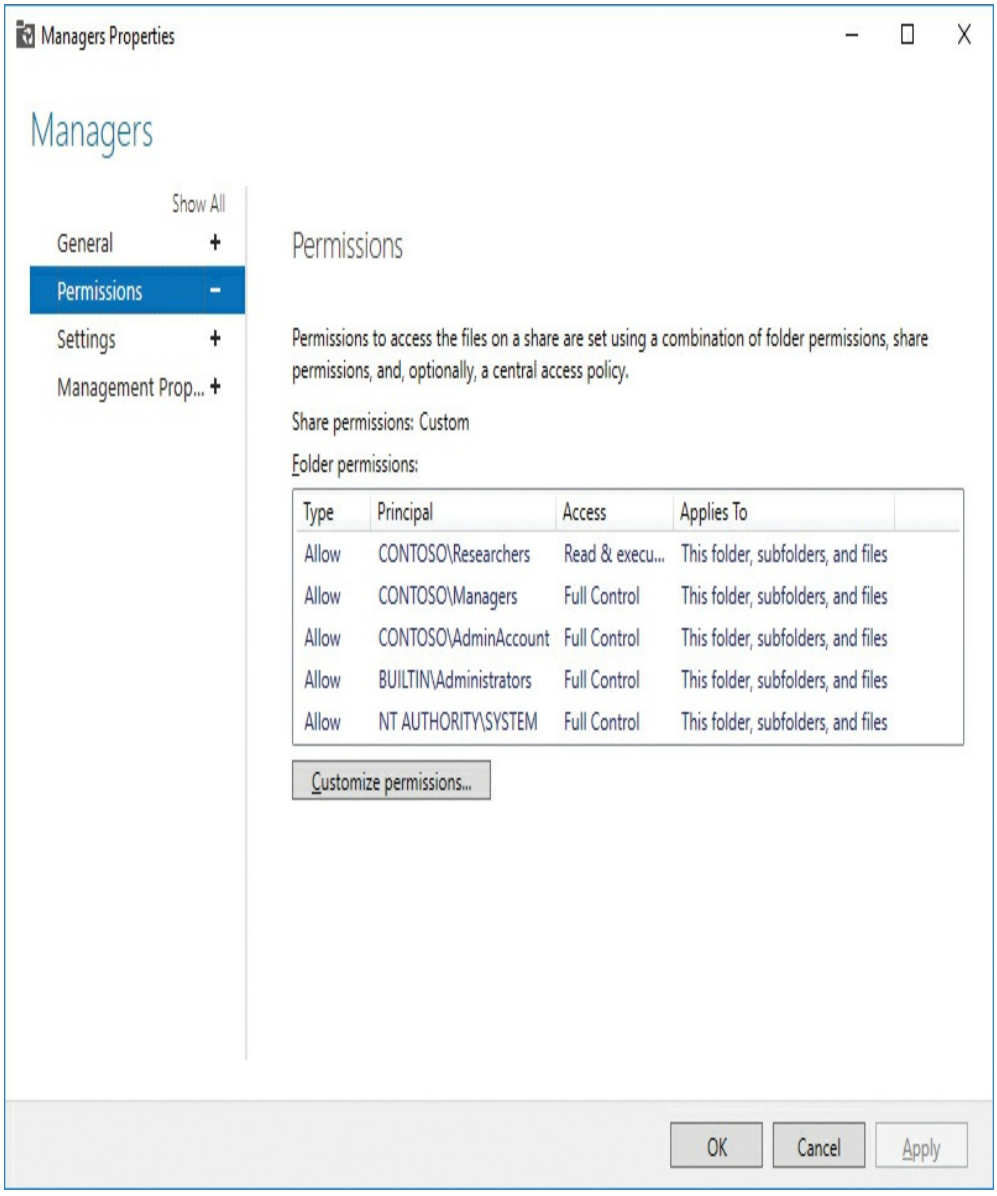


Figure 8-5 Server Manager Permissions management

You can also use the Server Manager's share properties interface when File Server Resource Manager is installed to edit the following settings:

- **Enable Access-Based Enumeration.** Enabled by default, this setting ensures that users can only see files and folders to which they have access. This is shown as enabled in Figure 8-6.

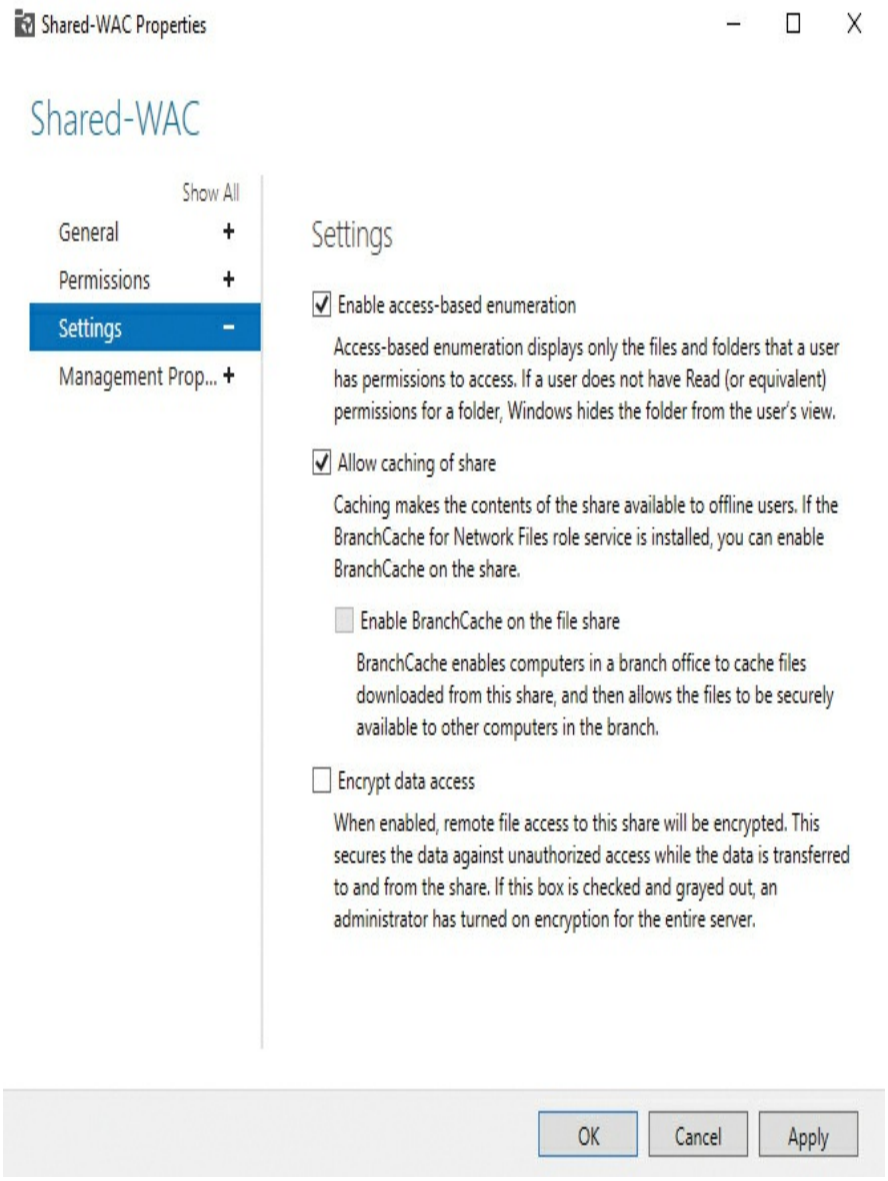


Figure 8-6 Server Manager share settings management

- **Allow Caching of Share.** Allows files to be used offline. An additional setting, Enable BranchCache on the file share allows use of BranchCache if the appropriate group policies are applied. You'll learn more about BranchCache later in this chapter. You can configure this option when the Settings tab is selected.
- **Encrypt Data Access.** When you enable this option, traffic to and from the shared folder is encrypted if the client supports SMB 3.0

Technet24

(Windows 8 and later). You can configure this option when the Settings tab is selected.

- **Specify Folder Owner Email Address.** Setting the folder owner's email address can be useful when resolving access-denied assistance requests. You can configure this when the Permission tab is selected.
- **Configure Folder Usage Properties.** Folder-usage properties allow you to apply metadata to the folder that specifies the nature of the files stored there. You can choose between User Files, Group Files, Application Files, and Backup and Archival files. You can use folder usage properties with data classification rules.

FILE SERVER RESOURCE MANAGER

File Server Resource Manager (FSRM) is a tool that allows you to perform advanced file server management. You can use FSRM to configure the following:

- Quotas at the folder level
- File screens
- Storage reports
- File classification
- File management tasks
- Access denied assistance

Folder level quotas

Quotas are important. If you don't use them, file shares tend to end up consuming all available storage. Some users dump as much as possible onto a file share unless

quotas are in place and unless you are monitoring storage; the first you'll hear about it is when the service desk gets calls about people being unable to add new files to the file share.

NTFS has had rudimentary quota functionality since the Windows NT days. The reason that most Windows server administrators don't bother with it is that it can't be applied only to individual user accounts and it applies at the volume level. Needless to say, if you have 500 users for whom you want to configure quotas, you don't want to have to individually configure a quota for each one. Even with command-line utilities, you still need to create an entry for each user.

Luckily, FSRM provides far more substantial quota functionality that makes quotas more practical to implement as a way of managing storage utilization on Windows file servers.

Quotas in FSRM can be applied on a per-folder basis and are not cumulative across a volume. You can also configure quotas in FSRM so that users are sent warning emails should they exceed a specific quota threshold, but before they are blocked from writing files to the file server. You manage quotas using FSRM by creating a quota template and then applying that quota template to a path.

Creating a quota template involves setting a limit,

specifying a quota type, and then configuring notification thresholds. You can choose between the following quota types:

- **Hard Quota. Do Not Allow Users To Exceed Limit.** A hard quota blocks users from writing data to the file share after the quota value is exceeded.
- **Soft Quota. Allow Users To Exceed Limit (Use For Monitoring).** A soft quota allows you to monitor when users exceed a specific storage utilization value, but it doesn't block users from writing data to the file share after the quota value is exceeded.

Notification thresholds allow you to configure actions to be taken after a certain percentage of the assigned quotas are reached. As shown in Figure 8-7, you can configure notifications via email, get an item written to an event log, run a command, or have a report generated.

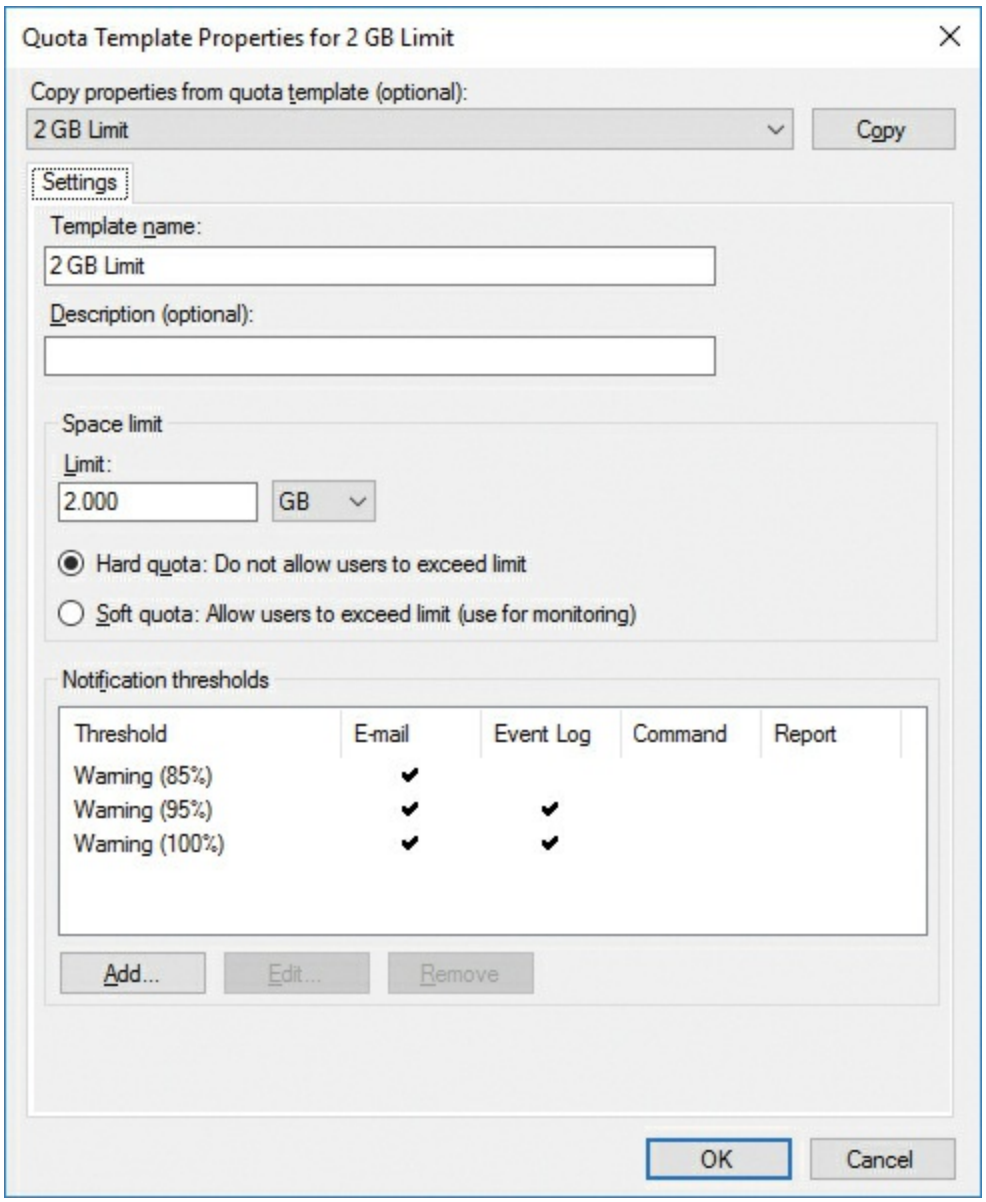


Figure 8-7 Quota template

After you create the quota template, you apply it to a folder. To do this, select the Quotas node under Quota Management, and from the Action menu, click Create Quota. In the Create Quota dialog box, shown in Figure

8-8, select the path to which the quota applies and the quota template to apply. You then choose between applying the quota to the whole path or setting up an auto-apply template. Auto-apply templates allow separate quotas to be applied to any new and existing quota path subfolders. For example, if you applied a quota to the *C:\Example* path using the 2 GB template, the quota would apply cumulatively for all folders in that path. If you chose an auto-apply template, a separate 2 GB quota would be configured for each new and existing folder under *C:\Example*.

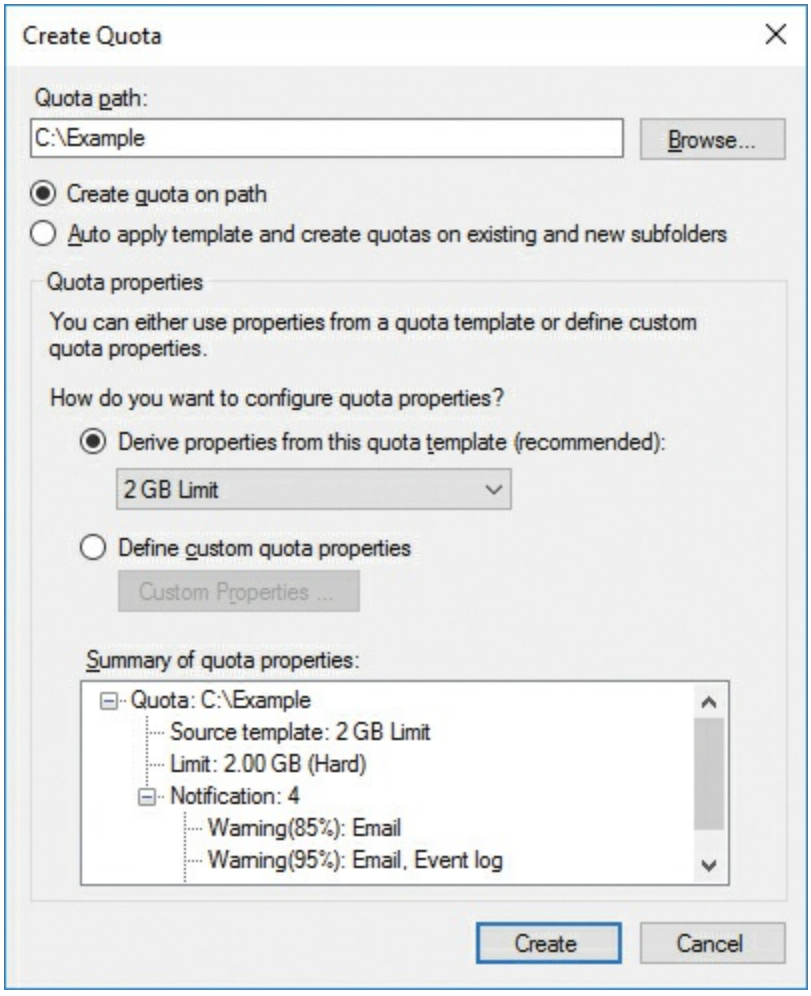


Figure 8-8 Create Quota dialog box

File screens

File screens allow you to block users from writing files to file shares based on file name extension. For example, you can use a file screen to stop people from storing video or audio files on file shares. File screens are implemented based on file name. Usually, this just means file screens are implemented by file extension, but you can configure file screens based on a pattern

match of any part of a file name. You implement file screens using file groups and file screen templates.

> **Inside OUT**
>
> **Doesn't block what is already there**
>
> A file screen doesn't block files that are already there; file screens just stop new files from being written to the share. File screens also only work based on file name. If you have users who are especially cunning, they might figure out that it's possible to get around a file screen by renaming files, so they don't get blocked by the screen.

File groups

A file group is a pre-configured collection of file extensions related to a specific type of file. For example, the Image Files file group includes file name extensions related to image files, such as *.jpg*, *.png*, and *.gif*. While file groups are usually fairly comprehensive in their coverage, they aren't always complete. Should you need to, you can modify the list to add new file extensions.

The file groups included with FSRM are shown in Figure 8-9 and include the following:

- **Audio And Video Files** Blocks file extensions related to audio and video files, such as *.avi* and *.mp3*
- **Backup Files** Blocks file extensions related to backups, including *.bak* and *.old* files
- **Compressed Files** Blocks file extensions related to compressed files, such as *.zip* and *.cab*
- **E-mail Files** Blocks file extensions related to email storage, including *.pst* and *.mbx* files

- **Executable Files** Blocks file extensions related to executable files and scripts, such as *.exe* or *.ps1* extensions
- **Image Files** Blocks file extensions related to images, such as *.jpg* or *.png* extensions
- **Office Files** Blocks file extensions related to Microsoft Office files, such as *.docx* and *.pptx* files
- **System Files** Blocks file extensions related to system files, including *.dll* and *.sys* files
- **Temporary Files** Blocks file extensions related to temporary files, such as *.tmp*. Also blocks files starting with the ~ character
- **Text Files** Blocks file extensions related to text files, including *.txt* and *.asc* files
- **Web Page Files** Blocks file extensions related to webpage files, including *.html* and *.htm* files

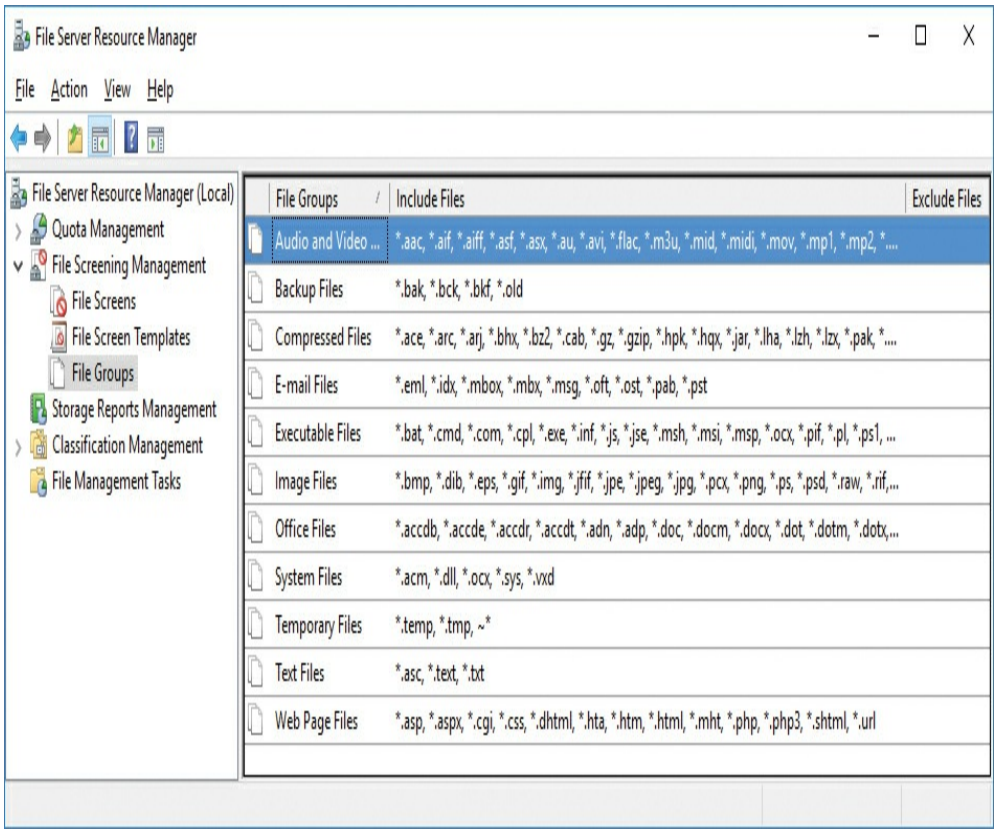


Figure 8-9 File Groups

To edit the list of files in a file group, right-click the file group and click Edit File Group Properties. Using the dialog box shown in Figure 8-10, you can modify the list of files to include and exclude files based on file name pattern. For example, you can do a simple exclusion or inclusion based on the file name suffix, such as *.*bak*. You also have the option of creating a more complex exclusion or inclusion based on the file name, such as backup *.*, which would exclude all files with the word backup at the start of any extension.

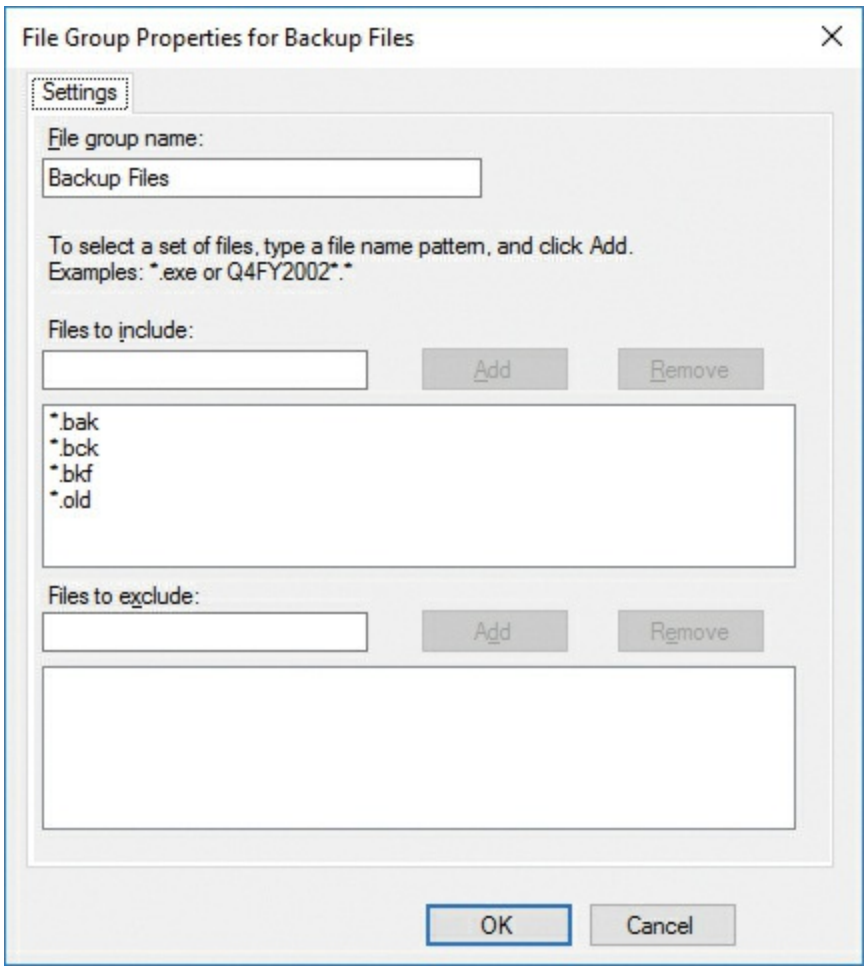


Figure 8-10 File Group Properties

Exclusions allow you to add exceptions to an existing block rule. For example, you could configure a file screen to block all files that have the extension *.vhdx*. You might then create an exception for the name *server2019.vhdx*. When implemented, all files with the *.vhdx* extension would be blocked from being written to the share, except for files with the name *server2019.vhdx*.

While the NTFS and ReFS file systems are case sensitive, file screens are not case sensitive.

To create a new file group, right-click the File Groups node in the File Server Resource Manager console and click Create File Group. Provide the following information:

- **File Group Name** The name for the file group.
- **Files To Include** Provide patterns that match the names of files you want to block from being written to the file server.
- **Files To Exclude** Provide patterns that match the names of files you want to exclude from the block.

File screen templates

File screen templates are made up of a screening type, a collection of file groups, and a set of actions to perform when a match is found. File screen templates support the following screening types:

- **Active Screening** An active screen blocks users from writing files to the file share that have names that match those patterns listed in the file group.
- **Passive Screening** A passive screen doesn't block users from writing files to the file share that have names that match patterns listed in the file group. Instead, you use a passive screen to monitor such activity.

The actions you can configure include sending an email, writing a message to the event log, running a command, or generating a report. Figure 8-11 shows the Create File Screen Template dialog box.

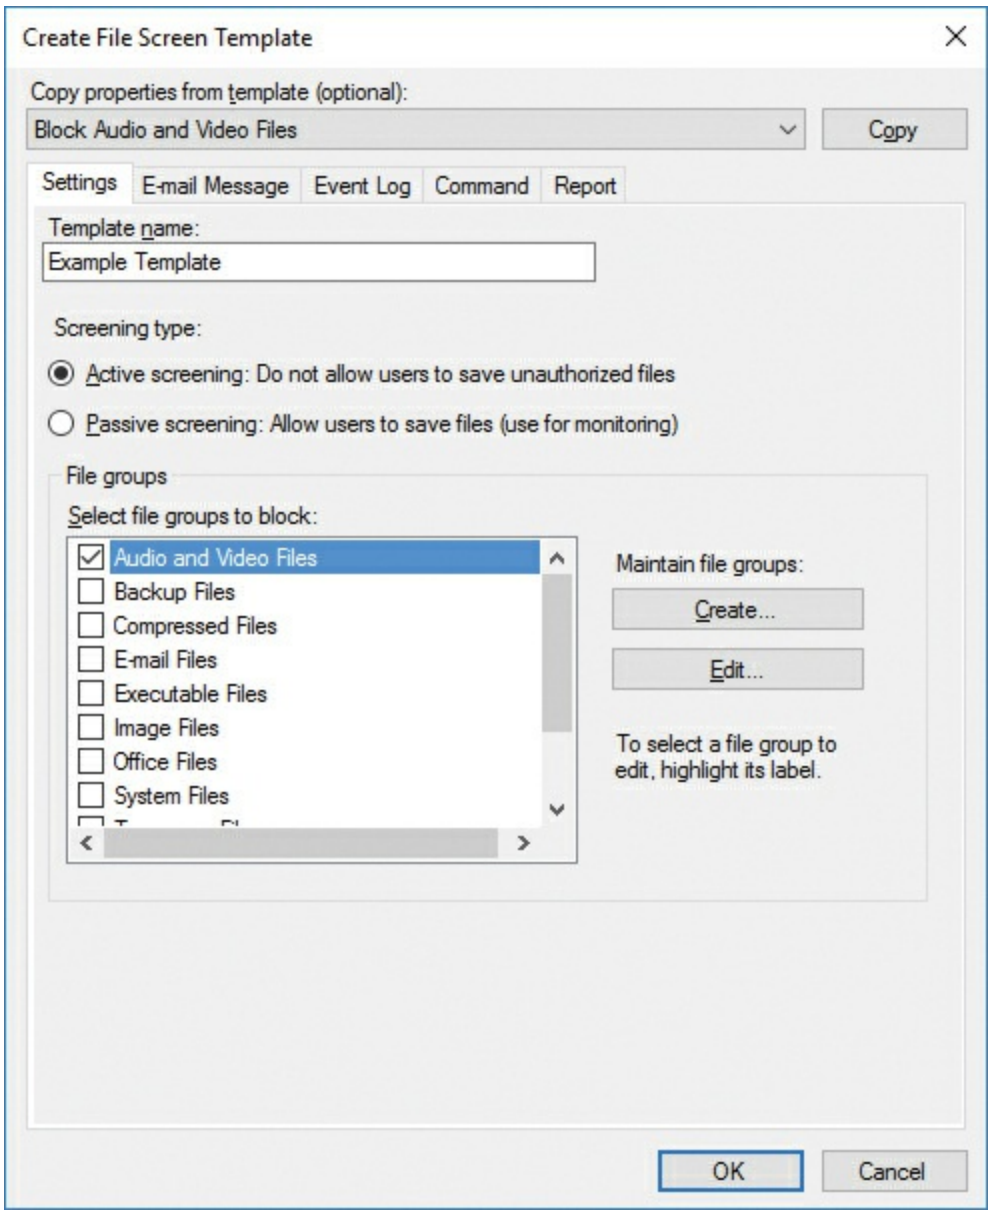


Figure 8-11 Creating a template

After you have configured the appropriate file screen template, create the file screen by applying the template to a specific path. You can also create file screen exceptions, which exempt specific folders from an

existing file screen. For example, you might apply a file screen at the root of a shared folder that blocks audio and video files from being written to the share. If you wanted to allow users to write audio and video files to one folder in the share, you could configure a file screen exception and apply it to that folder.

Storage reports

The Storage Reports functionality of FSRM allows you to generate information about the files stored on a particular file server. You can use FSRM to create the following storage reports:

- **Duplicate Files** This report locates multiple copies of the same file. If you've enabled deduplication on the volume hosting these files, these additional copies do not consume additional disk space as they are deduplicated.
- **File Screening Audit** This report allows you to view which users or applications are triggering file screens. For example, which users have tried to save music or video files to a shared folder.
- **Files By File Group** This report allows you to view files sorted by file group. You can view files by all file groups, or you can search for specific files. For example, a report on *.zip* files stored on a shared folder.
- **Files By Owner** This report allows you to view files by owner. You can search for files by all owners or run a report that provides information on files by one or more specific users.
- **Files By Property** Use this report to find out about files based on a classification. For example, if you have a classification named *Top_Secret*, you can generate a report about all files with that classification on the file server.

- **Large Files** This report allows you to find large files on the file server. By default, it finds files larger than 5 MB, but you can edit this to locate all files that are larger than a certain size.
- **Least Recently Accessed Files** This report allows you to identify files that have not been accessed for a certain number of days. By default, this report identifies files that have not been accessed in the last 90 days, but you can configure this to any number that is appropriate for your organization.
- **Most Recently Accessed Files** Use this report to determine which files have been accessed most recently. The default version of this report finds files that have been accessed in the last seven days.
- **Quota Usage** Use this report to view how a user's storage usage compares against the assigned quota. For example, you could run a report to determine which users have exceeded 90 percent of their quota.

You can configure storage reports to run and be stored locally on file servers. You also have the option of configuring storage reports to be emailed to one or more email addresses. You can generate storage reports in DHTML, HTML, XML, CSV, and text formats.

To run a storage report, perform the following steps:

1. In the File Server Resource Manager console, select the Storage Reports Management node, and then on the Action menu, click Generate Reports Now.
2. On the Settings page of the Storage Reports Task Properties dialog box shown in Figure 8-12, select the reports you want to run. You can click Edit Parameters to modify the properties of the report.

Technet24

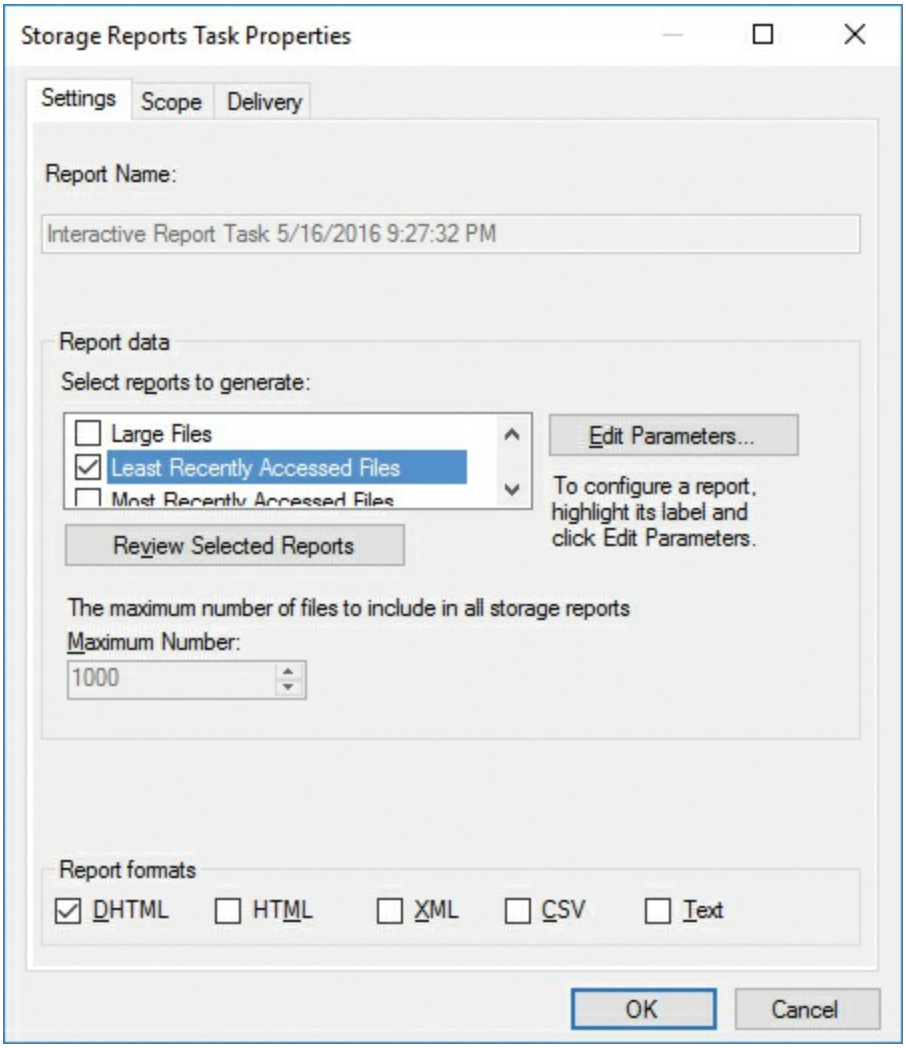


Figure 8-12 Storage reports

3. On the Scope tab of the Storage Reports Tasks Properties dialog box, you can configure which folders the report runs on.
4. On the Delivery tab of the Storage Reports Task Properties dialog box, you can configure the location where reports are saved and the email addresses to which the reports are sent.

File classification

File classification allows you to apply metadata to files based on file properties. For example, you can apply the

tag *Top_Secret* to a file that has specific properties, such as who authored it or whether a particular string of characters appeared in the file.

The first step to take when configuring file classification is to configure classification properties. After you've done this, you can create a classification rule to assign the classification property to a file. You can also allow users to manually assign classification properties to a file. By specifying the values allowed, you limit which classification properties the user can assign.

Classification properties

You can configure the following file classification properties:

- **Yes/No** Provide a Boolean value
- **Date-Time** Provide a date and time
- **Number** Provide an integer value
- **Multiple Choice List** Allow multiple values to be assigned from a list
- **Ordered List** Provide values in a specific order
- **Single Choice** Select one of a selection of options
- **String** Provide a text-based value
- **Multi-string** Assign multiple text-based values

To configure a classification property, perform the following steps:

1. In the File Server Resource Manager console, navigate to the

Classification Properties node under the Classification Management node.

2. On the Action menu, click Create Local Property.
3. In the Create Local Classification Property window, provide a Name, select a Property Type, and configure properties. Figure 8-13 shows a classification property named *Sensitivity* that gives a Single Choice property type and has the available values *Top_Secret* and *Not_Sensitive*. Click OK to save the new classification.

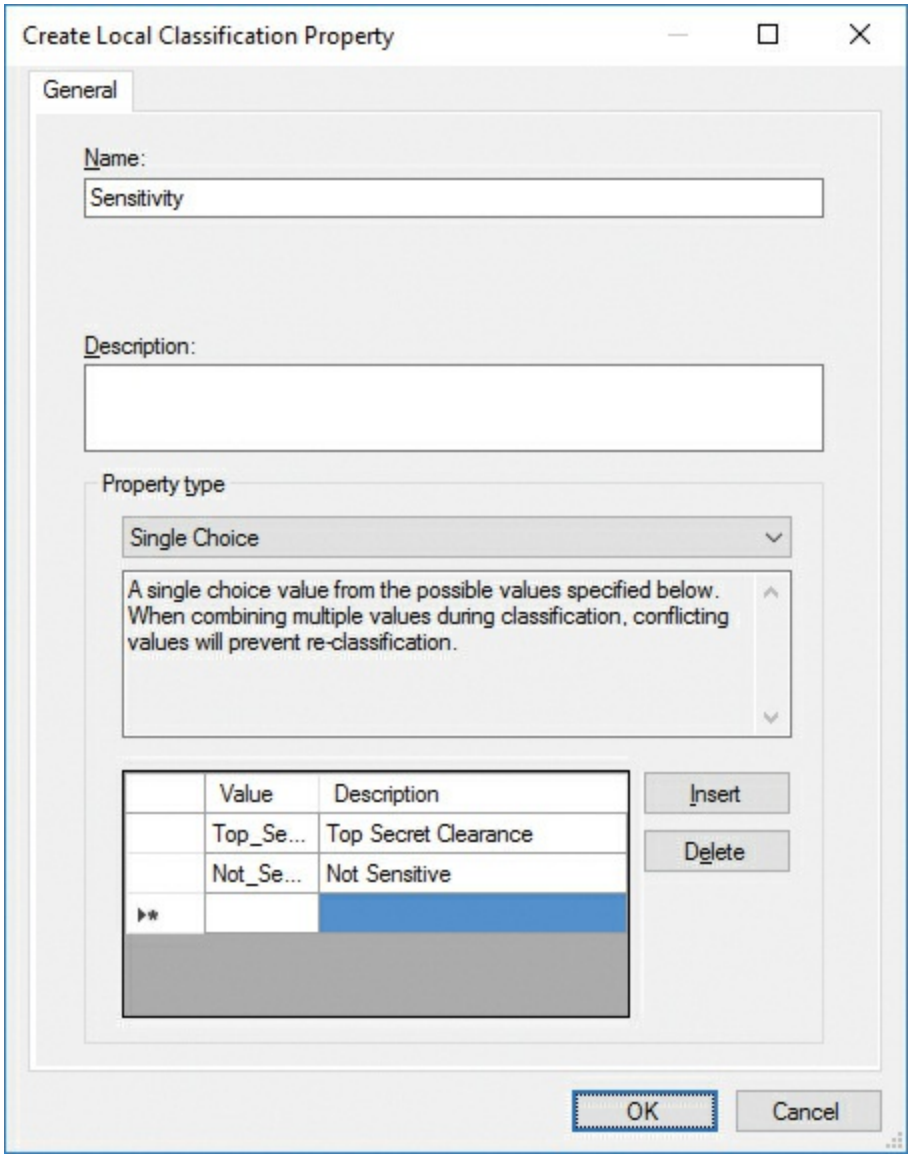


Figure 8-13 Create classification property

Classification rules

Classification rules allow you to assign classifications to files based on the properties of a file. You can configure three separate methods to classify a file:

- **Content Classifier** When you choose this method of classification, you configure a regular expression to scan the contents of a file for a specific string or text pattern. For example, you could use the content classifier to automatically assign the *Top_Secret* classification to any file that contained the text *Project_X*.
- **Folder Classifier** When you choose this method of classification, all files in a particular path are assigned the designated classification.
- **Windows PowerShell Classifier** When you choose this method of classification, a PowerShell script is run to determine whether a file is assigned a particular classification.

You can configure classification rules to run against specific folders. When configuring a classification rule, you can also choose to recheck files each time the rule is run. This allows you to change a file's classification in the event that the properties that triggered the initial classification change. When configuring reevaluation, you can also choose to remove user assigned classifications in case there is a conflict. Figure 8-14 shows the Evaluation Type tab of the Create Classification Rule dialog box, which allows you to configure these settings:

Technet24

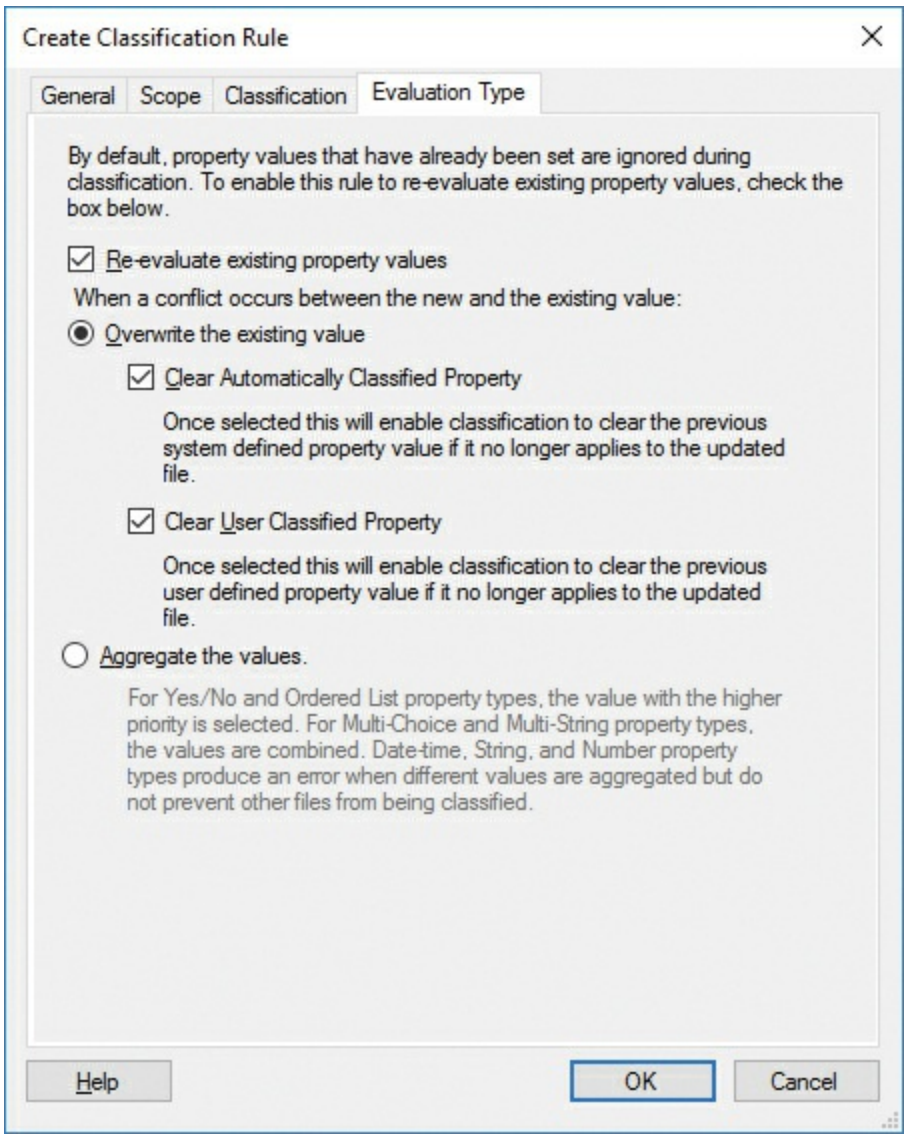


Figure 8-14 Evaluation Type

File management tasks

File management tasks are automated tasks that FSRM performs on files according to a schedule. FSRM supports three types of file management tasks. These are:

- **File Expiration** This moves all files that match the conditions to a specific directory. The most common usage of a file expiration task is to move files that haven't been accessed by anyone for a specific period, for example 365 days, to a specific directory.
- **Custom** Allows you to run a specific executable against a file. You can specify which executable is to be run, any special arguments to be used when running the executable, and the service account permissions, which can be Local Service, Network Service, or Local System.
- **RMS Encryption** Allows you to apply an RMS template or a set of file permissions to a file based on conditions. For example, you might want to automatically apply a specific set of file permissions to a file that has the *Top_Secret* classification or apply a specific RMS template to a file that has the *Ultra_Secret* classification.

When configuring a file management task, you also need to provide the following information:

- **Scope** The path where the task is run.
- **Notification** Any notification settings that you want to configure, such as sending an email, running a command, or writing an event to an event log. With file expiration, you can configure an email to be sent to each user who has files that are subject to the expiration task.
- **Report** Generating a report each time that the task is run. A notification is sent to the user who owns the file; reports are sent to administrators.
- **Schedule** When you want the file management task to be run.
- **Condition** The condition that triggers the management task. For example, Figure 8-15 shows a task that is triggered when the Sensitivity classification property is set to *Top_Secret*.

Technet24

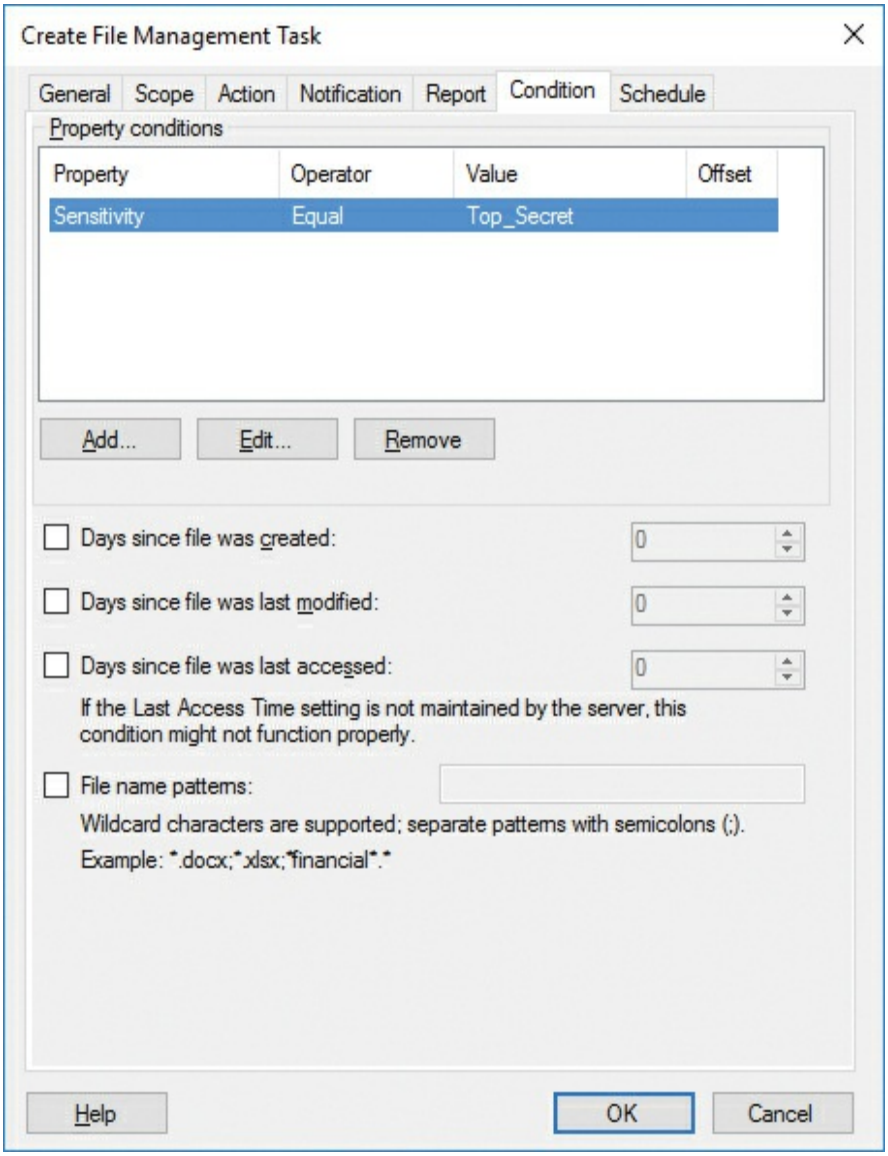


Figure 8-15 Create File Management Task

Inside OUT

Cleaning up file servers

Most files become stale after a few months. If a file hasn't been accessed for a year, there's a very good chance that it might never be accessed again. If you are responsible for managing file servers, you should come up with a way of identifying and removing these files on a periodic basis. Storage reports can help you identify them, and FSRM file management tasks allow you to move them off the file server in an automated way. Just remember that you might need to keep some files

around for a certain amount of time to meet your organization's compliance obligations. Another option is to adopt Azure File Sync, a solution that can automatically tier derelict files to the cloud where they can be synced back if required, but otherwise aren't actually present on the file server. You'll learn about Azure File Sync in Chapter 17, "Azure IaaS and hybrid services."

Access-Denied Assistance

Access-Denied Assistance allows users to be informed why they don't have access to a specific file. Access-Denied Assistance gives you the option of allowing the user to send an email message to the file owner so that they can, if appropriate, grant access to the file. You can configure Access-Denied Assistance using FSRM as shown in Figure 8-16, or by configuring group policy. You configure Access-Denied Assistance for a single server in FSRM by editing the FSRM options.

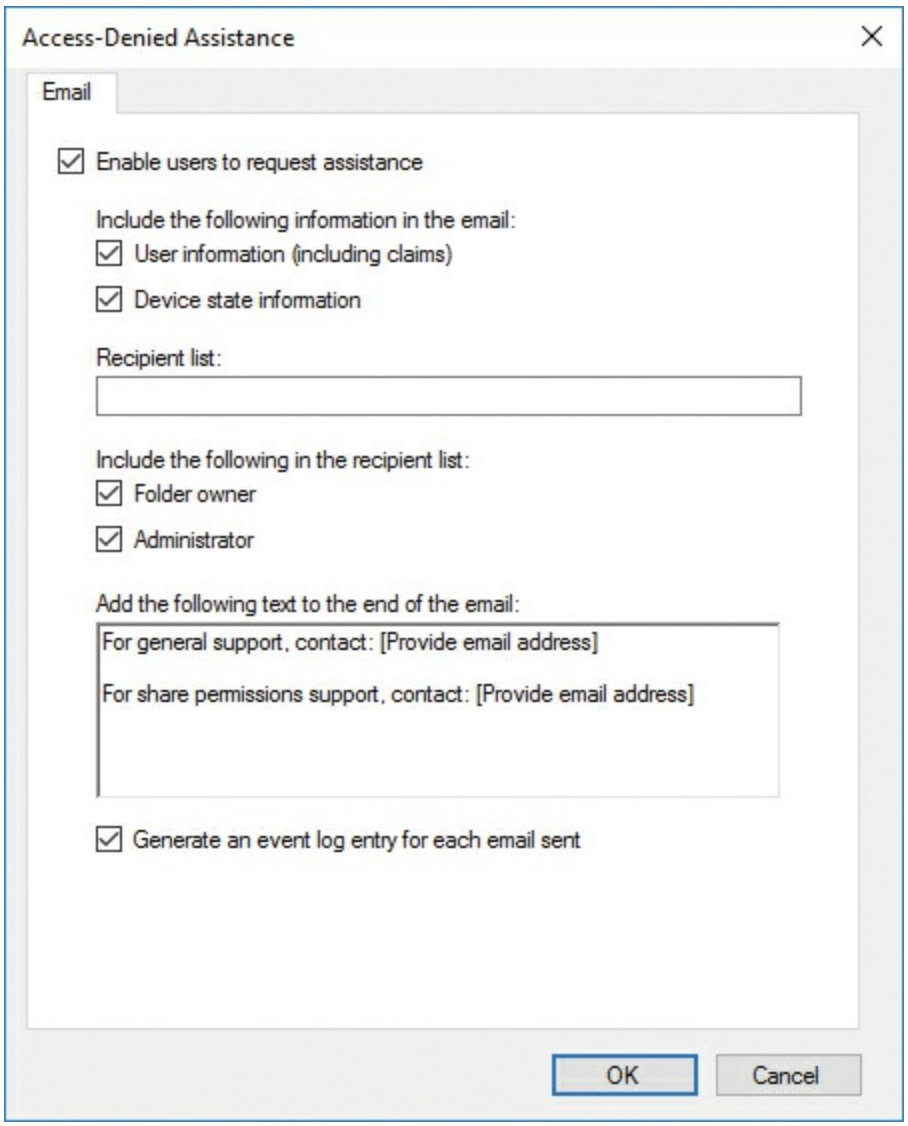


Figure 8-16 Access-Denied Assistance

If you want to use Access-Denied Assistance across all file servers in your organization, you can choose to use Group Policy. To configure Access-Denied Assistance using Group Policy, edit the policies located in the *Computer Configuration\Policies\Administrative Templates\System\Access-Denied Assistance* node. This

node contains the following policies:

- **Customize Message For Access Denied Errors** Use this policy to specify the message users see when they are blocked from accessing a file.
- **Enable Access-Denied Assistance On Client For All File Types** When enabled, Access-Denied Assistance functions for all file types where the user is blocked from accessing the file.

DISTRIBUTED FILE SYSTEM

Distributed File System (DFS) has two advantages over a traditional file share. The first is that DFS automatically replicates to create copies of the file share and its content on one or more other servers. The second is that clients connect to a single UNC address, with the client directed to the closest server and redirected to the next closest server in the event that a server hosting a DFS replica fails. Azure File Sync provides most of the first functionality but does not provide the second functionality. You'll learn more about Azure File Sync in Chapter 17, "Azure IaaS and hybrid services."

Using DFS, you can push a single shared folder structure out across an organization that has multiple branch offices. Changes made to files on one file share replica propagate across to the other file share replicas, with a robust and built-in conflict-management system present to ensure that problems do not occur when

Technet24

users are editing the same file at the same time.

DFS namespace

A DFS namespace is a collection of DFS shared folders. It uses the same UNC pathname structure, except instead of *\\ServerName\FileShareName with DFS*, it is *\\domainname* with all DFS shared folders located under this DFS root. For example, instead of

```
\\FS-1\Engineering
\\FS-2\Accounting
\\FS-3\Documents
```

You could have

```
\\Contoso.com\Engineering
\\Contoso.com\Accounting
\\Contoso.com\Documents
```

In this scenario, the *Engineering*, *Accounting*, and *Documents* folders could all be hosted on separate file servers and you could use a single namespace to locate those shared folders, instead of needing to know the identity of the file server that hosts them.

As shown in Figure 8-22, DFS supports the following types of namespaces:

- **Domain-Based Namespace** Domain-based namespaces store configuration data in Active Directory. You deploy a domain-based namespace when you want to ensure that the namespace remains available even if one or more of the servers hosting the namespace goes offline.

- **Stand-Alone Namespace** Stand-alone namespaces have namespace data stored in the registry of a single server and not in Active Directory as is the case with domain-based namespaces. You can only have a single namespace server with a standalone namespace. Should the server that hosts the namespace fail, the entire namespace is unavailable even if servers that host individual folder targets remain online.

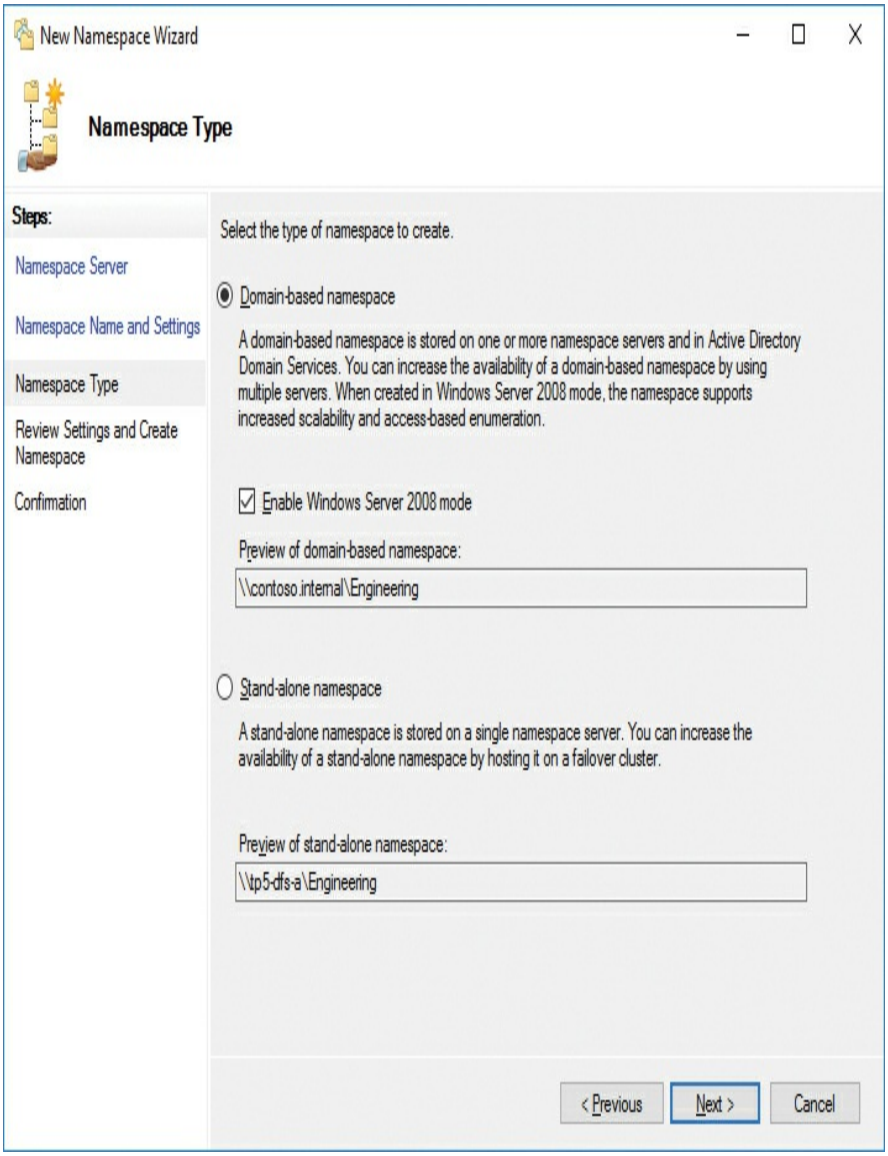


Figure 8-22 Namespace Type

Inside OUT

Windows Server 2008 mode

From the perspective of deploying Windows Server 2019, Windows Server 2008 mode might sound like an option that you'd use for backward compatibility. It's actually the opposite because if you don't select Windows Server 2008 mode, DFS uses Windows 2000 mode. Windows Server 2008 mode supports 50,000 folder targets and access-based enumeration, which means users can only see files and folders that they have access to. The requirements for Windows Server 2008 mode is that you have a domain functional level of Windows Server 2008 or higher.

To create a DFS namespace, perform the following steps:

1. In the DFS console, click the Namespaces node. In the Action menu, click New Namespace.
2. On the Namespace Server page, select a server that has the DFS Namespaces feature installed. You can install this feature with the following PowerShell cmdlet:

   Click here to view code image

   ```
   Install-WindowsFeature FS-DFS-Namespace
   ```

3. On the Namespace Name And Settings page, provide a meaningful name for the namespace. This is located under the Domain name. For example, if you added the name *Schematics* and you were installing DFS in the *contoso.internal* domain, the namespace would end up as *\\contoso.internal\Schematics*. By default, a shared folder is created on the namespace server, although you can edit settings on this page of the wizard and specify a separate location for the shared folder that hosts content you want to replicate.
4. On the Namespace Type page, you should generally select domain-based namespace as this gives you the greatest flexibility and provides you with the option of adding additional namespace servers later on for redundancy.

To add an additional namespace server to an existing namespace, ensure that the DFS Namespace role feature is installed on the server you want to add, and then perform the following steps:

1. In the DFS Console, select the namespace to which you want to add the additional namespace server, and on the Action menu, click Add Namespace Server.
2. On the Add Namespace Server page, specify the name of the namespace server, or browse and query Active Directory to verify that the name is correct, and then click OK. This creates a shared folder on the new namespace server with the name of the namespace.

DFS replication

A replica is a copy of a DFS folder. Replication is the process that ensures each replica is kept up to date. DFS uses block-level replication, which means that only blocks in a file that have changed are transmitted to other replicas during the replication process.

You install the DFS replication feature by running the following PowerShell command:

Click here to view code image

```
Install-WindowsFeature FS-DFS-Replication
```

In the event that the same file is being edited by different users on different replicas, DFS uses a “last writer wins” conflict-resolution model. In the unlikely event that two separate users create files with the same name in the same location on different replicas at

approximately the same time, conflict resolution uses "earliest creator wins." When conflicts occur, files and folders that "lose" the conflict are moved to the Conflict And Deleted folder, located under the local path of the replicated folder in the *DfsrPrivate\ConflictandDeleted* directory.

Replicated folders and targets

One of the big advantages of DFS is that you can create copies of folders across multiple servers that are automatically updated. Each copy of that replica is called a folder target. Only computers that have the DFS replication role feature installed can host folder targets. A replicated folder can have multiple folder targets. For example, you might have a replicated folder named *\\contoso.com\Engineering* that you have configured targets for in Sydney on *\\SYD-FS1\Engineering*, Melbourne on *\\MEL-FS1\Engineering*, and Auckland on *\\AKL-FS1\Engineering*.

To create a new folder to replicate, perform the following steps:

1. In the DFS console, select the namespace to which you want to add the folder.
2. On the Action menu, click New Folder. In the New Folder dialog box, provide a folder name and click Add.
3. If you've already created a shared folder to host the target, enter that address in the Add Folder Target dialog box. Otherwise, click Browse, select the server you want to host the folder, and click New Shared

Folder.

4. In the Create Share dialog box, shown in Figure 8-23, provide a share name, a local path, and the permissions that you want to apply. The advantage of using the DFS console to create the share is that it ensures all the appropriate folders are created and permissions are applied. You need an account that has local Administrator access on the server that hosts the share to perform these actions. Click OK until all dialog boxes are dismissed to create the new folder.

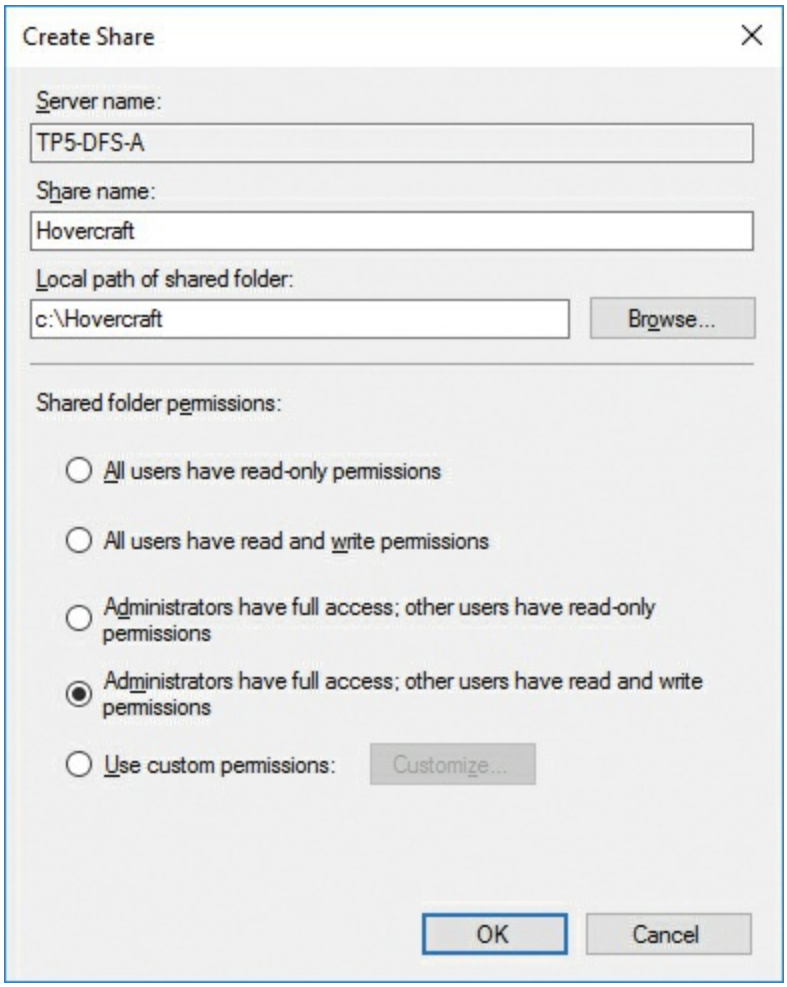


Figure 8-23 Create Share dialog box

To add a folder target to an existing folder, perform the

Technet24

following steps:

1. In the DFS console, select the folder under the namespace to which you want to add a folder target.
2. On the Action menu, click Add Folder Target.
3. If you've already created the share, enter the address of the shared folder. Otherwise, click Browse.
4. On the Browse For Shared Folders dialog box, click Browse again and use Active Directory to query for the address of the server that you want to have host the folder target. Remember that this server must have the DFS Replication role service installed on it.
5. Click New Shared Folder. In the Create Share dialog box, provide a share name, a local path, and the permissions that you want to apply. These should be the same as the permissions you've chosen for other replicas.
6. You need an account that has local Administrator access on the server that hosts the share to perform these actions. Click OK until all dialog boxes are dismissed to create the new folder target.

If this is the first additional target that you've created, you need to configure a replication topology and a replication group.

Replication topology

A replication group is a collection of servers that host copies of a replicated folder. When configuring replication for a replication group, you choose a topology and a primary member. The topology dictates how data replicates between the folders that each server hosts. The primary member is the seed from where file and folder data is replicated.

When creating a replication group, you can specify the following topologies as shown in Figure 8-24:

- **Hub And Spoke** This topology has hub members where data originates and spoke members to the location in which data is replicated. This topology also requires at least three members of the replication group. Choose this if you have a hub-and-spoke topology for your organizational WAN.
- **Full Mesh** In this topology, each member of the replication group can replicate with other members. This is the simplest form of replication group and is suitable when each member can directly communicate with the others.
- **No Topology** When you select this option, you can create a custom topology where you specify how each member replicates with others.

Technet24



Figure 8-24 Replication topology

Replication schedules

You use replication schedules to determine how replication partners communicate with each other. You use a replication schedule to specify when replication partners communicate and whether replication traffic is

throttled so that it doesn't flood the network.

You can configure replication to occur continuously and specify bandwidth utilization, with a minimum value of 16 Kbps and an upper value of 256 Mbps, with the option of setting it to Unlimited. If necessary, you can also set different bandwidth limitations for different periods of the day. Figure 8-25 shows that the cap set to 512 Kbps during work hours and to 64 Mbps after hours.

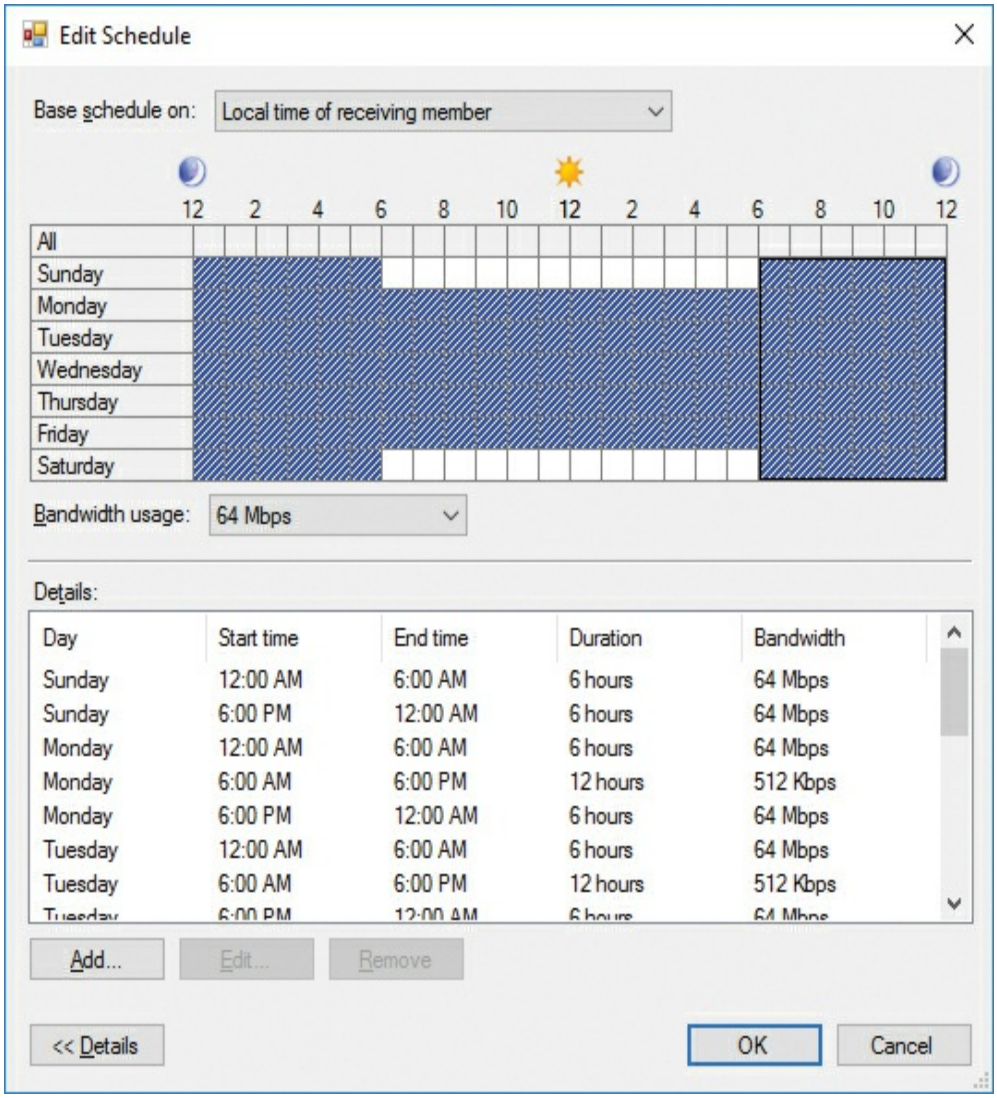


Figure 8-25 Edit Schedule dialog box

BRANCHCACHE

BranchCache speeds up access to files stored on shared folders that are accessed across medium to high latency WAN links. For example, several users in a company's Auckland, New Zealand branch office need to regularly access several files stored on a file server in the Sydney, Australia head office. The connection between the Auckland and Sydney offices is low bandwidth and high latency. The files are also fairly large and need to be stored on the Sydney file server. Additionally, the Auckland branch office is too small for a DFS replica to make sense. In a scenario such as this, you would implement BranchCache.

BranchCache creates a locally cached copy of files from remote file servers that can be accessed by other computers on the local network, assuming the file hasn't been updated at the source. In the example scenario, after one person in the Auckland office accesses the file, the next person to access the same file in the Auckland office accesses a copy that is cached locally, rather than retrieving it from the Sydney file server. The BranchCache process performs a check to verify that the cached version is up to date. If it isn't, the updated file is retrieved and stored in the Auckland network's BranchCache.

You add BranchCache to a file server by using the following PowerShell command:

Click here to view code image

```
Install-WindowsFeature FS-BranchCache
```

After installing BranchCache, you need to configure group policies that apply to file servers in your organization that allow them to support BranchCache. To do this, you need to configure the Hash Publication for BranchCache policy, located in the *Computer Configuration\Policies\Administrative Templates\Network\Lanman Server* node, as shown in Figure 8-26.

Technet24

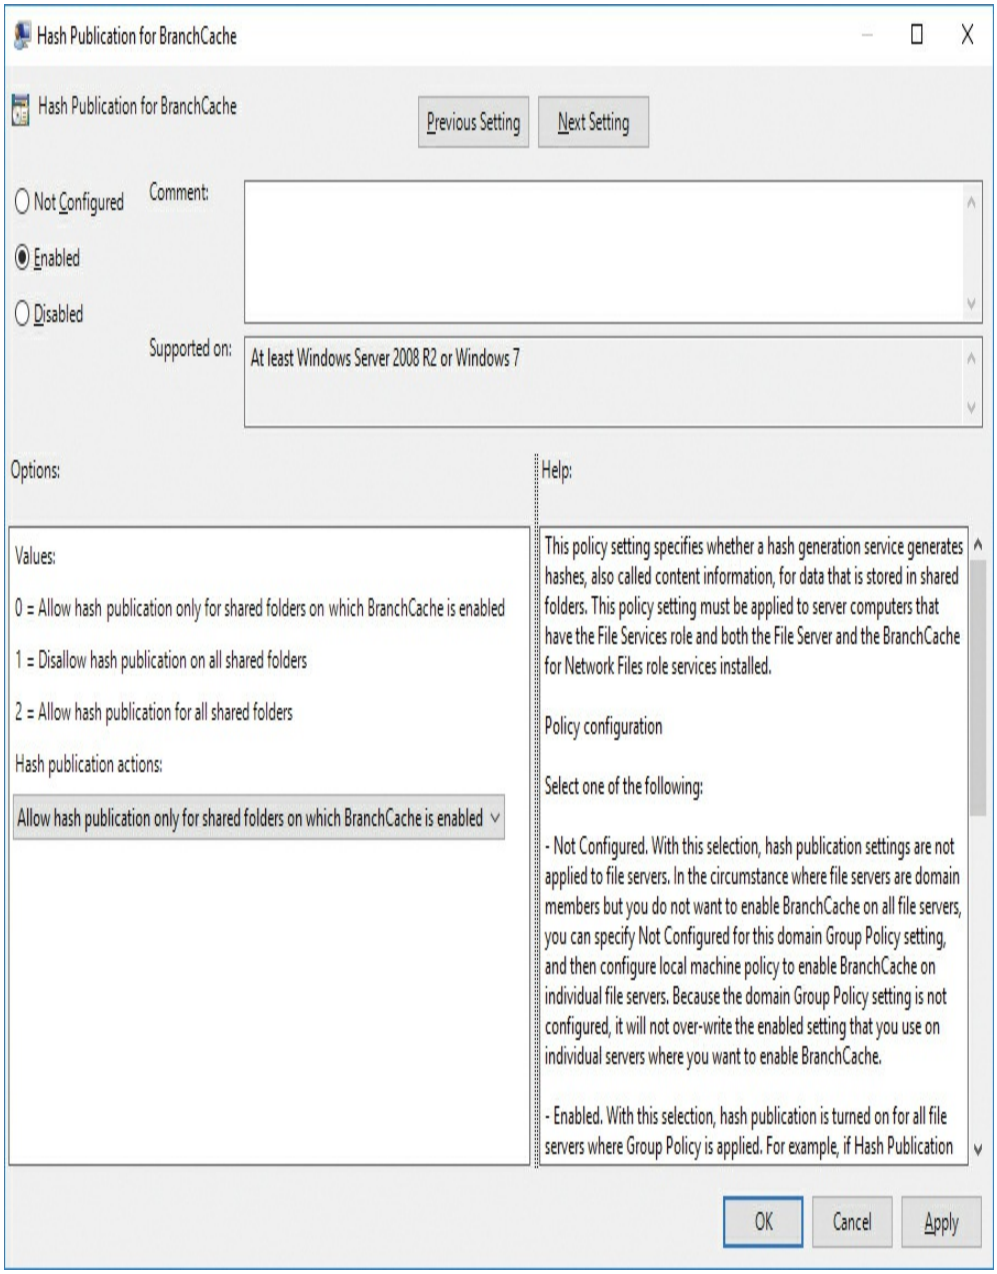


Figure 8-26 Hash Publication for BranchCache

You have three options when configuring this policy:

- **Allow Hash Publication Only For Shared Folders On Which BranchCache is enabled** This option allows you to selectively

enable BranchCache.

- **Disallow Hash Publication On All Shared Folders** Use this option when you want to disable BranchCache.
- **Allow Hash Publication For All Shared Folders** Use this option if you want to enable BranchCache on all shared folders.

Generally, there's rarely a great reason not to enable BranchCache on all shared folders, but should you want to be selective, you do have that option. Should you choose to be selective and only enable BranchCache on some shares, you need to edit the properties of the share and enable BranchCache. You do this by clicking Caching in the Advanced Sharing page and then using the Offline Settings dialog box as shown in Figure 8-27.

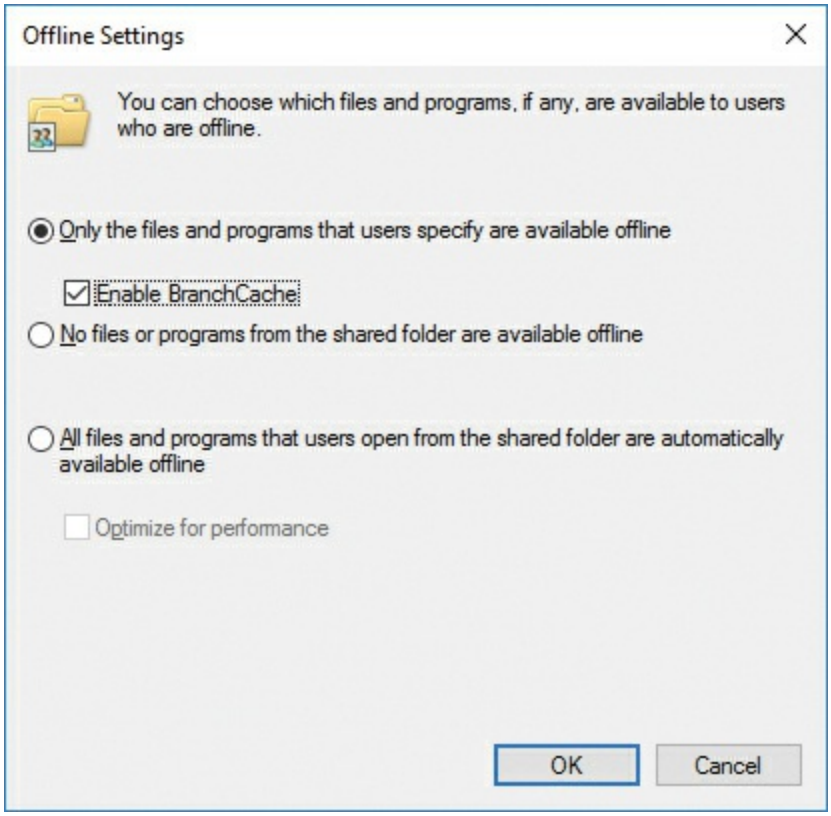


Figure 8-27 Offline Settings dialog box

After you've configured your file server to support BranchCache, you need to configure client computers to support branch cache. You do this through configuring group policy in the *Computer Configuration\Policies\Administrative Templates\Network\BranchCache* node of a GPO. Which policies you configure depends on how you want BranchCache to work at each branch office. You can choose between the following options:

- **Distributed Cache Mode** When client computers are configured for Distributed Cache mode, each Windows 7 or later computer hosts part of the cache.
- **Hosted Cache Mode** When you configure Hosted Cache mode, a server at the branch office hosts the cache in its entirety. Any server running Windows Server 2008 R2 or later can function as a hosted cache mode server.

To configure a branch office server to function as a hosted cache mode server, run the following PowerShell commands:

Click here to view code image

```
Install-WindowsFeature BranchCache
Start-Service BranchCache
Enable-BCHostedServer
```

POWERSHELL COMMANDS

While you can remotely manage file servers using the

Remote Server Administration Tools, if you need to manage a large number of servers it is more efficient to use PowerShell. In the next few pages, you'll learn which PowerShell cmdlets are available for managing Windows Server 2019 file servers.

Shared Folder cmdlets

Shared folder cmdlets are in the *SmbShare* PowerShell module. This is because traditional file shares use the SMB protocol. Sometimes, the documentation refers to SMB servers, which is another way of saying "file servers." Table 8-1 lists the shared folder related cmdlets available in Windows Server 2019.

Table 8-1 Shared Folder cmdlets

Noun	Verbs	Functionality
SMBBandwidthLimit	Get, Remove, Set	Manage shared folder bandwidth caps
SMBClientConfiguration	Get, Set	View client SMB configuration
SMBClientNetworkInterface	Get	View network interfaces used by client
SMBConnection	Get	View connections from client to shared folders

Technet24

SMBDelegation	Disable, Enable, Get	Manage constrained delegation authorization for SMB clients
SMBMapping	Get, New, Remove	Manage SMB mappings
SMBMultiChannelConnection	Get, Update	Manage multichannel connections to SMB shares
SMBMultichannelConstraint	Get, New, Remove	Manages SMB multichannel constraints on shared folders
SMBOpenFile	Get, Close	Manage files that are open from shared folders
SMBPathAcl	Set	Configure the ACL for the file system folder to match the ACL of the shared folder
SMBServerConfiguration	Get, Set	Manage SMB server configuration
SMBServerNetworkInterface	Get	View network interfaces used by SMB server
SMBSession	Get, Close	View and terminate SMB sessions
SMBShare	Get, New, Remove, Set	Manage shared folders

SMBShareAccess	Block, Get, Grant, Revoke, Unblock	Manage permissions on shared folders

File Server Resource Manager cmdlets

One of the challenges with FSRM is that you need to configure it on a per-server basis. Having complex FSRM configurations can make it challenging to deploy the same configuration across multiple servers. One way of simplifying the process of deploying FSRM across multiple file servers is to use PowerShell. By creating and running a PowerShell script, or by using DSC, you can apply the same FSRM settings across multiple file servers.

Table 8-2 lists the FSRM related cmdlets on the basis of PowerShell Noun, available verbs, and functionality.

Table 8-2 FSRM cmdlets

Noun	Verbs	Functionality
FSRMAction	New	Create a new action, such as sending an email, creating an event log entry, running a command, or generating a report
FsrmAdrSetting	Get, Set	Manage access-denied remediation settings, which display information for users when they are denied access to a file

Technet24

FSRMAutoQuota	Get, New, Remove, Set, Update	Manage auto apply quotas. Auto apply quotas are quota templates that you can apply at the volume or folder level
FSRMClassification	Get, Set, Start, Stop, Wait	Manage classification processes, such as starting and stopping them
FSRMClassificationPropertyDefinition	Get, New, Remove, Set, Update	Manage classification property definitions
FSRMClassificationPropertyValue	New	Manage classification property values.
FSRMClassificationRule	Get, New, Remove, Set	Manage classification property rules.
FSRMEffectiveNamespace	Get	Provides a list of paths that match FSRM namespaces
FSRMFileGroup	Get, New, Remove, Set	Manage FSRM file groups for file screens.
FSRMFileManagementJob	Get, New, Remove, Set, Start, Stop, Wait	Manage FSRM file management jobs

FSRMFileScreen	Get, New, Remove, Reset, Set	Manage FSRM file screens
FSRMFileScreenException	Get, New, Remove, Set	Manage FSRM file screen exceptions
FSRMFileScreenTemplate	Get, New, Remove, Set	Manage FSRM file screen templates
FSRMFmjAction	New	Create file management job actions
FSRMFmjCondition	New	Create file management job conditions
FSRMFmjNotification	New	Create file management job notifications
FSRMFmjNotificationAction	New	Create file management notification actions
FSRMMacro	Get	Lists FSRM macros
FSRMMgmtProperty	Get, Remove, Set	Manage FSRM management properties

FSRMQuota	Get, New, Remove, Reset, Set, Update	Manage FSRM quotas
FSRMQuotaTemplate	Get, New, Remove, Set	Manage FSRM quota templates
FSRMQuotaThreshold	New	Manage FSRM quota thresholds
FSRMRmsTemplate	Get	View available RMS templates that you can apply using FSRM
FSRMScheduledTask	New	Create an FSRM Scheduled Task
FSRMSetting	Get, Set	Manage FSRM settings
FSRMStorageReport	Get, Remove, New, Set, Start, Stop, Wait	Manage FSRM storage reports
FSRMTestEmail	Send	Send a test alert email using FSRM settings

BranchCache Cmdlets

You can use the BranchCache cmdlets, available when you install the BranchCache feature, to configure and

manage BranchCache on both file servers and clients. These cmdlets are listed in Table 8-3.

Table 8-3 BranchCache cmdlets

Noun	Verbs	Function
BC	Disable, Reset	Disable or reset BranchCache
BCAuthentication	Set	Configure BranchCache Computer Authentication mode
BCCache	Clear, Set	Manage cache files
BCCachePackage	Export, Import	Export and import cache packages
BCClientConfiguration	Get	View client computer settings
BCContentServerConfiguration	Get	View content server configuration
BCDataCache	Get	Get information about the data cache
BCDataCach	Set	Manage the amount of time data can be stored

Technet24

eEntryMaxAge		in the cache
BCDataCacheExtension	Add, Get	Manage cache storage space
BCDistributed	Enable	Configure a computer to operate in distributed cache mode
BCDowngrading	Enable, Disable	Manage whether computers can use Windows 7 BranchCache mode
BCFileContent	Publish	Creates hashes for files on a file server that has BranchCache enabled
BCHashCache	Get	View the BranchCache hash cache
BCHostedCacheServerConfiguration	Get	View hosted cache server settings
BCHostedClient	Enable	Configure BranchCache to function in hosted cache client mode
BCHostedServer	Enable	Configure BranchCache to function in hosted cache server mode
BCLocal	Enable	Enable BranchCache service to function in local caching mode

BCMinSMBLatency	Set	Configure minimum latency requirement for activation of BranchCache functionality
BCNetworkConfiguration	Get	View BranchCache network settings
BCSecretKey	Export, Import, Set	Manage cryptographic key used for generation of BranchCache segment secrets
BCServeOnBattery	Disable, Enable	Configure whether content discovery requests are responded to depending on whether a computer is running on battery
BCStatus	Get	View BranchCache status and configuration information
BCWebContent	Publish	Creates hashes for web content

DFS Cmdlets

DFS cmdlets are available across two separate PowerShell modules. You use the first set of cmdlets, listed in Table 8-4, to manage DFS namespaces and namespace servers. You use the second set of cmdlets, listed in Table 8-5 to manage DFS replication and folder targets.

Technet24

Table 8-4 DFS namespace cmdlets

Noun	Verbs	Function
DfsnAccess	Get, Grant, Remove, Revoke	Manage permissions for DFS namespace folders
DfsnFolder	Get, Move, New, Remove, Set	Manage DFS namespace folder settings
DfsnFolderTarget	Get, New, Remove, Set	Manage DFS namespace folder target settings
DfsnRoot	Get, New, Remove, Set	Manage settings for DFS roots
DfsnRootTarget	Get, New, Remove, Set	Manage settings for DFS root targets
DfsnServerConfiguration	Get, Set	Manage DFS namespace settings for DFSN root server

Table 8-5 DFS Replication cmdlets

Noun	Verbs	Function
DfsrBacklog	Get	Provides list of pending file updates between replication partners
DfsrClone	Export, Get, Import, Reset	Manage a cloned DFS replication database

DfsrConfigurationFromAD	Update	Updates the DFS Replication service by checking Active Directory database.
DfsrConnection	Add, Get, Remove, Set	Manage connections between members of a replication group
DfsrConnectionSchedule	Get, Set	Manage connection schedule between members of a replication group
DfsrDelegation	Get, Grant, Revoke	Manage replication group permissions
DfsReplicatedFolder	Get, New, Remove, Set	Manage replicated folders in a replication group
DfsReplicationGroup	Get, New, Remove, Set, Suspend, Sync	Manage replication groups
DfsrFileHash	Get	View a file hash
DfsrGroupSchedule	Get, Set	Manage replication group schedules

Technet24

DfsrGuid	ConvertFrom	Translates GUIDs to friendly names for a specific replication group
DfsrHealthReport	Write	Generates a replication health report that you can use for diagnostic purposes
DfsrIdRecord	Get	View ID records for replicated files or folders from the DFS replication database
DfsrMember	Add, Get, Set	Manage computer members of DFS replication group
DfsrMembership	Get, Set	View and configure replication group membership settings
DfsrPreservedFiles	Get, Restore	Manage preserved files and folders. Preserved files are ones where there has been a conflict or where files have been deleted
DfsrPropagationReport	Write	Create reports based on the propagation of test files
DfsrPropagationTest	Remove, Start	Manage propagation test files
DfsrServiceConfiguration	Get, Set	Manage DFS Replication service settings

DfsrState	Get	View the DFS Replication state for a member

Chapter 9

Internet Information Services

Managing sites

Application pools

IIS users and delegation

Managing FTP

Managing IIS using PowerShell

Many organizations use Internet Information Services (IIS) to host both their public-facing website as well as internal intranet services. IIS also supports many different roles, from Outlook Web Access in Exchange Server deployments, to SharePoint, Windows Server Update Services, and Certificate Services Web enrollment.

In this chapter, you learn about how to perform Web server–administration tasks, from basic tasks, such as

configuring new websites and virtual directories, to more advanced topics, such as configuring application pool recycling settings, setting FTP sessions to require TLS, and delegating administrative privileges.

MANAGING SITES

You perform website management in IIS through the IIS Manager console. The key to understanding the IIS Manager console, shown in Figure 9-1, is the Details pane that allows you to configure items depending on the configuration node selected. You are presented with the Details pane when selecting a particular node in the Connections pane, such as the Web server itself, the Applications Pool node, or individual websites. For example, if you open the Authentication node at the Server level, it configures authentication options for the server. If you have a specific website selected, opening the same Authentication item configures authentication for only that website. If you want to configure an item that isn't present, you need to verify that the relevant role service is installed.

Technet24

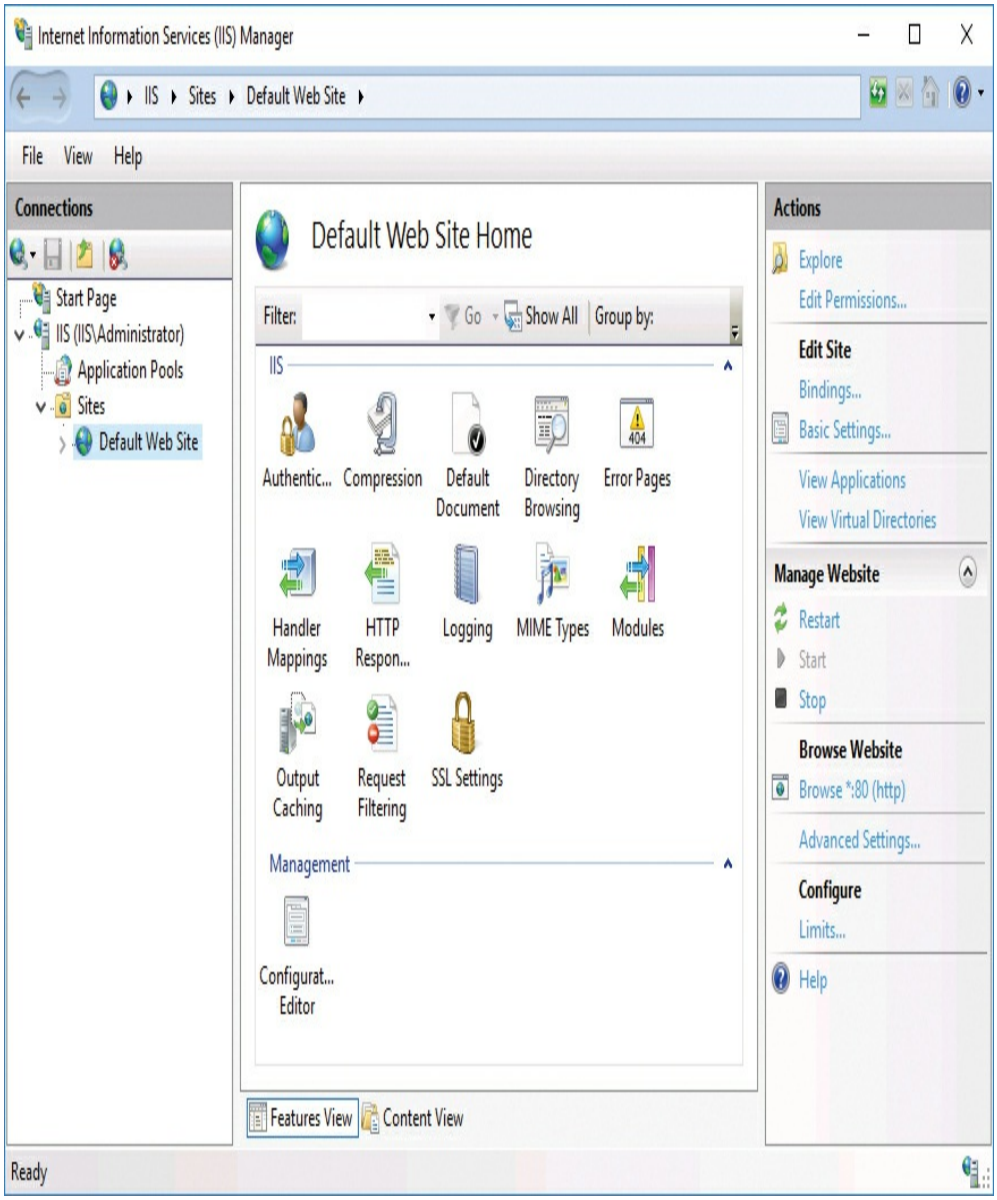


Figure 9-1 IIS Manager

Inside OUT

Missing items

Specific configuration items are only present within IIS manager when the appropriate role service is installed. IIS installs in a minimal configuration. If you can't find a specific option that you need to configure for a web application that you are hosting, check that the appropriate role service is installed.

Adding websites

The only practical limit for the number of websites that a server running IIS can host is hardware related. Although you can probably run 10,000 different sites off the same computer or virtual machine, the hardware resources needed to support that number of sites in order to run in a fast and responsive manner is prohibitively expensive.

To run more than one site on IIS, the sites need to be different in one of the following ways:

- **Unique IP address.** You can configure a different site for each IP address associated with the server. Although this isn't useful when you use one of the limited supply of IPv4 addresses, configuring different sites for each IP address associated with the server is useful when it comes to the more plentiful IPv6 addresses. You can use IP address differentiation to ensure that website traffic uses a specific network adapter.
- **Hostname.** If your server has only one IP address, you can still run multiple sites. You just need to ensure that each site uses a unique hostname. IIS parses the hostname in the HTTP request and directs the client to the appropriate site.
- **Port number.** You can also configure different sites that use the same IP address and hostname by assigning them different port addresses. The port address is a less popular way of differentiating websites; although it is an effective technique, most people have difficulty remembering to add a separate port address to a URL.

To add a website to an existing server, perform the following steps:

1. Open the IIS Manager console from the Tools menu of the Server Manager console.
2. Right-click the Server node, and then click Add Website.
3. In the Add Website dialog box, shown in Figure 9-2, enter the following information:
    1. **Site Name.** The name that represents the site on the server.
    2. **Application Pool.** The application pool with which the site is associated. The default settings create a new application pool for each new site. You can click the Select button to select a different application pool.
    3. **Physical Path.** This is where the website files are stored. You can configure site files to be stored on a remote network share. When you do this, you need to specify the authentication credentials used to connect to the network share or the file system. The default credentials use the web user's identity, so you'll need to configure authentication settings if you aren't using a specific account.
    4. **Type.** Allows you to specify whether the site uses HTTP or HTTPS.
    5. **IP Address**. Determines whether the site uses all unassigned IP addresses, or a specific IP address is assigned to the server.
    6. **Port.** Determines whether the site uses the default port 80 or 443 for HTTPS sites.
    7. **Host Name.** The hostname associated with the site.
    8. **SSL Certificate.** Allows you to select the TLS certificate associated with the site. This option is only available if you select HTTPS at the time. Certificates can be imported if necessary.

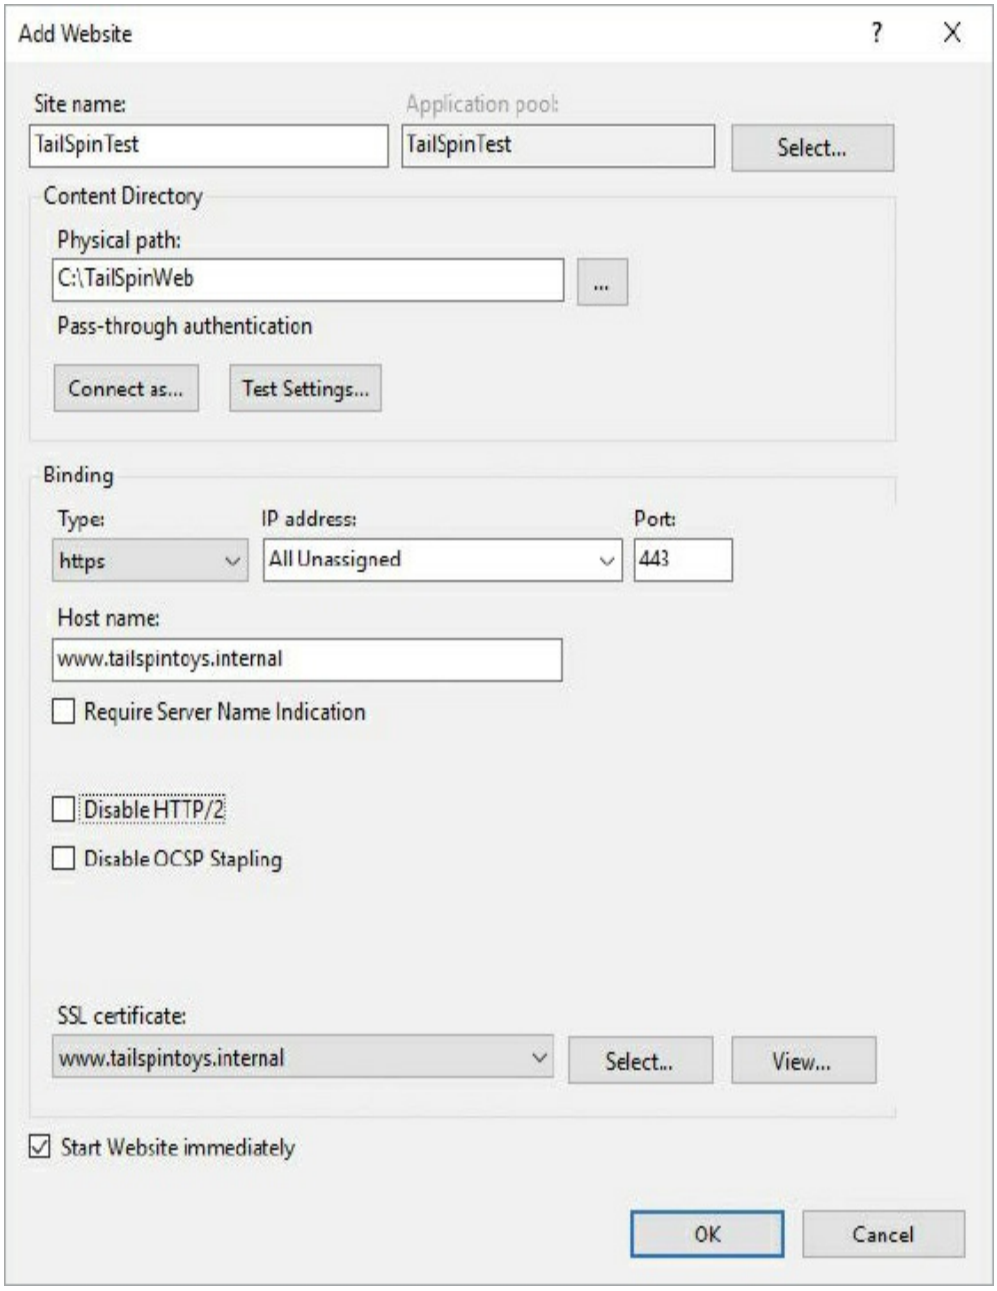


Figure 9-2 Add Website dialog box

Inside OUT

Test certificates

You can create a self-signed TLS certificate using the New-SelfSignedCertificate

cmdlet. For example, to create a self-signed certificate for the website www.tailspintoys.internal, run the command New-SelfSignedCertificate -DNSName www.tailspintoys.internal -CertStoreLocation “cert:\LocalMachine\My”.

Although you are more likely to use PowerShell to manage IIS on Windows Server 2019, long-time IIS administrators often have scripts that call the *appcmd.exe* utility, which was available in several previous versions of Windows Server and is located in the *c:\Windows\System32\inetsrv* folder, to perform administrative tasks. For example, you can use *appcmd.exe* to create a new website on port 80 associated with the hostname *www.contoso.internal* and using the *c:\contoso* folder using the following command:

Click here to view code image

```
Appcmd.exe add site /name:contoso /physicalPath:c:\contoso
/bindings:http/*:80:www.
contoso.internal
```

You can view a list of sites running on a server using the command:

```
Appcmd.exe list sites
```

Virtual directories

You use virtual directories when you want to create a directory on the website that does not map to a corresponding folder in the existing website folder

structure. For example, the site *www.fabrikam.com* might map to the c*:\fabrikam* folder on a computer running IIS. If you create the folder *c:\fabrikam\products*, people can access that folder by navigating to the URL *www.fabrikam.com/products*. Instead, if you put the *Products* folder on another volume, you can use a virtual directory to map the URL *www.fabrikam.com/products* to that alternate location.

To add a virtual directory using the IIS Manager console, perform the following steps:

1. Open the IIS Manager console and navigate to the site for which you wish to add the virtual directory.
2. Right-click the site, and then click Add Virtual Directory.
3. In the Add Virtual Directory dialog box, shown in Figure 9-3, enter the Alias for the virtual directory and the Physical Path to the virtual directory. If you are specifying a remote network share, you can configure which account is used to connect to that share.

Technet24

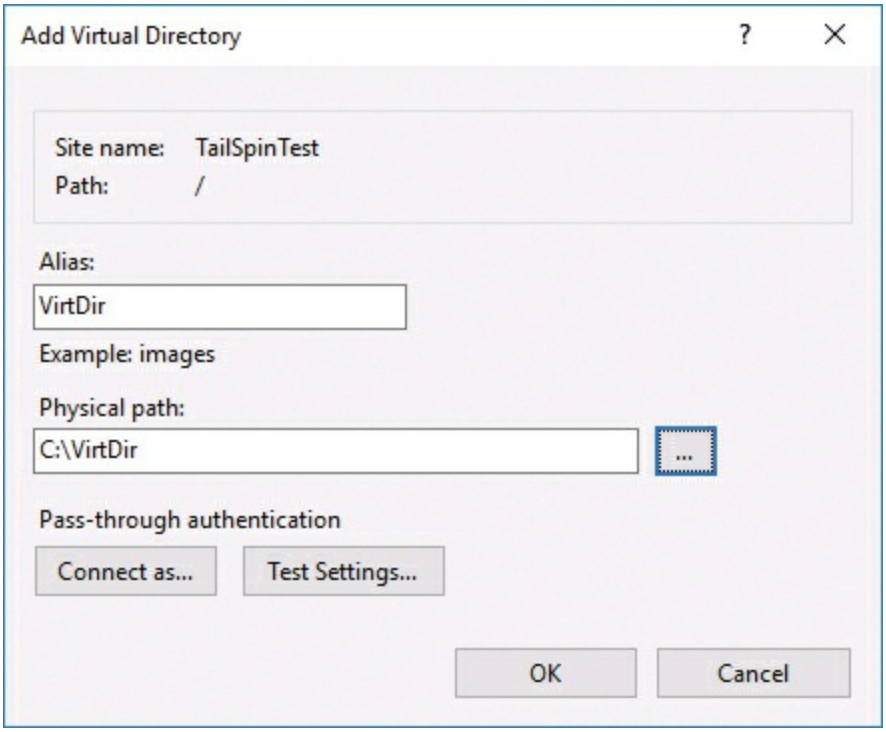


Figure 9-3 Add Virtual Directory dialog box

Modifying site settings

After you create a website, you can change the settings of the site by viewing its Advanced Settings, as shown in Figure 9-4. This can be done by clicking the site in the IIS Manager console and clicking the Advanced Settings item in the Actions pane.

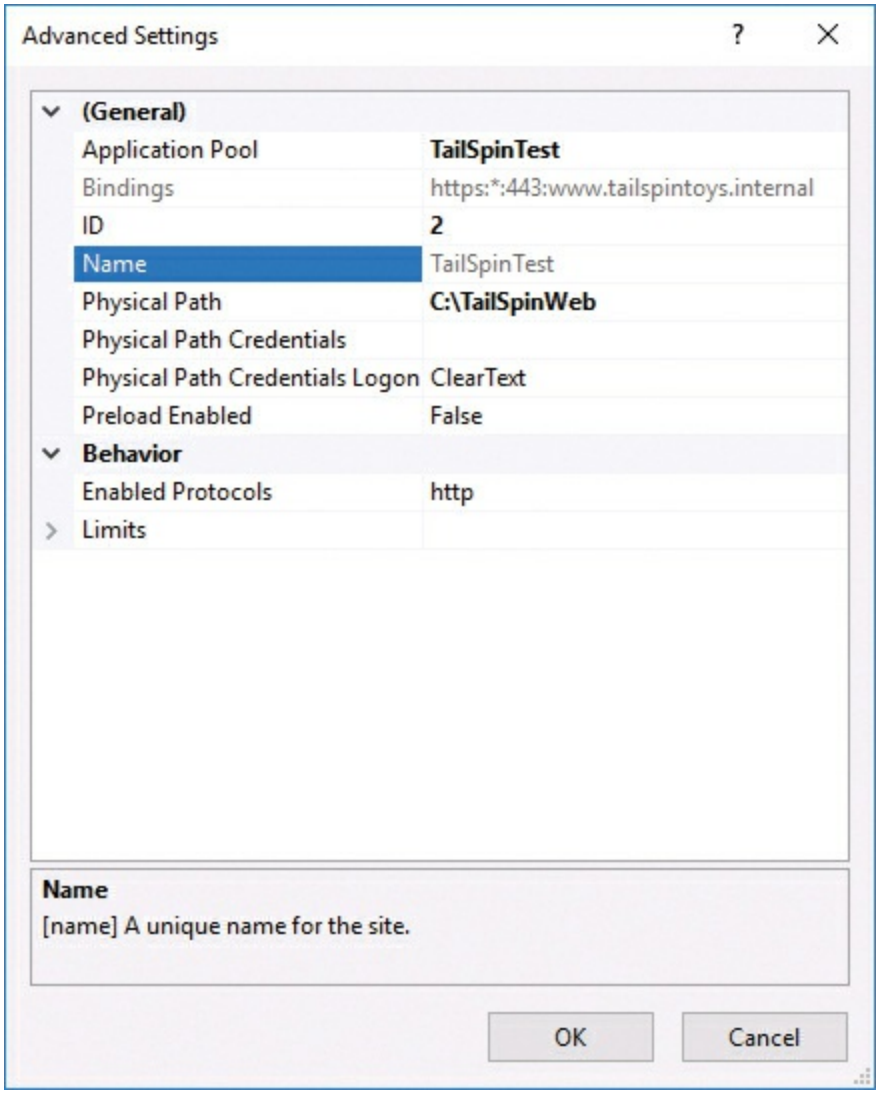


Figure 9-4 Advanced Settings

Adding web applications

A web application is a collection of content, either at a website's top level or in a separate folder under the website's top level. Applications can be collected in application pools. Application pools allow you to isolate one or more applications in a collection. When you

create a website, it automatically creates a web application associated with that website. This web application has the same name as the website.

To create a web application, perform the following steps:

1. Open the IIS Manager console and navigate to the site for which you want to create a web application.
2. Right-click the site, and then click Add Application.
3. In the Add Application dialog box, shown in Figure 9-5, enter the alias of the application, choose which Application Pool the application is a member of, and specify the Physical Path to the application on the server.

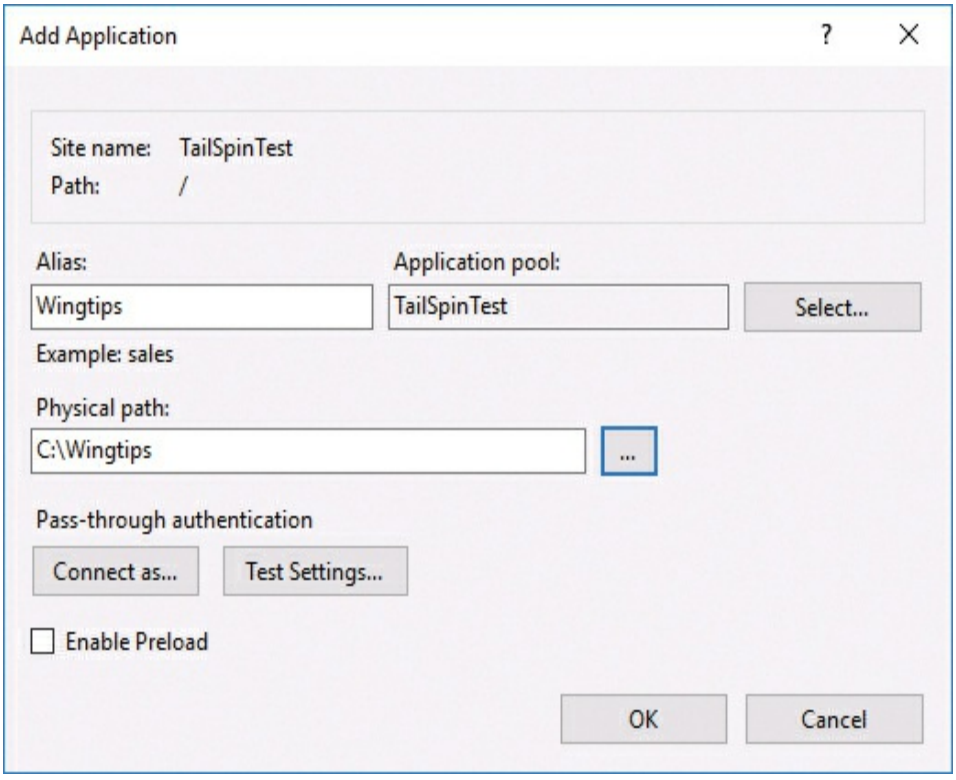


Figure 9-5 Add Application

Configuring TLS certificates

TLS certificates—often inaccurately referred to as “SSL certificates” because TLS is the successor protocol to SSL—allow clients to encrypt their session with the web server, minimizing the chance that communication between the client and the server can be intercepted; also, TLS certificates serve as an identity verification mechanism. You can configure TLS to be required at the server, site, web application, virtual directory, folder, and file level.

The first step to take prior to requiring TLS on a particular site is to request and then install a certificate. To do this, perform the following steps:

1. Open the IIS Manager console, select the Server node, and then double-click the Server Certificates item in the Details pane.
2. Click the Create Certificate Request item. In the Request Certificate dialog box, shown in Figure 9-6, enter the name for the certificate and your organization’s details and click Next.

Technet24

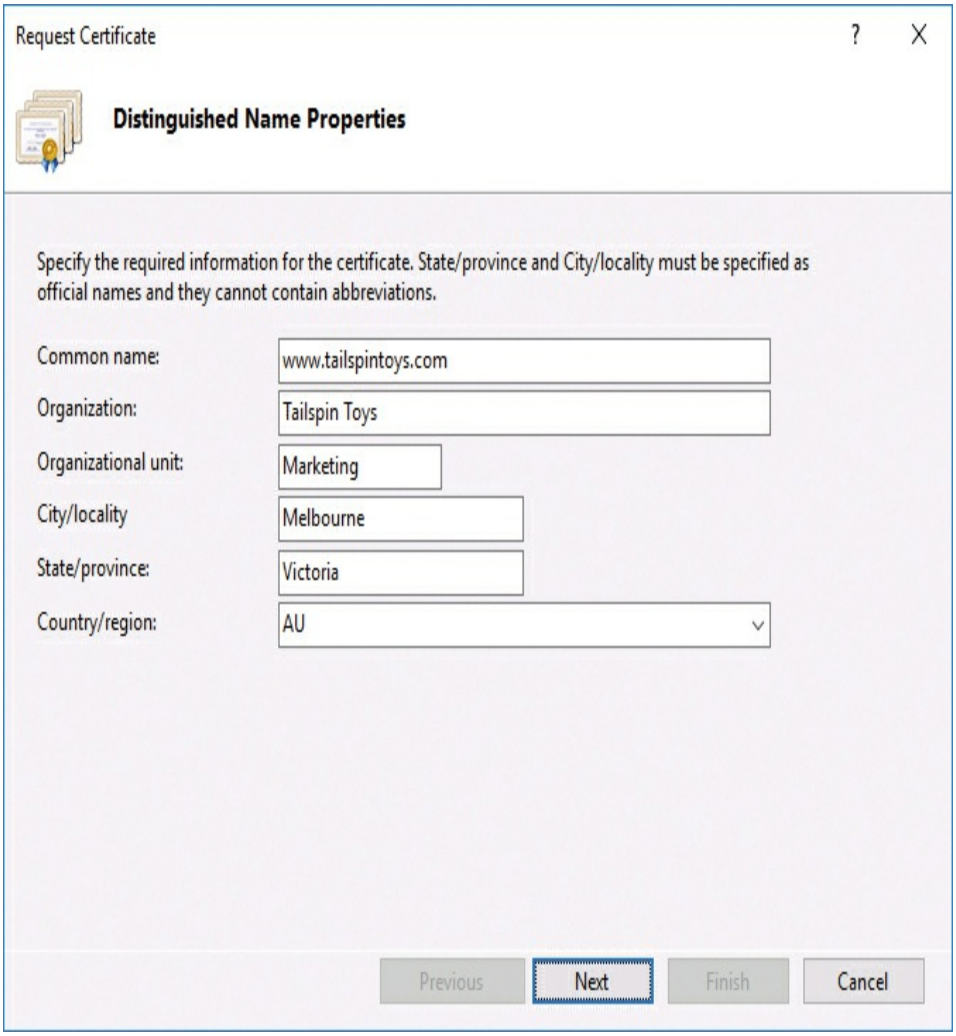


Figure 9-6 Request Certificate

3. On the Cryptographic Service Provider Properties page, enter the cryptographic service provider and the key length that you wish to use for the TLS certificate, and then click Next.
4. Enter a name for the certificate request file. After the certificate request file is created, forward it to a CA.
5. After the certificate has been issued, open the Server Certificates item, and then click Complete Certificate Request.
6. On the Specify Certificate Authority Response page of the Complete Certificate Request Wizard, specify the location of the certificate issued by the CA and a friendly name with which to identify the certificate, and

then click OK. This installs the certificate.

After the certificate is installed, you can configure the site, virtual directory, folder, or file to request or require HTTPS. To accomplish this goal, perform the following steps:

1. In IIS Manager, select the site on which you wish to configure TLS.
2. Click Edit Bindings, and then click Add. In the Edit Site Binding dialog box, ensure that HTTPs is selected in the Type drop-down menu, and the TLS certificate that you installed is selected in the SSL Certificate drop-down menu, as shown in Figure 9-7.

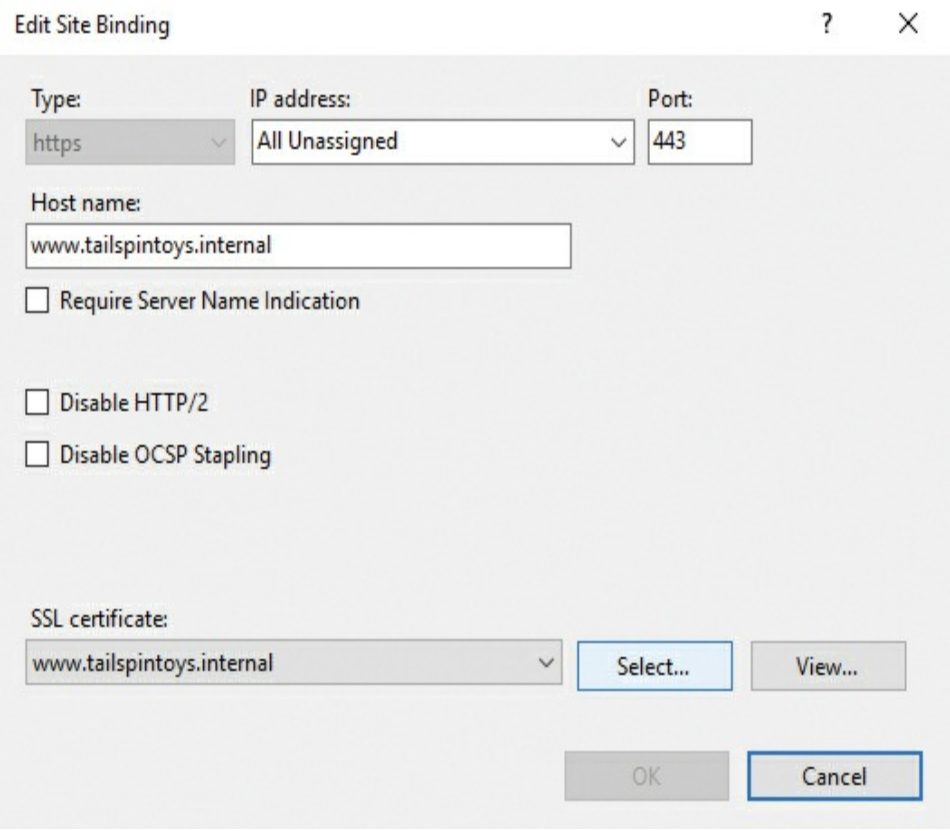


Figure 9-7 Edit Site Binding dialog box

3. Select the site, virtual directory, folder, or file on which you wish to enforce TLS settings.

Technet24

4. Double-click the SSL Settings item in the Details pane. (This should be named “TLS Settings,” but “SSL” remains the commonly used term).
5. On the SSL Settings page, ensure that Require SSL is selected, and then click Apply.

Site authentication

Authentication methods determine how a user authenticates with a server or a website. You can configure authentication at the server level or at the site level. If there is a conflict between the authentication methods configured at the server and site levels, the site-level authentication settings have precedence. IIS supports the following authentication methods:

- **Active Directory Client Certificate.** This form of authentication works by checking client certificates. These are almost always issued by an internal certificate authority.
- **Anonymous Authentication.** This is the most typical form of authentication for a web server. Clients can access the webpage without entering credentials.
- **ASP.NET Impersonation.** Use ASP.NET impersonation when it is necessary to execute an ASP.NET application under a different security context.
- **Basic Authentication.** Basic authentication works with all web browsers but has the drawback of transmitting unencrypted base64-encoded passwords across the network. If you use Basic Authentication, you should also configure the site to require an encrypted TLS connection.
- **Digest Authentication.** Digest authentication occurs against a Domain Controller and is used when you want to authenticate clients that may be accessing content through a proxy, something that can be problematic for clients if you configure a server to use

Windows authentication.

- **Forms Authentication.** This method is used when you redirect users to a custom webpage on which they enter their credentials. Once they've been authenticated, they are returned to the page that they were attempting to browse.
- **Windows Authentication.** Use this method for intranet sites when you want clients to authenticate using Kerberos or NTLM. It works best when client computers are members of the same forest as the computer hosting the web server role.

Inside OUT

Authentication methods

Authentication options are available only if you have installed them. If you can't find the method that you want to enable, ensure that you have installed it.

Remember that you need to disable Anonymous Authentication if you want to force clients to use a different authentication method. Web browsers always request content from a server anonymously the first time you attempt to access content. If Anonymous Authentication is enabled, other forms of authentication are ignored. To enable a specific form of authentication for a website, perform the following steps:

1. Open the IIS Manager console from the Administrative Tools menu.
2. Click the Web Server node. In the Content pane, scroll down and double-click the Authentication item.
3. In the list of available authentication technologies, shown in Figure 9-8, click the authentication technology that you want to enable, and then click Enable in the Actions pane.

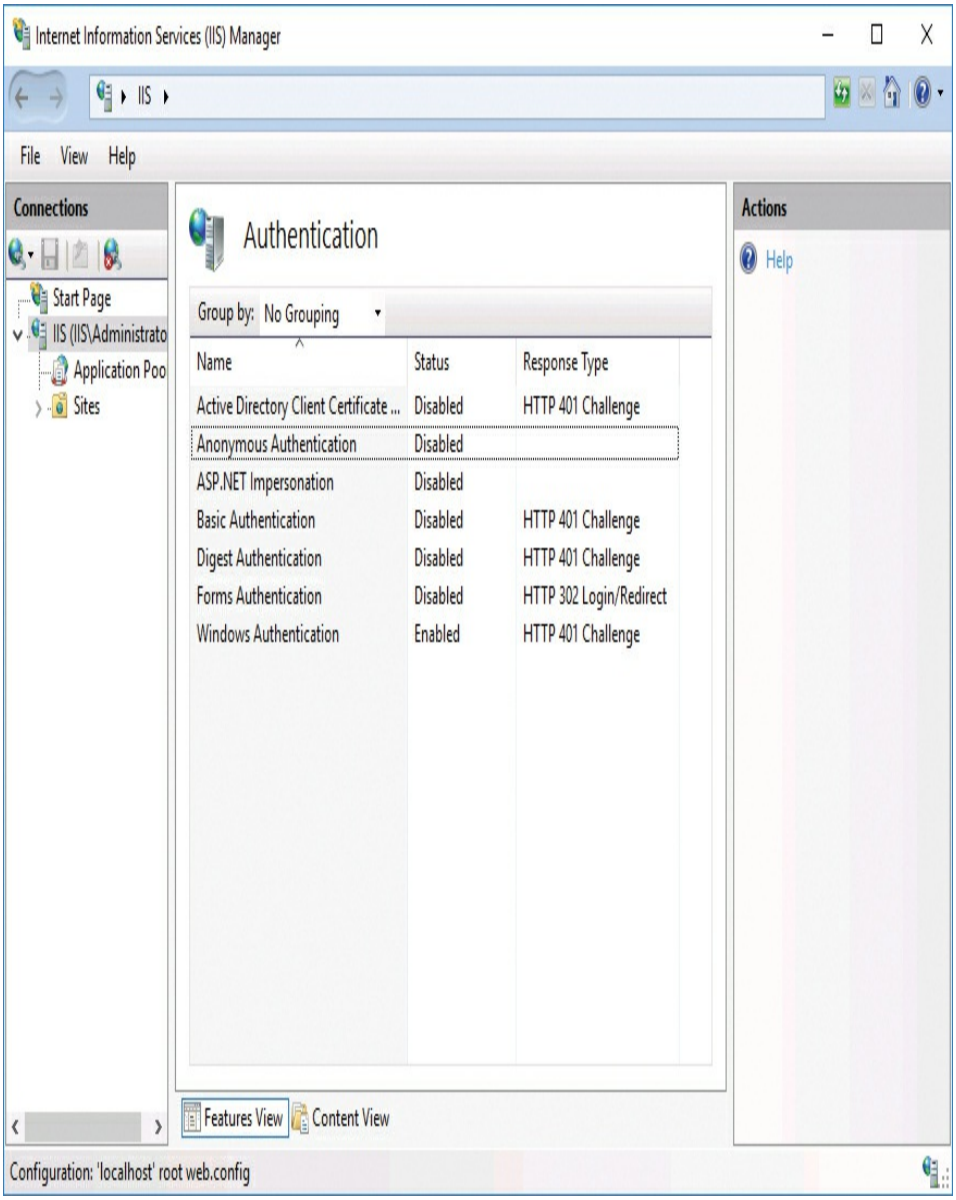


Figure 9-8 Configure authentication

Modifying custom error response

Many sites customize their error responses so that users are provided with something more meaningful than a simple error message. Although the default error messages are perfectly serviceable, many organizations

think they lack character and like to customize the most popular error messages, such as the 404 error message displayed by the web server when a page is not found.

To modify custom error message settings, perform the following steps:

1. Open the IIS Manager console and select the Server node.
2. Scroll down the Details pane and double-click the Error Pages item.
3. On the Error Pages pane, select the error code that you want to modify, and then click Edit.
4. In the Edit Custom Error Page dialog box, shown in Figure 9-9, choose to forward clients to a static page based on file location, forward clients to a URL on the server, or to forward clients to an absolute URL, which can be a page on another website.

Technet24

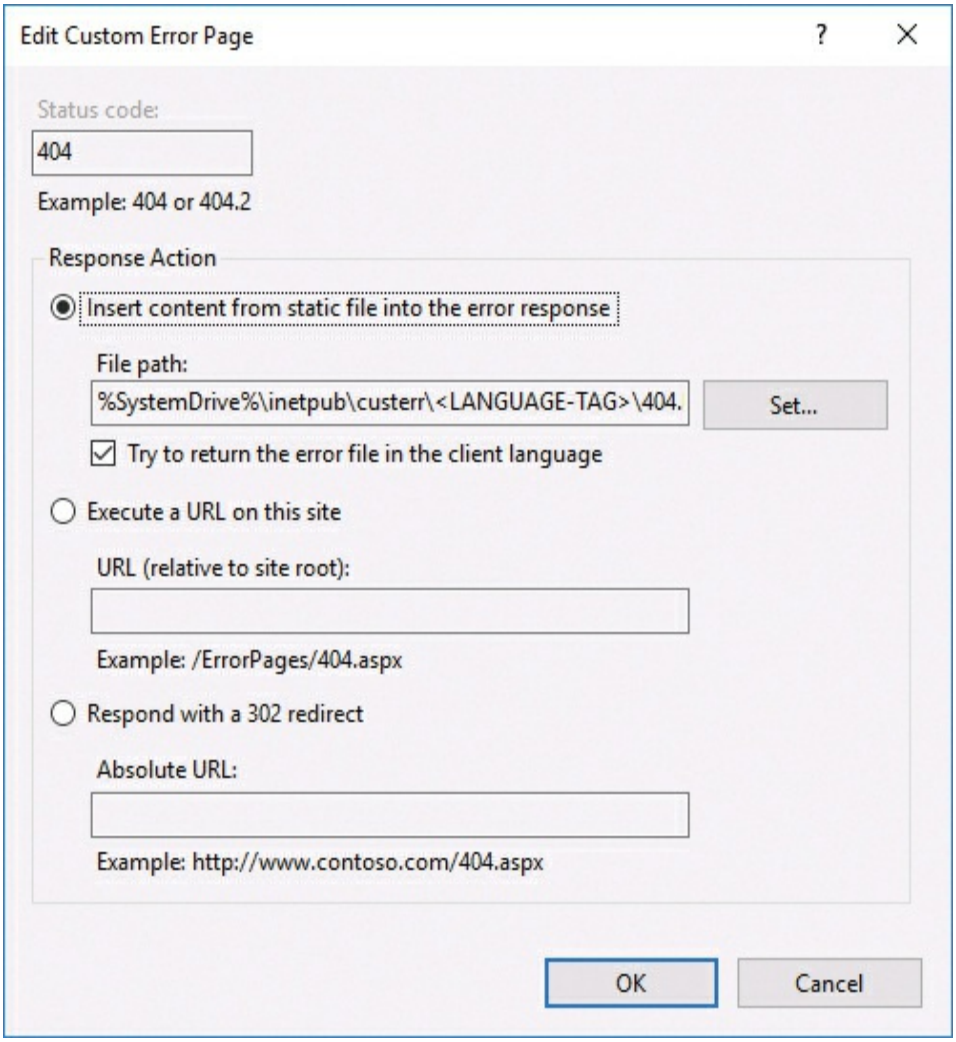


Figure 9-9 Edit Custom Error Page

Adding or disabling the default document

The default document is the document that loads when a client navigates to a webpage but doesn't specify which page to load. The Default Documents role service is installed in a default installation of IIS, though this can be removed if you so choose. On servers running IIS, the following default documents are used:

- default.htm
- default.asp
- index.htm
- index.html
- iisstart.htm
- default.aspx

If more than one of these documents is present in the same folder, the document that is higher on the list overrides the documents that are lower on the list. You can add, reorder, or disable default documents using the IIS Manager console by double-clicking the Default Document item, either at the server or the website level, as shown in Figure 9-10.

Technet24

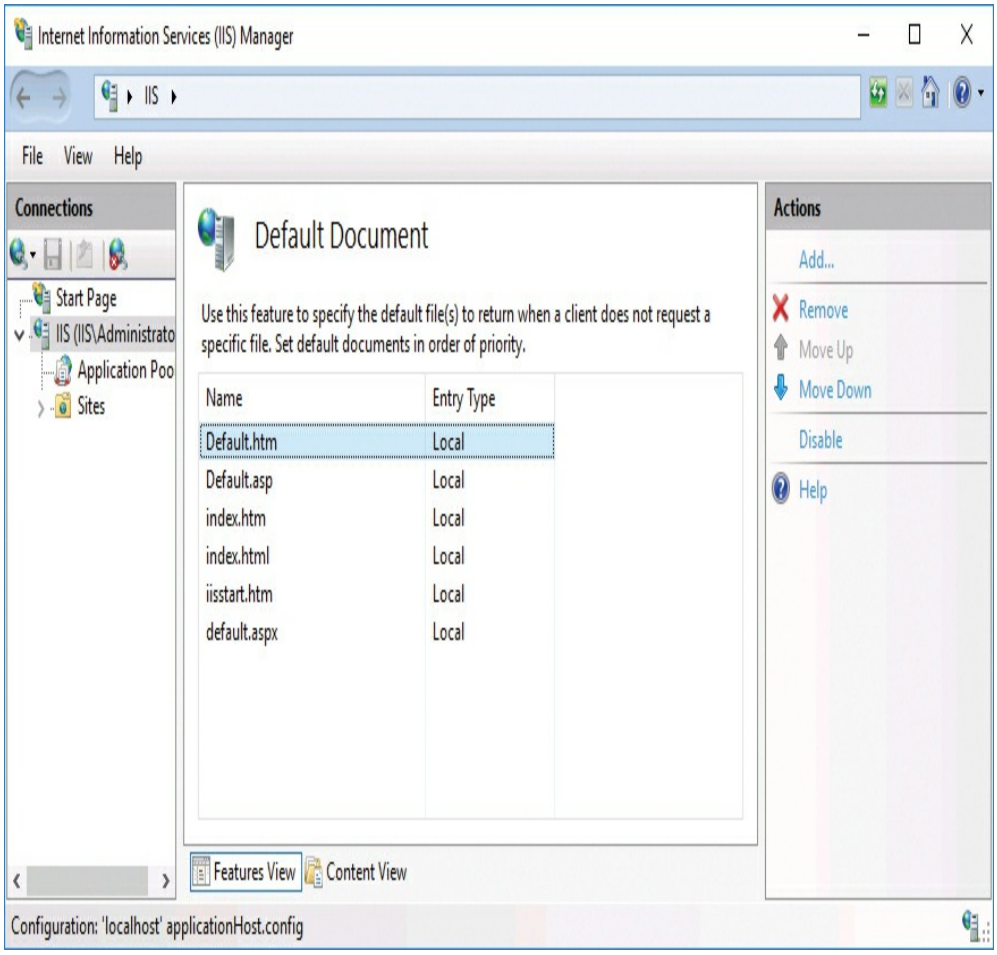


Figure 9-10 Default Document

Directory browsing

Directory browsing allows clients to view the files that are stored in a folder. Directory browsing works only if the default document is not present or has been disabled. To enable directory browsing, the Directory Browsing role service must be installed. To configure directory browsing, perform the following steps:

1. Open the IIS Manager console and select the site on which you want to

enable directory browsing.

2. Scroll down in the Details pane and double-click the Directory Browsing item.
3. When Directory Browsing is shown in the Details pane, click Enable, and select the items that you want to allow to be displayed when the client is browsing. By default, the file name, time, size, extension, and date information are displayed.

IP address and domain name filtering

IP address restrictions and domain name restrictions allow you to block or allow access to a site based on a client's address. You can configure two types of rules: Allow Rules and Block Rules. When you configure an Allow Rule, only hosts that are on the Allow List can access your website. When you configure a Block Rule, hosts with an address on your Block List are blocked from accessing your site. You can configure Allow and Block rules at the server, site, application, virtual directory, folder, and file level in IIS. You must have the IP And Domain Restrictions role service installed to use IP address and domain name filtering.

To configure an IP address and domain name restriction, perform the following steps:

1. Open the IIS Manager console and select the site on which you want to implement an Allow List.
2. Double-click the IP Address And Domain Restrictions item.
3. If you want to configure an Allow List entry, click Add Allow Entry, and in the Add Allow Restriction Rule dialog box, shown in Figure 9-11, enter the IP Address Range or Specific IP Address that you want to

Technet24

allow.

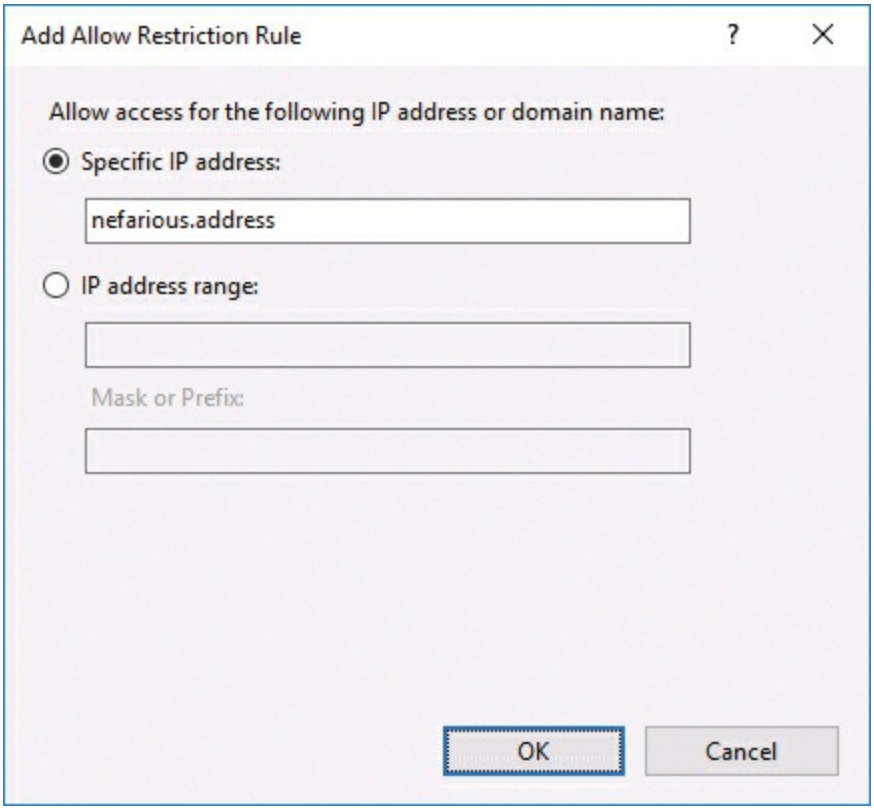


Figure 9-11 Add Allow Restriction Rule

4. If you want to configure a Deny entry, click Add Deny Entry, and in the Add Deny Restriction Rule dialog box, enter the Specific IP Address range or IP Address Range that you want to block.

Dynamic IP restriction settings allow you to configure IIS to automatically block traffic based on suspicious request behavior, such as what can occur during a denial of service attack. As Figure 9-12 shows, you can choose to block traffic based on a maximum number of concurrent requests or a maximum number of requests over a specific period of time.

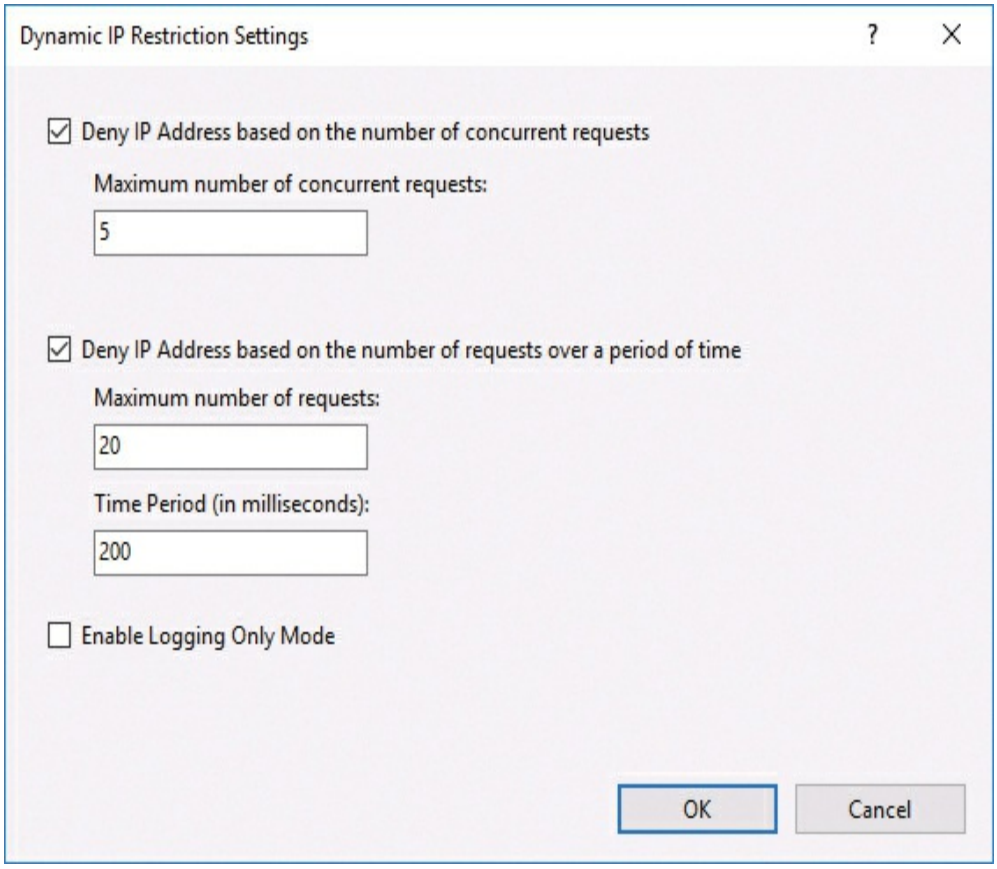


Figure 9-12 Dynamic IP Restriction Settings

URL authorization rules

URL authorization rules allow you to grant or deny access to specific computers, groups of computers, or domain access to sites, web applications, folders, or specific files. For example, you can use URL authorization rules to block everyone except the Managers group from accessing specific pages on the organization's intranet server. You can also configure URL authorization rules to apply only when the client attempts to use specific HTTP verbs, such as GET or

Technet24

POST. To use URL authorization rules, ensure that the URL Authorization role service is installed on the server. There are two types of authorization rules: Allow Rules and Deny Rules.

To create an authorization rule, perform the following steps:

1. Open IIS Manager and navigate to the site, virtual directory, folder, or file on which you want to create the authorization rule.
2. Double-click the .NET Authorization Rules item in the Details pane.
3. Click Add Allow Rule or Add Deny Rule, depending on the type of authorization rule that you wish to create.
4. In the Add Allow Authorization Rule or Add Deny Authorization Rule dialog box, specify the users or groups for which you wish to configure the rule and any specific HTTP verbs that you wish to trigger the rule. Figure 9-13 shows the Add Allow Authorization Rule dialog box.

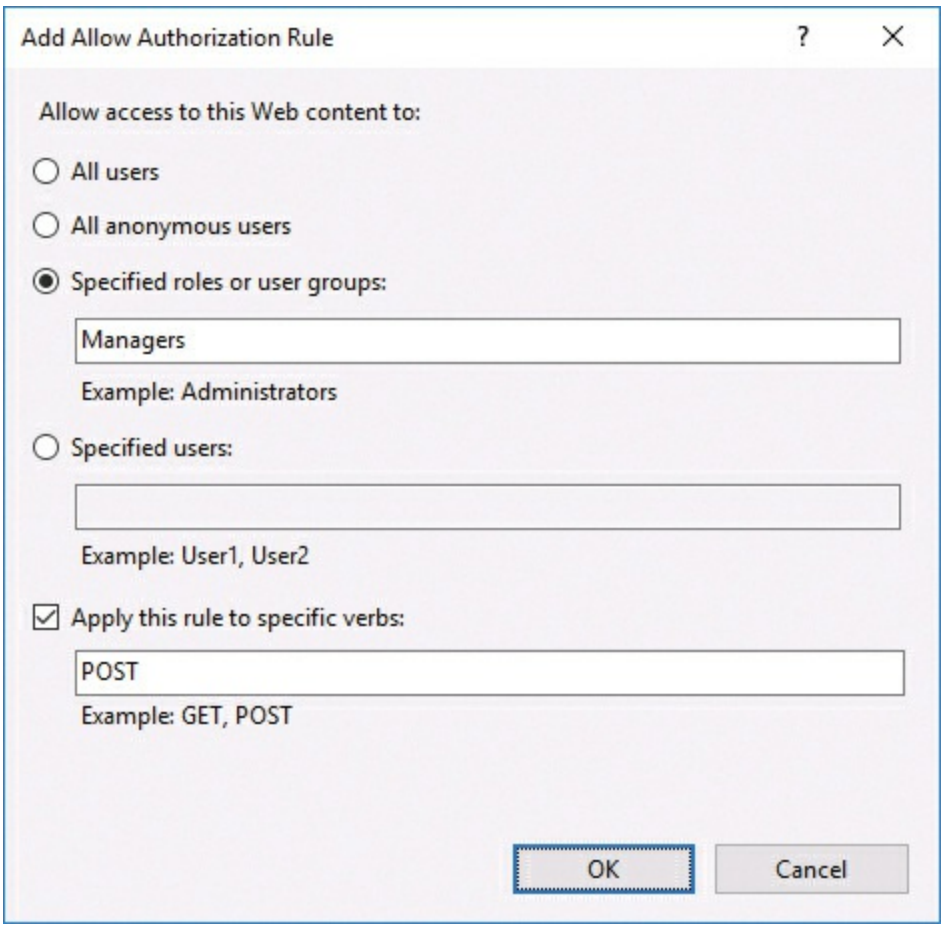


Figure 9-13 Add Allow Authorization Rule dialog box

If you want to block access to everyone except a specific group of users, remove the default All Users Allow rule and replace it with the rule that grants specific access to the security principals in question.

Request filters

Request filters allow you to block clients from making certain types of requests of your web server. Request filters allow you to block a common form of attack against web applications—one where the attacker enters

Technet24

a specially formatted HTTP request in order to elicit an unplanned response from the web application. SQL injection attacks are one such kind of specially formatted HTTP request. There are several different types of request filters you can configure using the Request Filtering item. Request filters can be configured at the server, site, and file level. To use request filtering, ensure that you install the Request Filtering role service on the server.

You can configure the following types of request filters:

- **File Name Extensions.** This type of filter allows you to specify which type of files clients are not allowed to request from the web server on the basis of a file extension. The default request filtering settings block 43 different file extensions, and it is relatively easy to add additional extensions by clicking Deny File Name Extension on the File Name Extensions tab of the Request Filtering item.
- **Rules.** This allows you to configure specific rules that check either the URL, query string, or both. If a match occurs, the request is blocked.
- **Hidden Segments.** Hidden Segments allow you to set a list of URL segments that will be blocked to requesting clients. A URL segment is the section of a URL path that lies between slash (/) characters.
- **URL.** You can specify a set of URL sequences (such as *setup/config.xml*) that the Request Filtering role service blocks to clients. All instances of the URL sequence are blocked within the scope of the filter.
- **HTTP Verbs.** You can restrict which HTTP verbs can be used in requests by clients. For example, you might block clients from using the PUT verb in requests.

- **Headers.** You can specify a maximum URL size for a specific HTTP request header.
- **Query.** This allows you to filter based on the content of query strings.

You can modify the general request filtering settings, as shown in Figure 9-14, by clicking Edit Feature Settings in the Actions pane. Using this dialog box, you can configure the following options:

- **Allow Unlisted File Name Extensions.** If this box is not checked, file extensions that are not explicitly allowed cannot be requested by clients.
- **Allow Unlisted Verbs.** If this box is not checked, HTTP verbs that are not explicitly allowed are blocked.
- **Allow High-Bit Characters.** If this box is not checked, unusual non-ASCII characters are blocked.
- **Allow Double Escaping.** If this box is not checked, URLs can contain double escape characters.
- **Maximum Allowed Content Length.** This text box specifies the maximum size of content that a request can process. Keep this figure small if you do not allow HTTP uploads to your site.
- **Maximum URL Length.** This allows you to restrict the maximum size of a URL request sent to your site.
- **Maximum Query String.** This allows you to restrict the maximum query string size in a URL request sent to your site.

Technet24

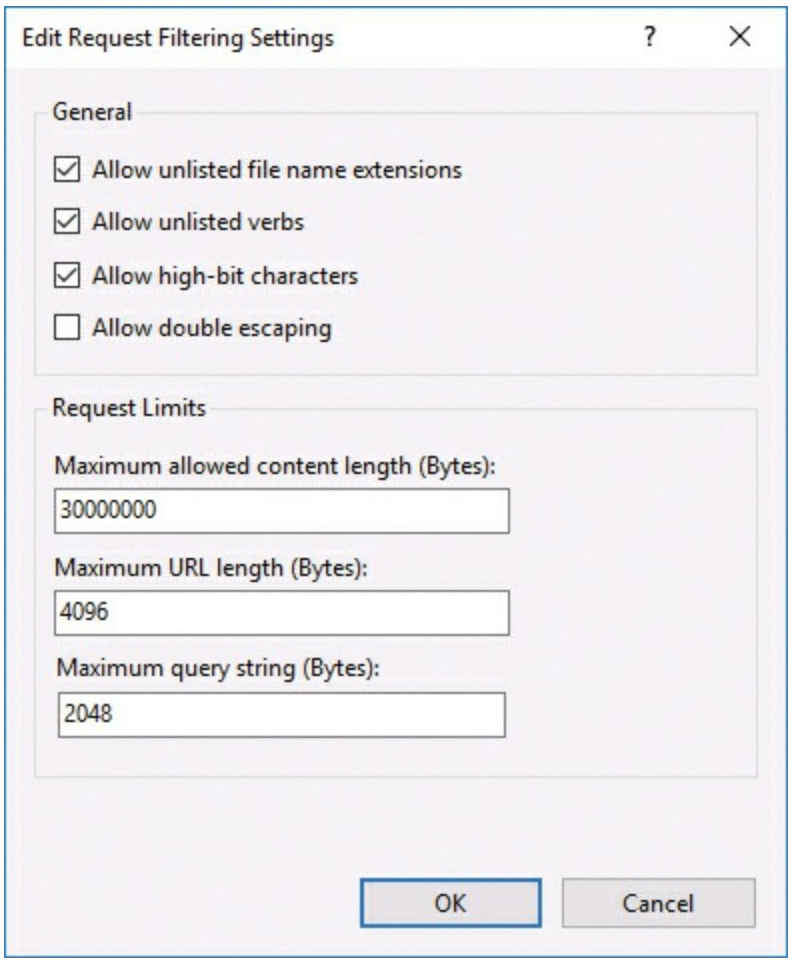


Figure 9-14 Edit Request Filtering Settings

APPLICATION POOLS

Application pools are collections of web applications that are served by a worker process or a group of worker processes. When you create a new website in IIS, IIS automatically creates an associated application pool. One of the ideas behind application pools is that by separating web applications into their own pools, the failure of one web application does not impact the

functioning of other web applications. Application pools allow you to do the following:

- Collect sites and applications that use the same configuration settings.
- Improve security by allowing the use of a custom security account to run an application.
- Silo resources in such a way that one web application cannot address or influence resources in another web application.

Creating application pools

Although new application pools are created each time you create a new site, there are reasons why you want to configure a new application pool. Prior to creating a new application pool, you need to know what pipeline mode you want to use to process requests for managed code. Managed pipeline mode determines how IIS processes requests for managed code. The difference between the two managed pipeline modes is as follows:

- **Integrated.** This mode allows the request-processing pipelines of IIS and ASP.NET to process the request. Most web applications are likely to use this mode.
- **Classic.** This mode forwards requests for managed code through *aspnet_isapi .dll*. This is the same method used for processing requests on servers running older versions of IIS. Use this mode if an older web application that you have migrated to IIS on Windows Server 2019 does not function well in integrated mode.

To create a new application pool, perform the following steps:

1. Open the IIS Manager console, and then click the Application Pools node.
2. Click Add Application Pool in the Actions menu. This opens the Add Application Pool dialog box shown in Figure 9-15.

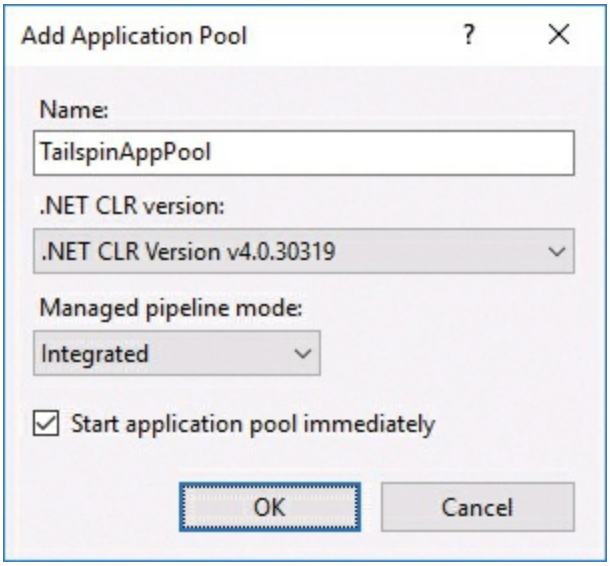


Figure 9-15 Add Application Pool

3. In this dialog box, provide a Name for the application pool, select the .NET CLR Version for the application pool, and choose the Managed Pipeline Mode.

Configuring application pool recycling settings

Not all developers write code that functions correctly over an extended period of time. If you find that a web application that you need to host on a computer running Windows Server 2019 starts to develop problems after it has run for a lengthy amount of time, you can deal with this by configuring the application pool recycling settings to recycle the processes related to the

application on a regular basis. You configure application pool recycling settings in the Edit Application Pool Recycling Settings dialog box shown in Figure 9-16. You can configure an application pool to be recycled when the following thresholds are met:

- After a certain amount of time has elapsed. By default, all application pools automatically recycle every 1,740 minutes
- After an application pool has processed a certain number of requests
- At specific times of the day
- After an application pool exceeds a certain virtual memory threshold
- After an application pool exceeds a specific private memory threshold

Technet24

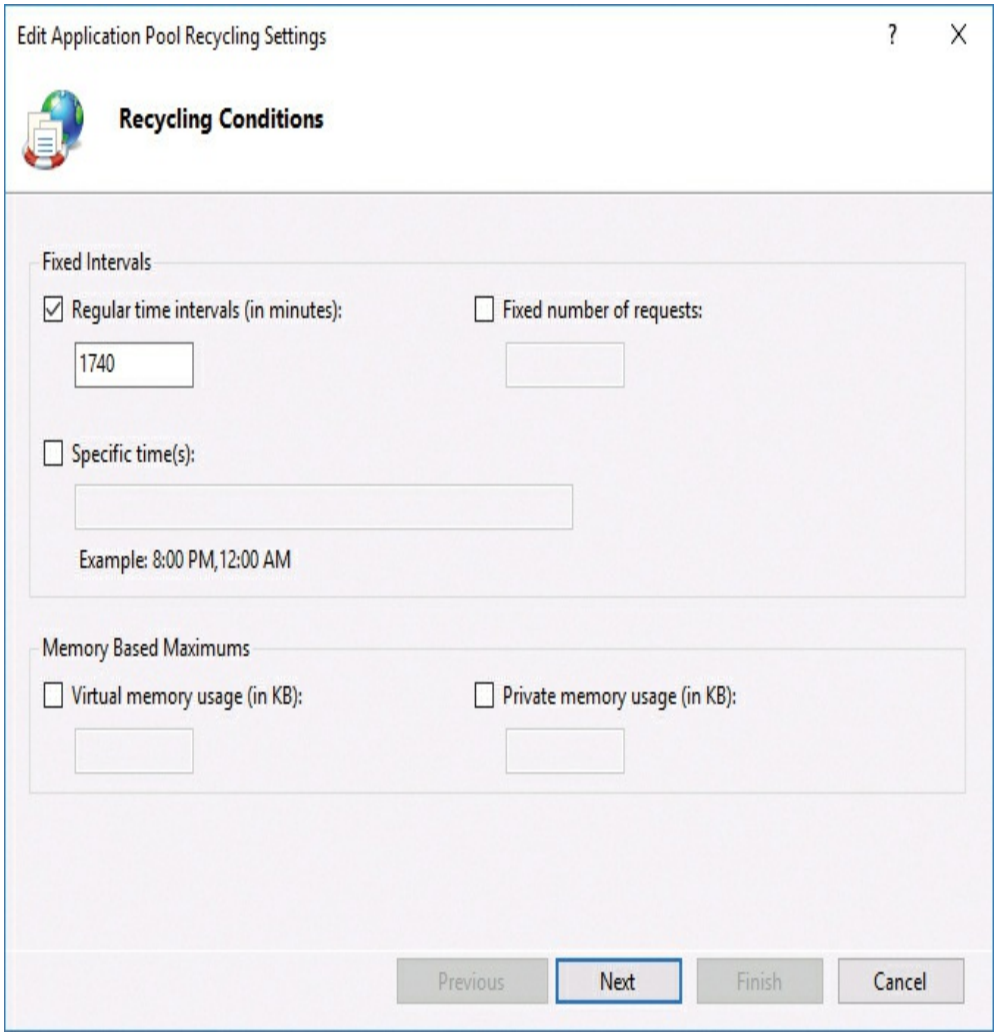


Figure 9-16 Application Pool Recycling Settings

Through the Advanced Settings dialog box, shown in Figure 9-17, you can configure basic settings such as the pipeline mode and recycling settings. You can also configure settings including processor affinity, CPU limits, process orphaning, request limits, and rapid fail protection settings.

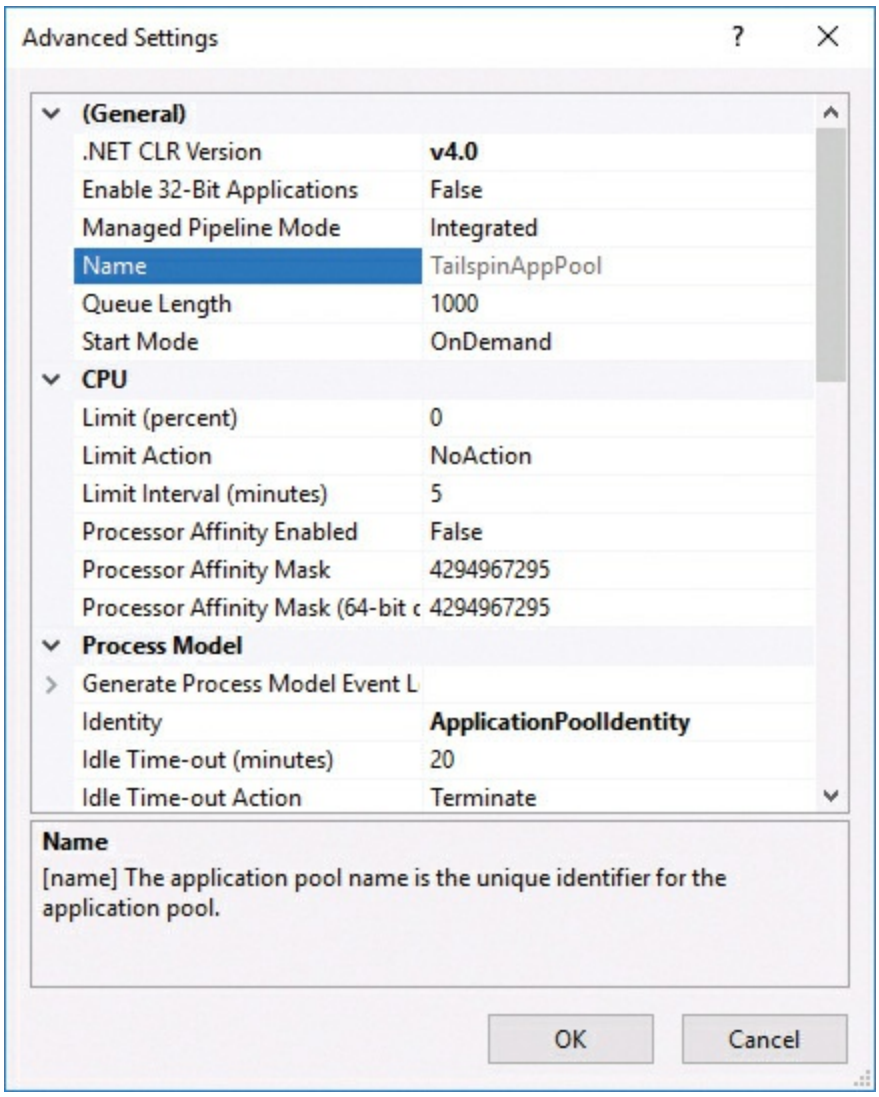


Figure 9-17 Application pool Advanced Settings

IIS USERS AND DELEGATION

IIS allows you to delegate the management of specific websites and web applications. This means that you can allow a specific user to manage a website without giving him logon privileges to the server or control over any other site or web application hosted by IIS.

Technet24

When you install IIS, it is configured so that it does not allow remote management using the IIS Manager console. You configure IIS to allow remote management using the Management Service item at the server level in the IIS Manager console. This involves installing the Management Service role service and ensuring that the service is configured to start automatically. Using the Management Service item, shown in Figure 9-18, you can configure the following remote management properties:

- **Enable Remote Connections.** This item must be enabled for remote administration connections using IIS Manager Console to work.
- **Identity Credentials.** This allows you to specify whether administrators can connect using Windows credentials or IIS Manager credentials. The default setting only allows Windows user accounts.
- **Connections.** This allows you to specify whether remote management connections can occur on any IP address interface, or whether they are limited to a specific IP address interface. You can also specify which SSL certificate will be used to protect remote administration sessions.
- **IP Address Restrictions.** These allow you to specify which IP addresses or networks may connect using IIS Administrator.

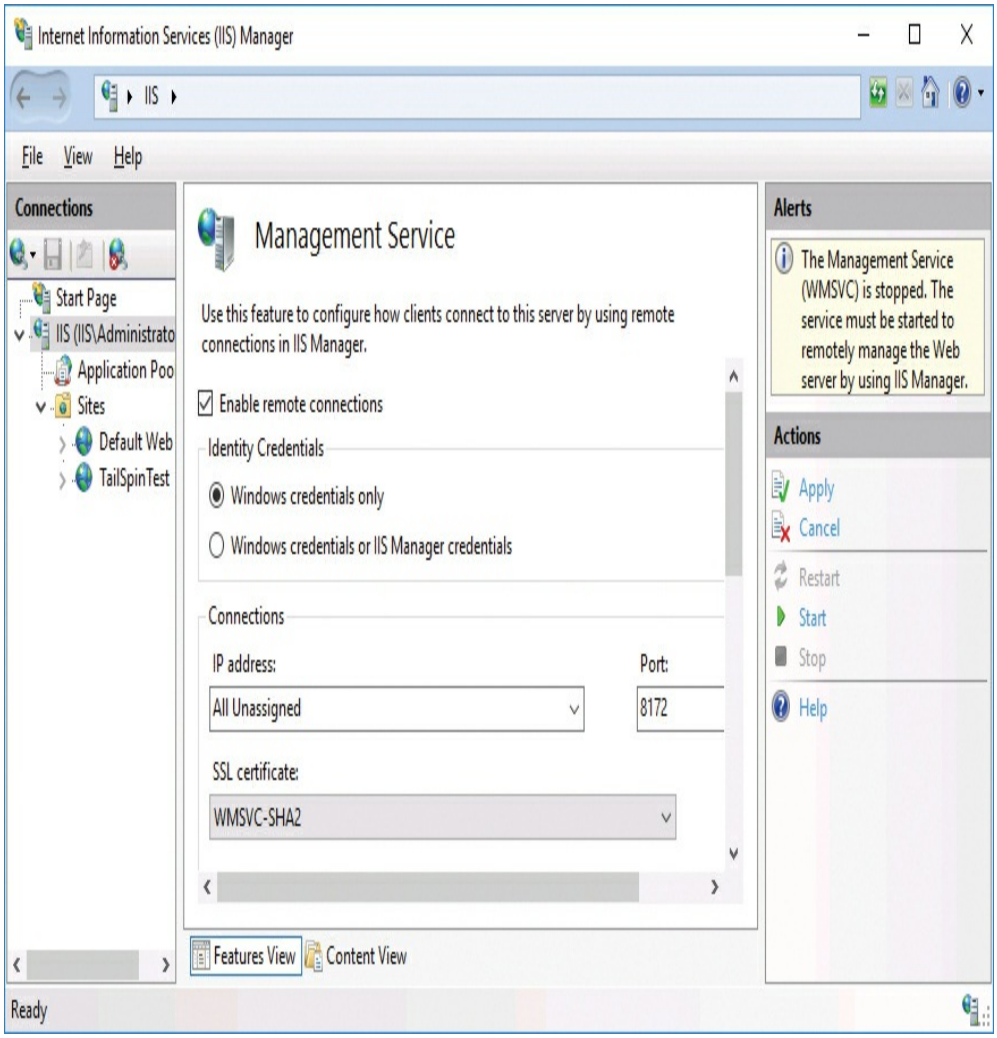


Figure 9-18 Management Service configuration

IIS user accounts

IIS allows you to create user accounts that exist only in IIS. These user accounts do not have any privileges outside IIS and cannot be used to log on locally to the server. This allows you to assign administrative permissions to users without having to create a corresponding machine local or domain user account.

To create an IIS Manager user, perform the following steps:

1. Open the IIS Manager console and double-click the IIS Manager Users item when the Server level is selected in the IIS Manager console.
2. Click Add User in the Actions pane. This opens the Add User dialog box. Enter the user name and password for the IIS user.

Delegating administrative permissions

After you have enabled remote management and, if you choose to create IIS users, configured the appropriate IIS user accounts, you can delegate management permissions on a site or web application through the IIS Manager Permissions item. You can check precisely which permissions users have been delegated by viewing the Feature Delegation item at the server level. You can also modify which configuration items are delegated using the Feature Delegation item.

To delegate administrative permissions over a site, perform the following steps:

1. Open the IIS Manager console and select the website for which you want to delegate administrative permissions.
2. Double-click the IIS Manager Permissions item in the Details pane.
3. In the Actions pane, click Allow User. In the Allow User dialog box, select the Windows or IIS Manager User, and then click OK. Figure 9-19 shows user *Rooslan* delegated permissions over the *TailSpinTest* site.

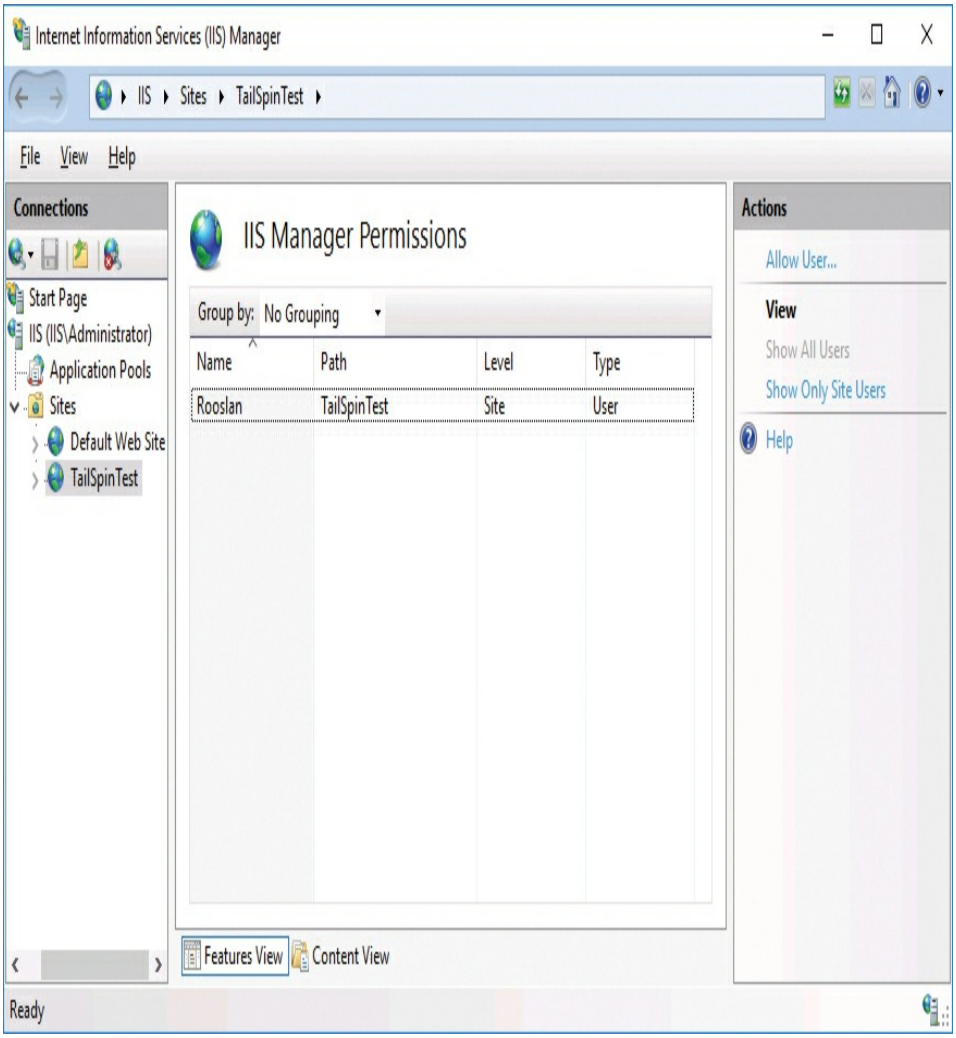


Figure 9-19 IIS Manager Permissions

MANAGING FTP

Although the FTP protocol is almost as old as the Internet itself, and newer ways of transferring files are becoming more popular, many organizations still use FTP as their primary method of uploading data to web servers such as IIS.

The FTP server in Windows Server 2019 supports TLS

connections and is fully managed from the IIS Manager console.

From the IIS Management console, you can configure the following:

- **FTP Authentication.** This allows you to enable or disable Anonymous or Basic authentication. The default settings have both forms of authentication disabled.
- **FTP Authorization.** This allows you to specify which users or groups are allowed to connect to the FTP server, and whether they have read and write permissions when using FTP.
- **FTP Directory Browsing.** This allows you to specify whether MS-DOS or UNIX-style directory listings are provided. It also allows you to specify whether virtual directory, available bytes, and four-digit years information are displayed.
- **FTP Firewall Support.** This allows you to configure the FTP server to accept passive connections from a perimeter network firewall.
- **FTP IPv4 Address and Domain Name Restrictions.** This allows you to block or allow FTP clients based on IPv4 address or client domain name.
- **FTP Logging.** This allows you to configure the FTP logging settings, including how often log files are rotated.
- **FTP Messages.** This allows you to configure FTP banner settings. This is the informational text that is presented to new FTP sessions.
- **FTP Request Filtering.** Using this item, you can configure the following:
    - Configure FTP to allow or block specific file name extensions from transfers.
    - Block segments, such as sensitive folders, from being accessible in FTP.

- Block particular FTP URLs.
- Block specific FTP commands from being used during a session.

- **FTP SSL Settings.** This allows you to specify which server certificate is used for TLS (though still called SSL) sessions. It also allows you to force the use of TLS as well as enforce 128-bit encryption on FTP SSL sessions.
- **FTP User Isolation.** This allows you to limit which folders an FTP user can navigate. You can configure a user to start in the FTP root directory or a specific user name–related directory.

To add a new FTP site, perform the following steps:

1. Open the IIS Manager console, right-click the Sites node, and then click Add FTP Site.
2. On the Site Information page, enter a name for the site and the directory path that is used to host the site's files.
3. On the Binding And SSL Settings page, shown in Figure 9-20, specify the following:
    1. Which IP addresses the FTP site listens on
    2. Whether the FTP site uses a virtual host name
    3. Whether the site starts automatically
    4. Whether SSL/TLS is allowed or required
    5. Which server certificate is used in conjunction with the FTP site

Technet24

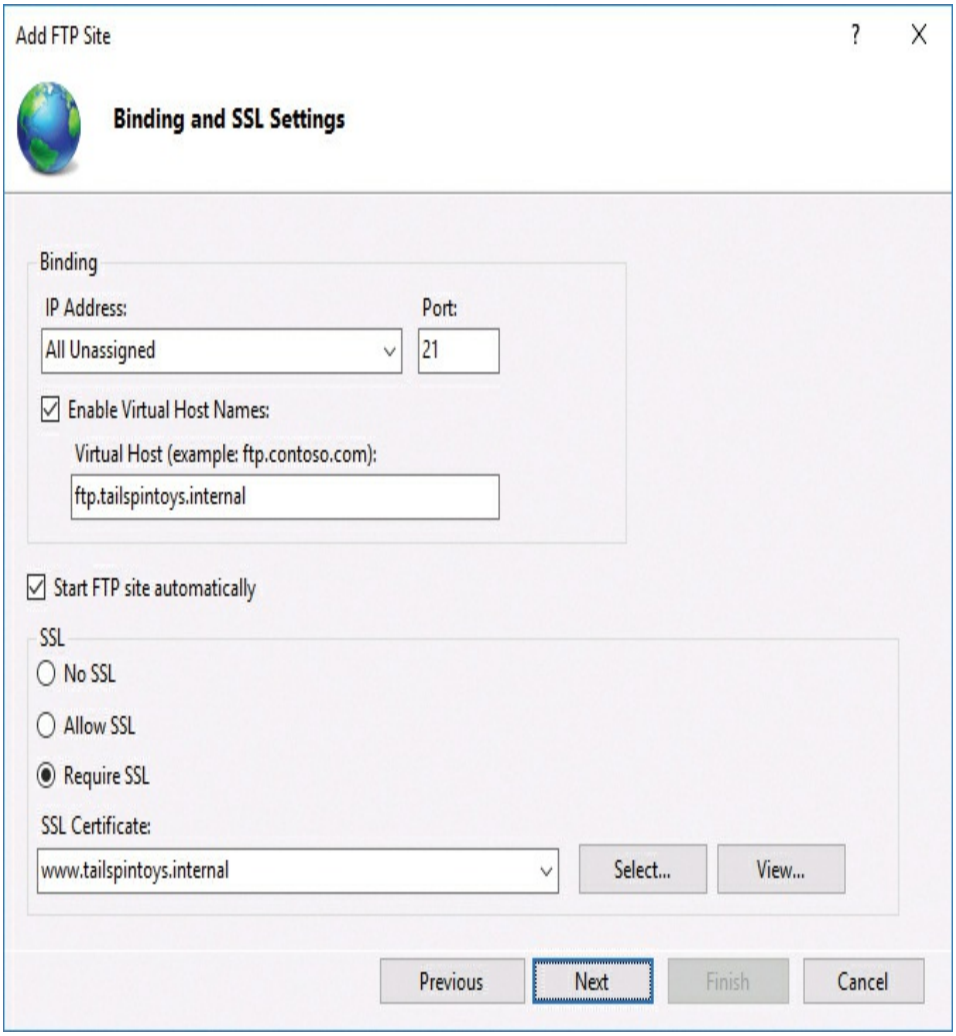


Figure 9-20 FTP site Binding and SSL Settings

4. On the Authentication And Authorization Information page, specify whether you'll allow Anonymous and/or Basic authentication. On the Authorization drop-down menu, specify whether you want to allow access for the following:

    1. All Users
    2. Anonymous Users
    3. Specified Roles Or User Groups
    4. Specified Users

If you choose to allow access only to specified roles or user groups, you need to specify which roles, users, or user groups you want to grant access.

MANAGING IIS USING POWERSHELL

You can use cmdlets from two separate PowerShell modules to manage IIS and websites running on Windows Server 2019. The *IIS Administration* PowerShell module contains a large number of cmdlets that you can use to manage IIS on Windows Server 2019. These cmdlets are described in Table 9-1. The *WebAdministration* PowerShell module contains a large number of cmdlets that you can use to manage web and FTP sites running on Windows Server 2019. These cmdlets are described in Table 9-2.

Table 9-1 IIS PowerShell cmdlets

Noun	Verbs	Function
IISAppPool	Get	View App Pool information
IISCentralCertProvider	Clear, Get, Set, Enable, Disable	Manage the IIS central certificate store
IISCentralC	Set	Manage user account credentials for the

Technet24

ertProvider Credential		IIS certificate store
IISCommit Delay	Stop, Start	Manage delay of commitment of changes
IISConfigAt tribute	Remove	Remove configuration attribute from IIS configuration section or configuration element attribute
IISConfigAt tributeValue	Set, Get	Manage configuration attribute value for an IIS configuration section or configuration element attribute
IISConfigC ollection	Get, Clear	Manage configuration element objects in an IIS configuration collection
IISConfigC ollectionEle ment	New, Remove, Get	Manage configuration element objects
IISConfigEl ement	Remove, Get	Manage configuration elements
IISConfigSe ction	Get	View configuration section object
IISConfigur ation	Export	Export IIS configuration
IISServerM anager	Reset, Get	Manage IIS Server Manager

IISSharedConfig	Disable, Enable, Get	Manage IIS shared config
IISSite	Stop, Start, New, Get, Remove	Manage IIS Site

Table 9-2 Web Administration PowerShell cmdlets

Noun	Verbs	Function
WebAppDomain	Get	View application domains in which the IIS worker process is running
WebApplication	New, Convert, Remove, Get	Manage web applications in IIS websites
WebAppPool	Start, New, Restart, Remove, Stop	Manage application pools
WebAppPoolState	Get	View application pool states
WebBinding	New, Set, Remove, Get	Configure web bindings
WebCentralCertProvider	Clear, Get, Enable, Set, Disable	Manage configuration settings for central certificate provider

Technet24

WebCentralCertProviderCredential	Set	Configure user account credentials for central certificate provider
WebCommitDelay	Start, Stop	Manage delay of commitment of changes
WebConfigFile	Get	View file system path of *web.config* file
WebConfiguration	Backup, Add, Select, Set, Restore, Clear, Get	Manage IIS configuration elements
WebConfigurationBackup	Get, Remove	Manage IIS configuration backups
WebConfigurationLocation	Remove, Get, Rename	Manage location of configuration settings
WebConfigurationLock	Add, Get, Remove	Manage lock status of IIS configuration
WebConfigurationProperty	Set, Add, Remove, Get	Manage configuration properties
WebFilePath	Get	View path to location of specific IIS module
WebFtpSite	New	Create a new FTP site
WebGlobalMo	Disable, New, Get,	Manage IIS global modules

dule	Remove, Enable, Set	
WebHandler	Get, New, Remove, Set	Manage IIS request handlers
WebItem	Restart, Start, Stop	Manage application pools or websites
WebItemState	Get	View state of application pool or website
WebManaged Module	Get, Set, New, Remove	Manage IIS managed modules
WebRequest	Get	View IIS requests currently being run
WebRequestTracing	Disable, Enable	Manage request tracing for a website
WebRequestTracingSetting	Clear	Clears request tracing configuration
WebRequestTracingSettings	Clear	Clears request tracing configuration (not a duplicate)
Website	Start, Get, Stop, Remove, New	Manage websites
WebsiteState	Get	View the state of a website

Technet24

WebURL	Get	View a website URL
WebVirtualDirectory	Remove, Get, New	Manage web virtual directories

Chapter 10

Containers

Container concepts

Isolation modes

Managing containers with Docker

Managing containers

Applying updates

Container networking

Linux containers on Windows

Container orchestration

A container is a portable isolated application execution environment. Containers allow developers to bundle an application and its dependencies in a single image. This image can easily be exported, imported, and deployed on different container hosts, from a developer's laptop computer, to Server Core on bare-metal hardware, and eventually to being hosted and run on Azure. Because an

Technet24

application's dependencies are bundled with the application within a container, IT operations can deploy a container as soon as it is handed off without worrying whether an appropriate prerequisite software package has been installed or if a necessary setting has been configured. Because a container provides an isolated environment, a failure that occurs within one container will only impact that container and will not impact other containerized applications running on the container host.

CONTAINER CONCEPTS

Containers can be conceptually challenging. To understand how containers work, you first must understand some of the terminology involving containers. While these concepts will be fleshed out more completely later in this chapter, you should understand the following terms at a high level:

- **Container image.** A container image is a template from which a container is generated. There are two general types of container images, which are usually created for a workload: a base OS image and a specific image. The difference between these containers on Windows Server is as follows:
    - **Container base OS image.** This is an image of the operating system upon which other container images are built. Four Windows OS container images are regularly published and updated by Microsoft. These images are available through an official public container registry.

- **Windows Server Core.** This is an image of the Windows Server operating system in the Server Core configuration. Containers don't run with a GUI, so Server with Desktop Experience is not available. This container image is suitable for traditional .NET Framework applications.
- **Nano Server.** This is a Windows Server image with all unnecessary elements stripped out. It is suitable for hosting .NET Core applications. The Nano Server image container image is what became of the Nano Server installation option that was available with the RTM version of Windows Server 2016. Because the image is stripped down to the essentials, it is far smaller than the Windows Server Core image and can be deployed and run more quickly.
- **Windows IoT Core.** This is an image for Internet of Things (IoT) applications.
- **Windows.** This is an image that provides the full Windows API set but doesn't include all the server roles and features that are available in the Server Core image. You should only use this option if the application you are trying to host has a dependency that is not included in the Windows Server Core container image.
- **Container image.** A container image stores changes made to a running container base OS image or another container image. For example, you can start a new container from a container base OS image, make modifications such as installing Java or a Windows feature, and then save those modifications as a new container image. The new container image only stores the changes you make, and therefore, it is much smaller than the parent container base OS image. You can then create an additional container image that stores modifications made to the container image that has an installed Java and Windows feature. Each container image only stores the changes made to the image from which it was run.

- **Sandbox.** The sandbox is the environment in which you can make modifications to an existing container image before you commit the changes to create a new container image. If you don't commit those changes to a new image, those changes will be lost when the container is removed.
- **Image dependency.** A new container image has the container image from which it was created as a dependency—for example, if you create a container image named WebServer-1.0 that has the IIS role installed from the Server Core base OS image, and the Server Core base OS image is a dependency for the WebServer-1.0 image. This dependency is very specific, and you can't use an updated version of the Server Core base OS image as a replacement for the version that you used when creating the WebServer-1.0 image. You can then export this container image that only records the changes made to another container host. You can start the image as long as the dependency container OS base image is present on that container host. You can have multiple images in a dependency chain.
- **Container host.** A container host is a computer that runs containers. A container host can be virtual, physical, or even a cloud-based Platform as a Service (PaaS) solution.
- **Container registries.** Container registries are central storehouses of container images. While it's possible to transfer containers or copy them to file shares using tools like FTP, common practice is to use a container registry as a repository for container images. Public container registries, such as the one that hosts Microsoft's base OS images, also store previous versions of the container base OS images. This allows you to retrieve earlier versions of the container base OS image that other images may depend upon. Docker Hub is the primary public container registry used by Microsoft to publish images, though the container images themselves are hosted in the Microsoft Container Registry at *mcr.microsoft.com*. Container registries can be public or private. When working with your own organization's container images, you have the option of creating and maintaining a private container registry.

Inside OUT

Data and Containers

Containers offer a non-persistent, isolated application execution environment. An application should run in the container, but it shouldn't store application data in a container. Application data should be stored in a location separate from the container, such as in a back-end database server, Azure storage blob, or other persistent alternative. Containers are good at serving in stateless roles, such as load balanced front-end web servers; this is especially true because you can start new containers quickly to handle additional loads as required.

ISOLATION MODES

Windows Server supports two container isolation modes: Process Isolation mode and Hyper-V Isolation mode. In previous versions of Microsoft's documentation, these isolation modes were occasionally called "Windows Server containers" and "Hyper-V containers." Windows 10 supports only Hyper-V Isolation mode. Windows Server 2016 and Windows Server 2019 support Process Isolation and Hyper-V Isolation modes.

Process Isolation mode

Process Isolation mode provides a container with process and namespace isolation. Containers running in this Isolation mode share a kernel with all other containers running on the container host. This is similar to the manner in which containers run on Linux container hosts. Process Isolation is the default mode used when running a container. If you want to ensure

that the container is being used, start the container using the *--isolation=process* option.

Hyper-V Isolation mode

A container running in Hyper-V Isolation mode runs in a highly optimized virtual machine that also provides an isolated application execution environment. Hyper-V Isolation mode containers don't share the kernel with the container host, nor do they share the kernel with other containers on the same container host. You can only use Hyper-V Isolation mode containers if you have enabled the Hyper-V role on the container host. If the container host is a Hyper-V virtual machine, you will need to enable Nested Virtualization. By default, a container uses the Process Isolation mode. You can start a container in Hyper-V Isolation mode by using the *--isolation=hyperv* option.

For example, to create a Hyper-V container from the *microsoft/windowsservercore* image, issue this command:

Click here to view code image

```
Docker run -it --isolation=hyperv
mcr.microsoft.com/windows/servercore cmd
```

More Info

Container isolation

You can learn more about container isolation at https://docs.microsoft.com/en-us/virtualization/windowscontainers/manage-containers/hyperv-container.

MANAGING CONTAINERS WITH DOCKER

Containers on Windows Server are managed using the Docker engine. The advantage is that the command syntax of Docker on Windows is almost identical to the command-line tools in Docker on Linux. While there is a community-maintained PowerShell module for managing containers on Windows Server available on GitHub, PowerShell is not the primary tool for Windows Server container management, and very few people use PowerShell to manage containers.

For Windows Server administrators unfamiliar with Docker syntax, the commands include extensive help support. Typing **Docker** at the command prompt will provide an overview of the high-level Docker functionality. You can learn more about specific commands by using the *--help* command parameter. For example, the following command will provide information about the *Docker Image* command:

```
docker image --help
```

More Info

Docker Engine on Windows

You can learn more about Docker Engine on Windows at https://docs.microsoft.com/en-us/virtualization/windowscontainers/manage-

docker/configure-docker-daemon.

Installing Docker

Docker is not included with the Windows Server installation media, and you don't install it as a role or feature. Instead, you need to install Docker from the PowerShell Gallery. Although this is unusual for an important role on a Windows Server operating system, it does have the advantage of ensuring that you have the latest version of Docker.

The simplest way to install Docker is to ensure that your Windows Server computer has Internet connectivity. You then run the following command from an elevated PowerShell prompt, as shown in Figure 10-1, to install the Docker Microsoft Provider:

Click here to view code image

```
Install-Module -Name DockerMsftProvider -Repository PSGallery -
Force
```

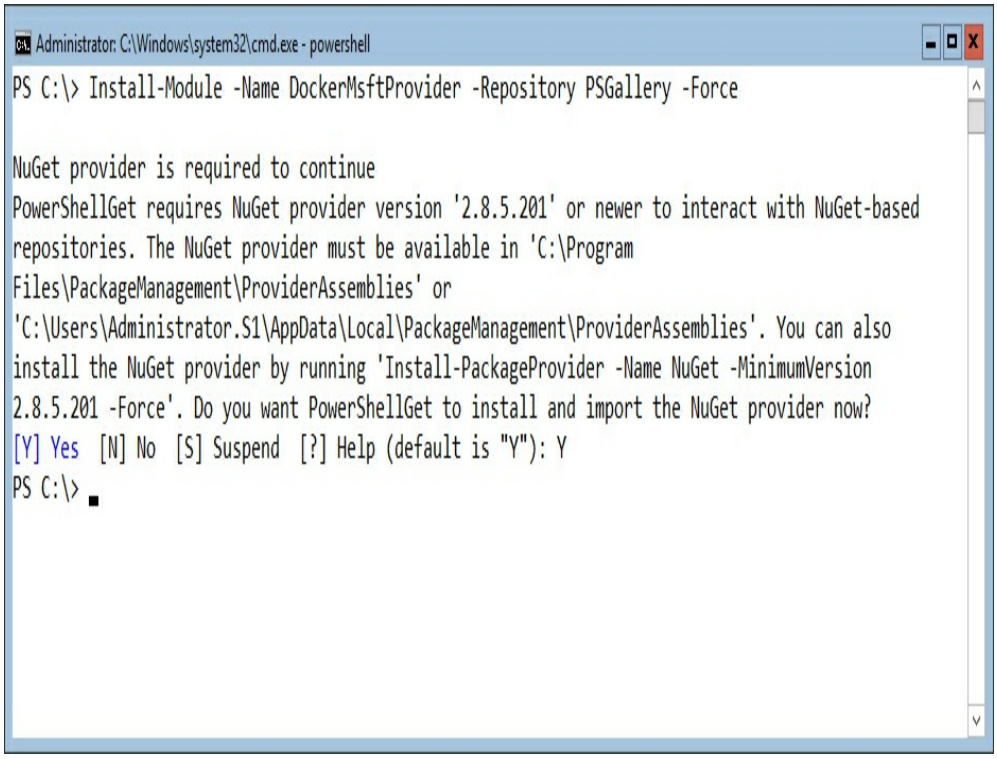


Figure 10-1 Install DockerMsftProvider

Next, install the most recent version of Docker. You do this by executing the following command, as shown in Figure 10-2; choose Yes when prompted to install software from a source not marked as Trusted.

Click here to view code image

```
Install-Package -Name docker -ProviderName DockerMsftProvider
```

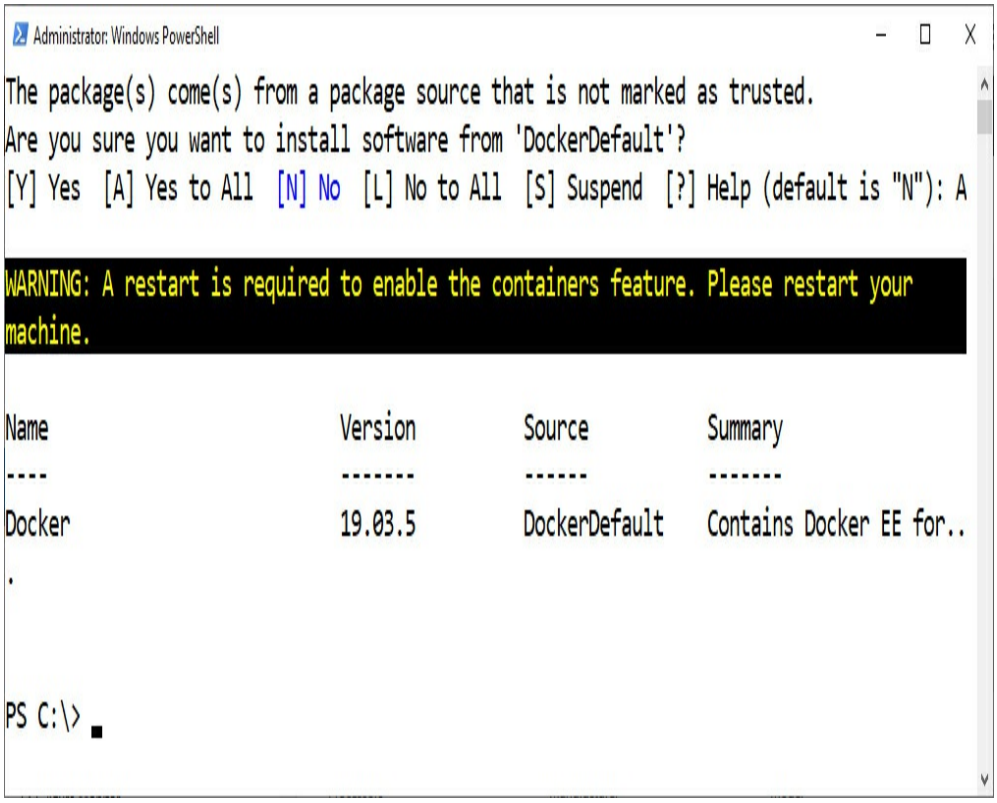


Figure 10-2 Install Docker

You'll then need to restart the computer that will function as the container host to complete the installation of Docker. You can update the version of Docker when a new one becomes available by rerunning the following command:

Click here to view code image

```
Install-Package -Name docker -ProviderName DockerMsftProvider
```

Demon.json

If you want to change the default Docker Engine

settings, such as whether to create a default NAT network, you need to create and configure the Docker Engine configuration file. This file doesn't exist by default. When it is present, the settings in the file override the Docker Engine's default settings.

You should create this file in the *C:\ProgramData\Docker\config* folder. Before editing the *Demon.json* file, you'll need to stop the Docker service using the following command:

```
Stop-Service docker
```

You only need to add settings that you want to change to the configuration file. For example, if you only want to configure the Docker Engine to accept incoming connections on port 2701, you add the following lines to *demon.json*:

```
{
    "hosts": ["tcp://0.0.0.0:2701"]
}
```

The Windows Docker Engine doesn't support all possible Docker configuration file options. The ones that you can configure are shown below:

Click here to view code image

```
{
    "authorization-plugins": [],
    "dns": [],
    "dns-opts": [],
    "dns-search": [],
```

```
    "exec-opts": [],
    "storage-driver": "",
    "storage-opts": [],
    "labels": [],
    "log-driver": "",
    "mtu": 0,
    "pidfile": "",

    "cluster-store": "",
    "cluster-advertise": "",
    "debug": true,
    "hosts": [],
    "log-level": "",
    "tlsverify": true,
    "tlscacert": "",
    "tlscert": "",
    "tlskey": "",
    "group": "",
    "default-ulimits": {},
    "bridge": "",
    "fixed-cidr": "",
    "raw-logs": false,
    "registry-mirrors": [],
    "insecure-registries": [],
    "disable-legacy-registry": false
}
```

These options allow you to do the following when starting the Docker Engine:

- **authorization-plugins.** Which authorization plug-ins the Docker Engine should load
- **dns.** Which DNS server the containers should use for name resolution
- **dns-opts.** Which DNS options to use
- **dns-search.** Which DNS search domains to use
- **exec-opts.** Which runtime execution options to use
- **storage-driver.** Specify the storage driver
- **storage-opts.** Specify the storage driver options

- **labels.** Docker Engine labels
- **log-driver.** Default driver for the container logs
- **mtu.** Container network MTU
- **pidfile.** Path to use for daemon PID file
- **group.** Specify the local security group that has permissions to run Docker commands
- **cluster-store.** Cluster store options
- **cluster-advertise.** Cluster address to advertise
- **debug.** Enable Debug mode
- **hosts.** Daemon sockets to connect to
- **log-level.** Logging detail
- **Tlsverify.** Use TLS and perform verification
- **tlscacert.** Specify which Certificate Authorities to trust
- **tlscert.** Location of the TLS certificate file
- **tlskey.** Location of the TLS key file
- **group.** UNIX socket group
- **default-ulimits.** Default ulimits for containers
- **bridge.** Attach containers to network bridge
- **fixed-cidr.** IPv4 subnet for static IP address
- **raw-logs.** Log format used
- **registry-mirrors.** Preferred registry mirror
- **insecure-registries.** Allow insecure registry communication
- **disable-legacy-registry.** Block contacting legacy registries

Once you have made the necessary modifications to the *daemon.json* file, you should start the Docker service by

running this PowerShell command:

```
Start-Service docker
```

Retrieving container OS image

You can retrieve the Server Core base OS container image by running the following command, as shown in Figure 10-3:

Click here to view code image

```
docker pull mcr.microsoft.com/windows/servercore:ltsc2019
```

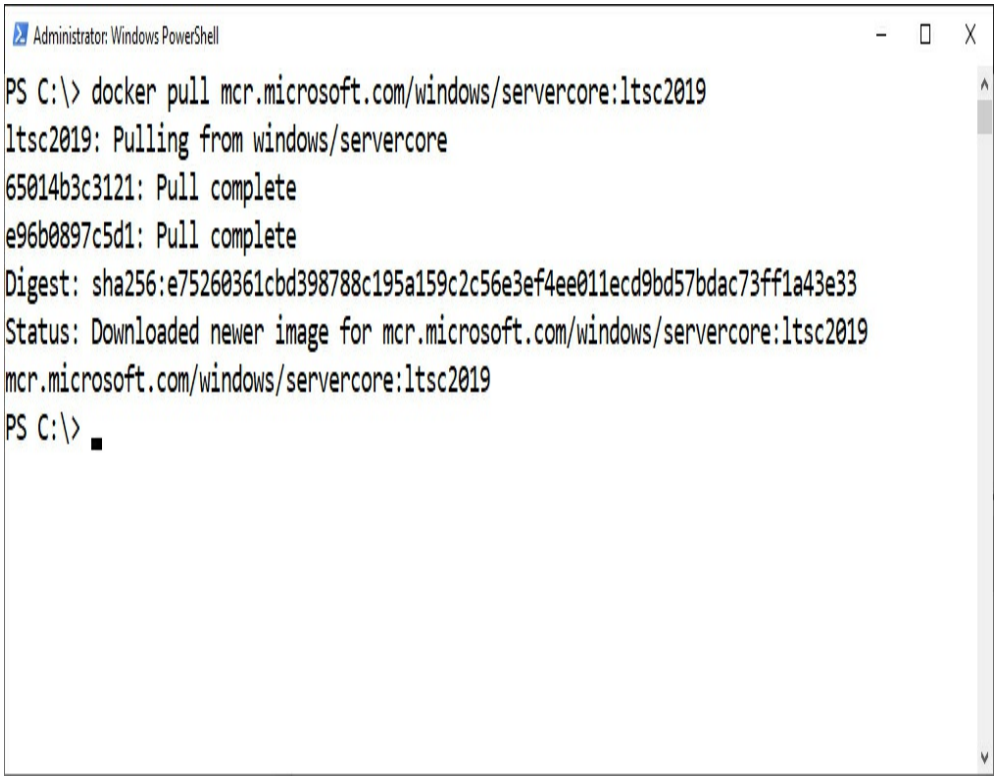


Figure 10-3 Pull OS image

You can retrieve the Nano Server base OS container

image by running this command:

Click here to view code image

```
docker pull mcr.microsoft.com/windows/nanoserver:1909
```

Unlike the Server Core image, where you will want to pull the Long Term Servicing Channel (LTSC) version, with Nano Server you'll need to determine the most recent supported version, keeping in mind that Nano Server images are only supported for 18 months. You can do that by looking at the image page on Docker Hub at *https://hub.docker.com/_/microsoft-windows-nanoserver*.

The latest version of the Server Core image is updated frequently. Because of dependency issues when working with older containers, you might need an earlier version of the base OS image. To retrieve all versions of the *WindowsServerCore* image from the public registry, run this command:

Click here to view code image

```
docker pull -a mcr.microsoftcom/windows/servercore
```

To retrieve all versions of the Nano Server image, as shown in Figure 10-4, run this command:

Click here to view code image

```
Docker pull -a mcr.microsoft.com/windows/nanoserver
```

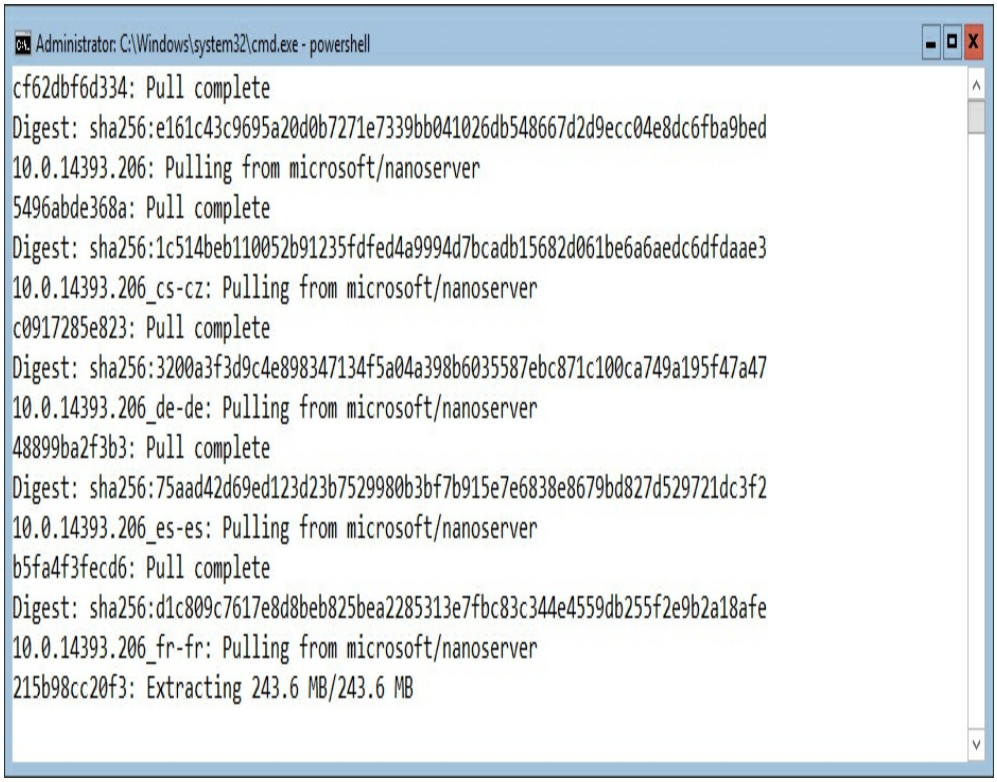


Figure 10-4 Pulling all versions of an OS image

Container registries and images

Container registries are repositories for the distribution of container images. The main container registry that will be of interest to Windows Server administrators is the Docker Hub public repository. Microsoft posts container images on the Docker Hub registry, including the base OS images, images that include evaluation editions of SQL Server, and technologies such as the latest builds of ASP.NET on containers and Azure CLI. You can view all Microsoft's published container images at *https://hub.docker.com/u/microsoft*.

From the Docker Hub registry, you can retrieve the following published Microsoft container images:

- **Nanoserver.** This is the base image for the Nano Server container operating system.

  Click here to view code image

  ```
  docker pull
  mcr.microsoft.com/windows/nanoserver:1909
  ```

- **Windows Server Core.** This is the base image for the Windows Server Core container operating system.

  Click here to view code image

  ```
  docker pull
  mcr.microsoft.com/windows/servercore:ltsc20
  19
  docker pull
  mcr.microsoft.com/windows/servercore:1909
  ```

- **Windows IoT Core.** This is the base image for Windows IoT containers.

  Click here to view code image

  ```
  docker pull
  mcr.microsoft.com/windows/iotcore:1809
  ```

- **Windows IIS.** This includes the Internet Information Services on Windows Server Core container operating system.

  Click here to view code image

  ```
  docker pull
  mcr.microsoft.com/windows/servercore/iis:wi
  ```

```
ndowsservercore-ltsc2019
docker pull
mcr.microsoft.com/windows/servercore/iis:wi
ndowsservercore-1909
```

- **Microsoft SQL Server.** This container image contains Microsoft SQL Server on Linux and will only run on Linux container hosts, rather than on Windows Server container hosts, unless you configure LCOW.

  Click here to view code image

  ```
  docker pull
  mcr.microsoft.com/mssql/server:2019-latest
  ```

- **Azure Cli.** This includes Azure CLI on a Linux container operating system.

  Click here to view code image

  ```
  docker pull mcr.microsoft.com/azure-cli
  ```

- **ASP.NET Core Runtime.** This includes ASP.NET Core on the Windows Server Core container operating system.

  Click here to view code image

  ```
  docker pull
  mcr.microsoft.com/dotnet/core/aspnet:3.1
  ```

- **Microsoft PowerShell.** This includes PowerShell on cross-platform container images.

  Click here to view code image

  ```
  docker pull mcr.microsoft.com/powershell
  ```

In cases in which multiple images exist, such as Windows Server Core and Nano Server, you can use the *-a* option with the Docker *pull* command to retrieve all images. This can be helpful if you don't know the image ID of the specific image that you wish to retrieve. For example, use this command when pulling the *Microsoft PowerShell* container image, which includes both Windows Server Core and Nano Server images: *docker pull -a mcr.microsoft.com/powershell*.

When you pull an image from a registry, the action will also download any dependency images that are required. For example, if you pull an image that was built on a specific version of the Windows Server Core base image and you don't have that image on the container host, that image will also be downloaded from a container registry.

You can view a list of images that are installed on a Windows Server container host by using the following command:

```
Docker image list
```

You can use Windows Admin Center to view a list of container images that are installed on a Windows Server container host, as shown in Figure 10-5. You can also use Windows Admin Center to delete images that are installed on a Windows Server container host.

Technet24

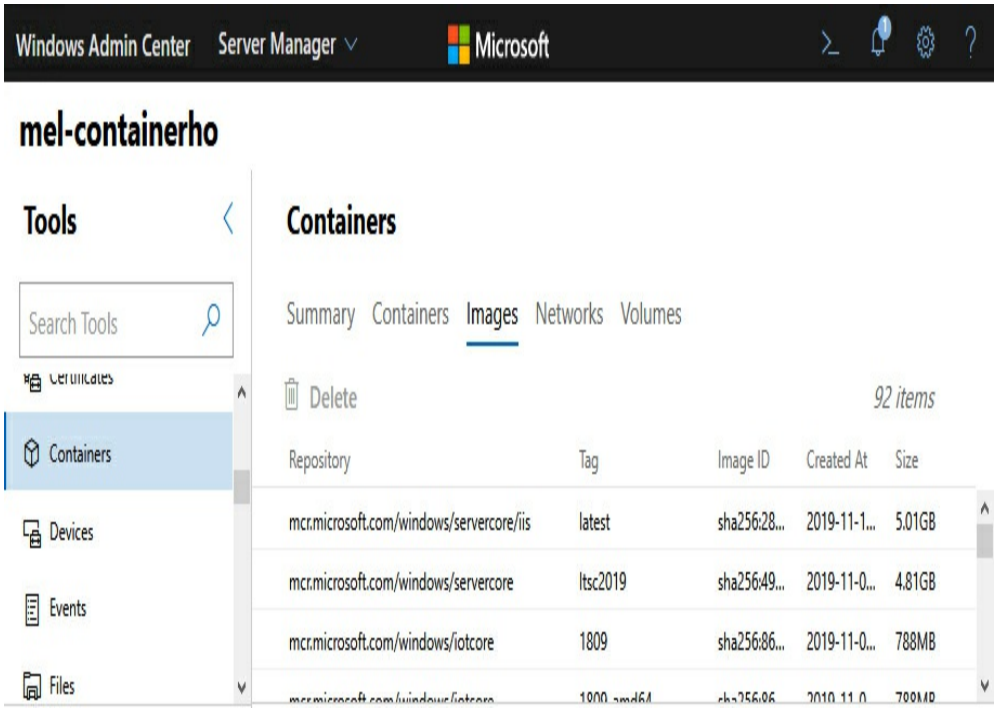


Figure 10-5 List of container images

More Info

Official Microsoft Container Images

You can see which container images Microsoft has made available at https://hub.docker.com/u/microsoft/.

MANAGING CONTAINERS

You use Docker to perform all container management tasks on computers running Windows Server. At present, the container management functionality available in Windows Admin Center is limited, but it's likely that over time, most tasks that you can perform from the Docker prompt will be available in WAC.

Starting a container

You create a new container by specifying the container image from which you wish to create the container. You can start a container and run an interactive session either by specifying *cmd.exe* or *PowerShell.exe* by using the *-it* option with *docker run*. Interactive sessions allow you to directly interact with the container through the command line from the moment the container starts. Detached mode starts a container, but it doesn't start an interactive session with that container.

For example, to start a container from the *Microsoft/windowsservercore* image and to enter an interactive PowerShell session within that container once it is started, use this command:

Click here to view code image

```
Docker run -it mcr.microsoft.com/windows/servercore:ltsc2019
powershell.exe
```

Also important is that, by default, containers use network address translation. This means that if you are running an application or service on the container that you want to expose to the network, you'll need to configure port mapping between the container host's network interface and the container. For example, you would run the following command if you had downloaded the *Microsoft/iis* container image and you wanted to start a container in Detached mode and map

port 8080 on the container host to port 80 on the container (see Figure 10-6).

Click here to view code image

```
Docker run -d -p 8080:80 mcr.microsoft.com/windows/servercore/
iis:windowsservercore-ltsc2019
```

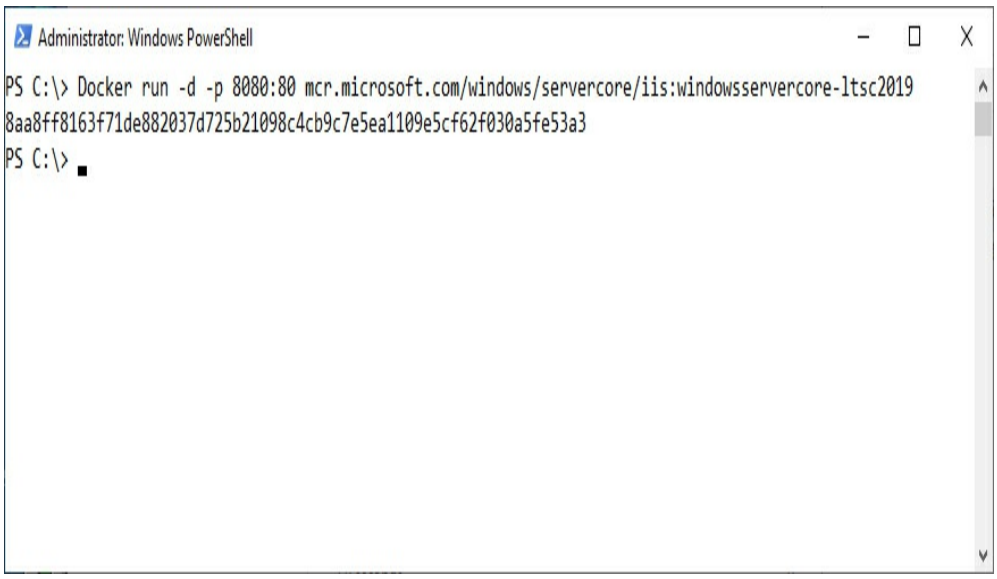


Figure 10-6 Port mapping

This is only the very basic sort of information you'd need to get started with a container, and you will learn more about container networking later in this chapter.

You can verify which containers are running by using the *docker ps* command, as shown in Figure 10-7, or by using Windows Admin Center. The problem with the simple *docker ps* command option is that this will only show you the running containers and won't show you any that are in a stopped state. You can see which

containers are on a container host, including containers that aren't currently running, by using the *docker ps -a* command.

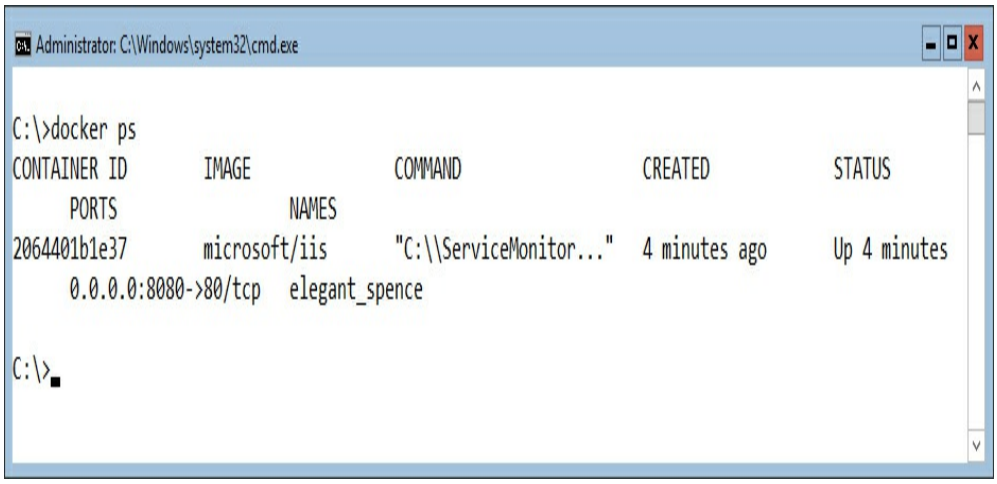


Figure 10-7 Listing containers

You can also view the containers that are on a Windows Server container host using Windows Admin Center, if you have loaded the Containers extension, as shown in Figure 10-8.

Technet24

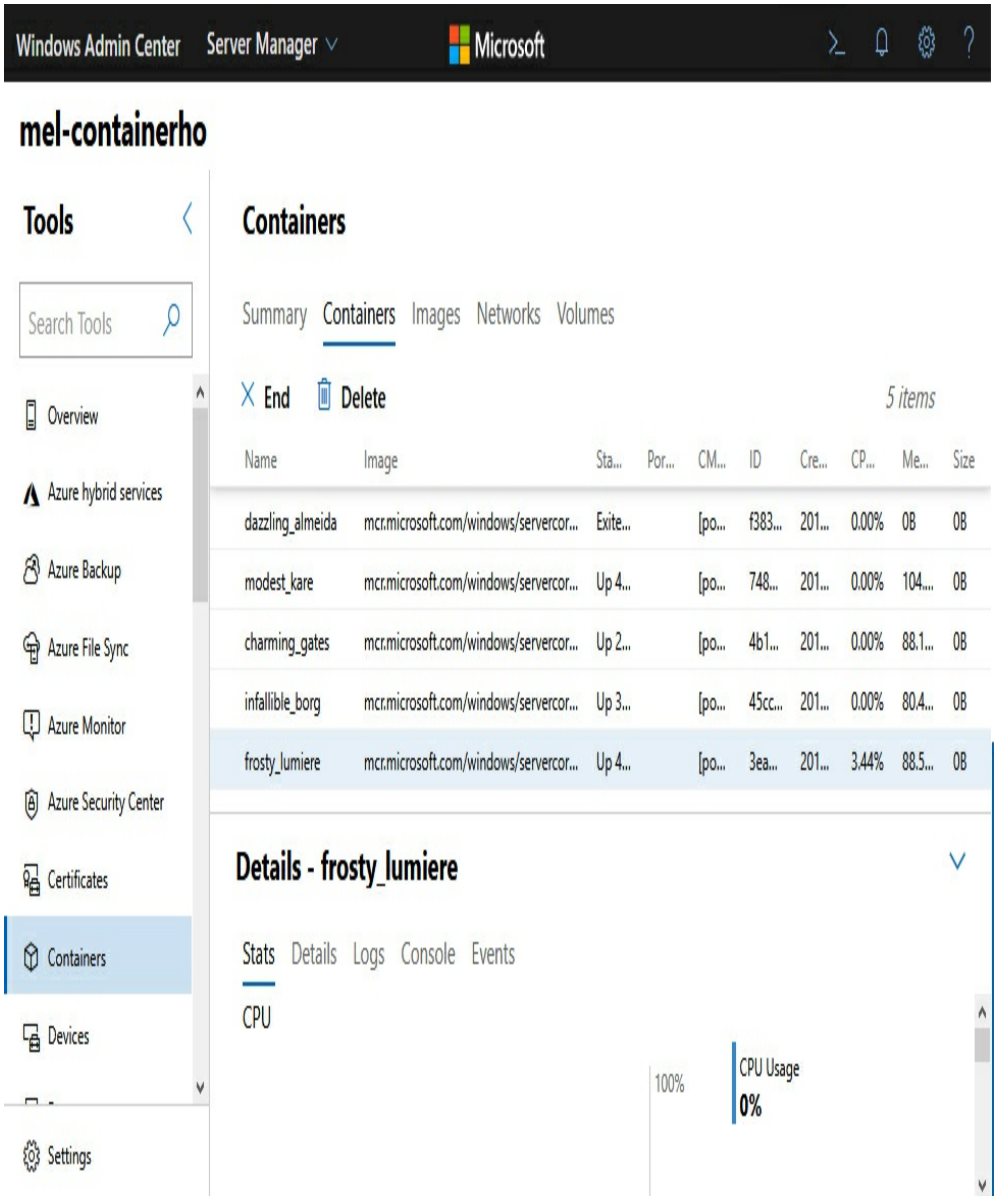


Figure 10-8 Windows Admin Center container list

One thing that you'll notice about containers is that they appear to be assigned random names, such as *sarcastic_hedgehog*, *dyspeptic_hamster*, and *sententious_muppet*. Docker assigns random names,

rather than asking you for one, because containers are a more ephemeral type of application host than a VM; because they are likely to only have a short lifespan, it isn't worth assigning any name that you'd need to remember. The reason for the structure of the random names is that they are easy to remember in the short term, which makes containers that you must interact with on a short-term basis easier to address than when using hexadecimal container IDs. Figure 10-9 shows the output of the *docker ps -a* command and shows the names of containers including: *charming_gates*, *infallible_borg*, *frosty_lumiere*, *modest_kare*, and *dazzling_alemida*.

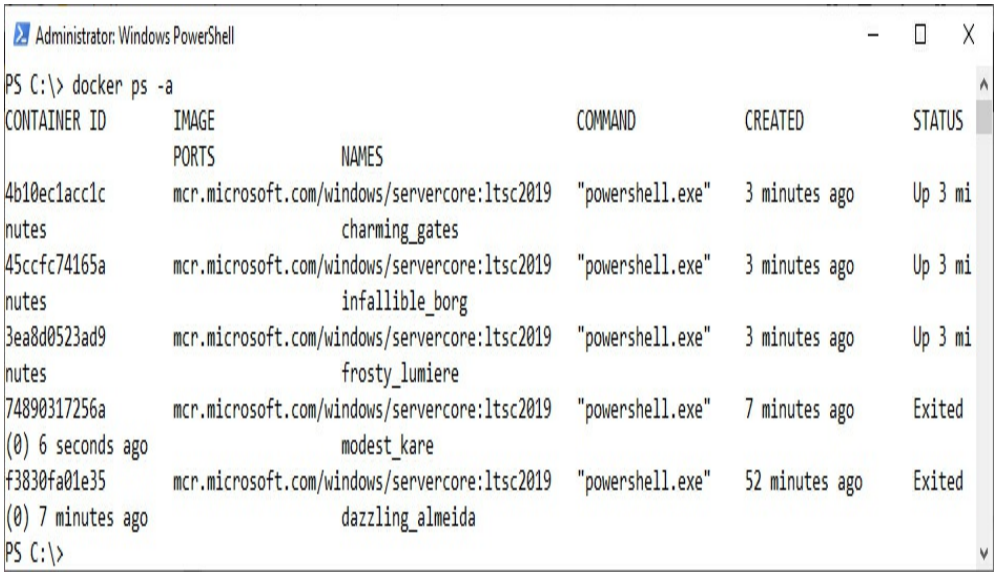


Figure 10-9 List all images

Earlier in the chapter, you learned that it was possible to start a container in Detached mode. If you want to start

Technet24

an interactive session on a container that you started in Detached mode, use the *docker exec -i <containername> powershell.exe* command. Figure 10-10 shows entering an interactive session on a container started in Detached mode. The *hostname* command displays the container's ID.

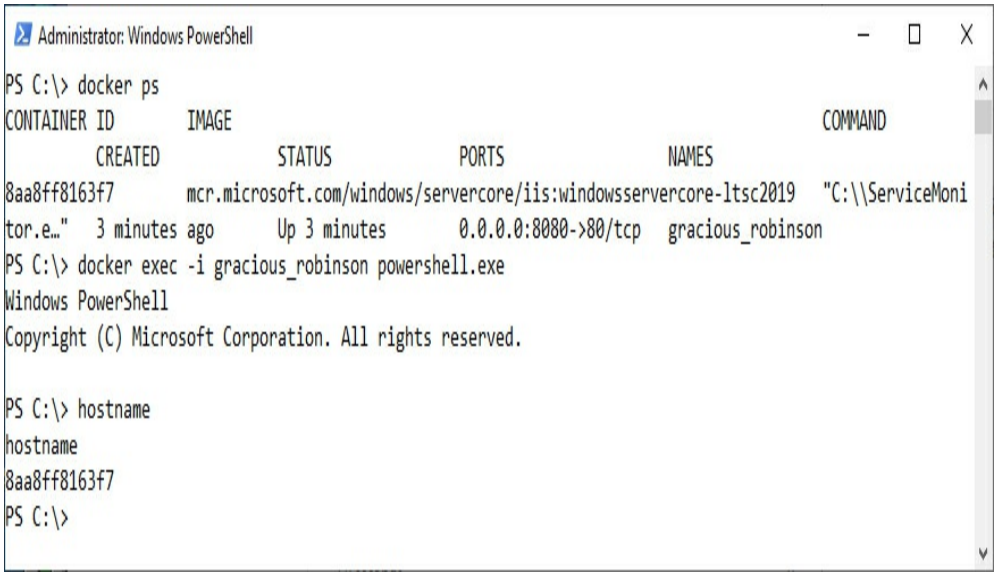


Figure 10-10 Starting an interactive session

In some cases, it will be necessary to start a stopped container. You can start a stopped container using the *Docker start <containername>* command.

Modifying a running container

Once you have a container running, you can enter the container and make the modifications that you want to make to ensure that the application the container hosts will run correctly. This might involve creating

directories, using the *dism.exe* command to add roles, or using the *wget* PowerShell alias to download and install binaries such as Java. For example, the following code, when run from inside a container, downloads and installs an older version of Java into a container based on the Server Core base OS container:

Click here to view code image

```
wget -Uri "http://javadl.sun.com/webapps/download/AutoDL?
BundleId=107944" -outfile javaInstall.exe -UseBasicParsing
REG ADD HKLM\Software\Policies\Microsoft\Windows\Installer /v
DisableRollback /t REG_DWORD /d 1 | Out-Null ./javaInstall.exe /s
INSTALLDIR=C:\Java REBOOT=Disable | Out-Null
```

Once you are finished modifying the container, you can type **Exit** to exit the container. A container must be in a shut-down state before you can capture it as a container image. You use the *docker stop <containername>* cmdlet to shut down a container.

Inside OUT

Containers and Chocolatey

Chocolatey offers another method of installing software on containers. You can spin up a new container, install Chocolatey, and then create a new image from that container. The process of installing additional software will be simpler each time you spin up a new container from that new image that includes Chocolatey.

Creating a new image from a container

Once the container is in the desired state and shut down, you can capture the container to a new container image. You do this using the *docker commit*

<container_name> *<new_image_name>* command. Figure 10-11 shows committing a modified container to a new container image named *new_iis_java* and verifying the properties of the image using the *docker images* command.

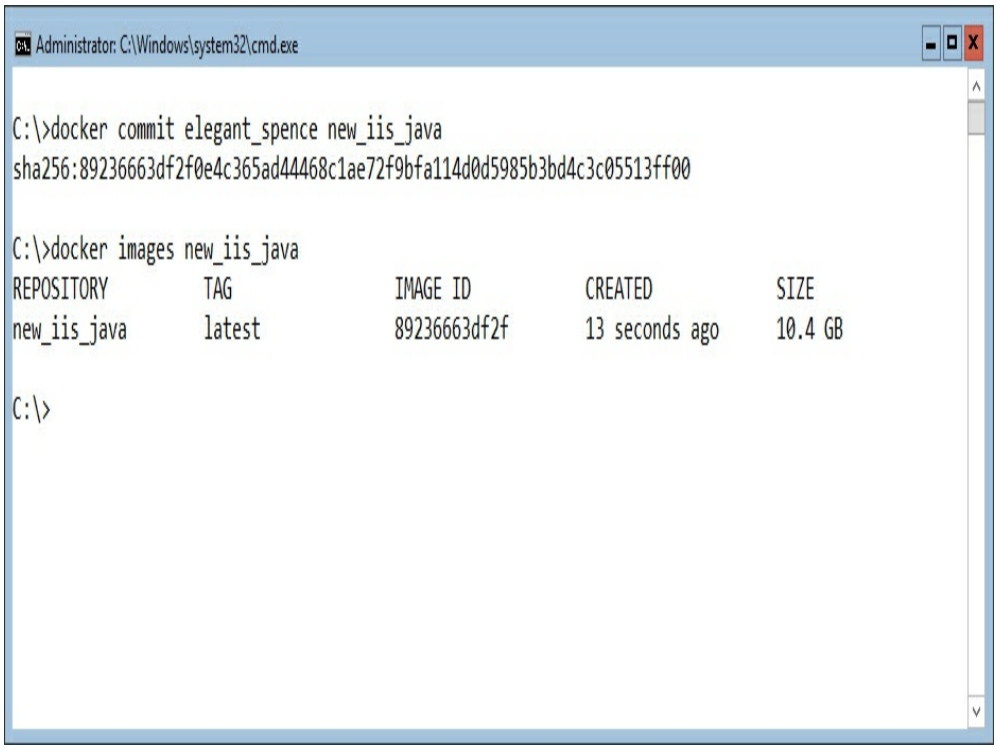


Figure 10-11 Docker commit

Once you have committed a container to a container image, you can remove the container using the *docker rm* command. For example, to remove the *elegant_spence* container, issue this command:

```
Docker rm elegant_spence
```

Using Dockerfiles

Dockerfiles are text files that allow you to automate the process of creating new container images. You use a Dockerfile with the *docker build* command to automate container creation, which is very useful when you need to create new container images from regularly updated base OS container images.

Dockerfiles have the elements shown in Table 10-1.

Table 10-1 Dockerfile Element

Instruction	Description
FROM	Specifies the container image used in creating the new image creation. For example: `FROM mcr.microsoftcom/windows/servercore:ltsc2019`

RUN	Specifies commands to be run and captures them into the new container image.
	For example:
	`RUN wget -Uri "http://javadl.sun.com/webapps/download/AutoDL?BundleId=107944" -outfile javaInstall.exe -UseBasicParsing`
	`RUN REG ADD HKLM\Software\Policies\Microsoft\Windows\Installe`

Technet24

	`r /v DisableRollback /t REG_DWORD /d 1 | Out-Null`
	`RUN ./javaInstall.exe /s INSTALLDIR=C:\Java`
	`REBOOT=Disable | Out-Null`

COPY	Copies files and directories from the container host filesystem to the filesystem of the container. For windows containers, the destination format must use forward slashes. For example: `COPY example1.txt c:/temp/`
ADD	Can be used to add files from a remote source, such as a URL. For example: *ADD https://www.python.org/ftp/python/3.5.1/python-3.5.1.exe*
WORKDIR	Specifies a working directory for the *RUN* and *CMD* instructions.
CMD	A command to be run when deploying an instance of the container image.

For example, the following Dockerfile will create a new container from the *mcr.microsoft.com/windows/servercore:ltsc2019* image, create a directory named *ExampleDirectory*, and then install the IIS Webserver feature.

Click here to view code image

```
FROM mcr.microsoftcom/windows/servercore:ltsc2019
```

```
RUN mkdir ExampleDirectory
RUN dism.exe /online /enable-feature /all /featurename:iis-
webserver /NoRestart
```

To create a container image named *example_image*, change into the directory that hosts the *Dockerfile* (no extension) file and run the following command:

```
Docker build -t example_image
```

More Info

Dockerfile with Windows

You can learn more about using Dockerfile with Windows at https://docs.microsoft.com/en-us/virtualization/windowscontainers/manage-docker/manage-windows-dockerfile.

Managing container images

To save a Docker image for transfer to another computer, use the *docker save* command. When you save a Docker image, you save it in *.tar* format. Figure 10-12 shows exporting the *new_iis_java* image to the *c:\archive\new_iis_java.tar* file.

Technet24

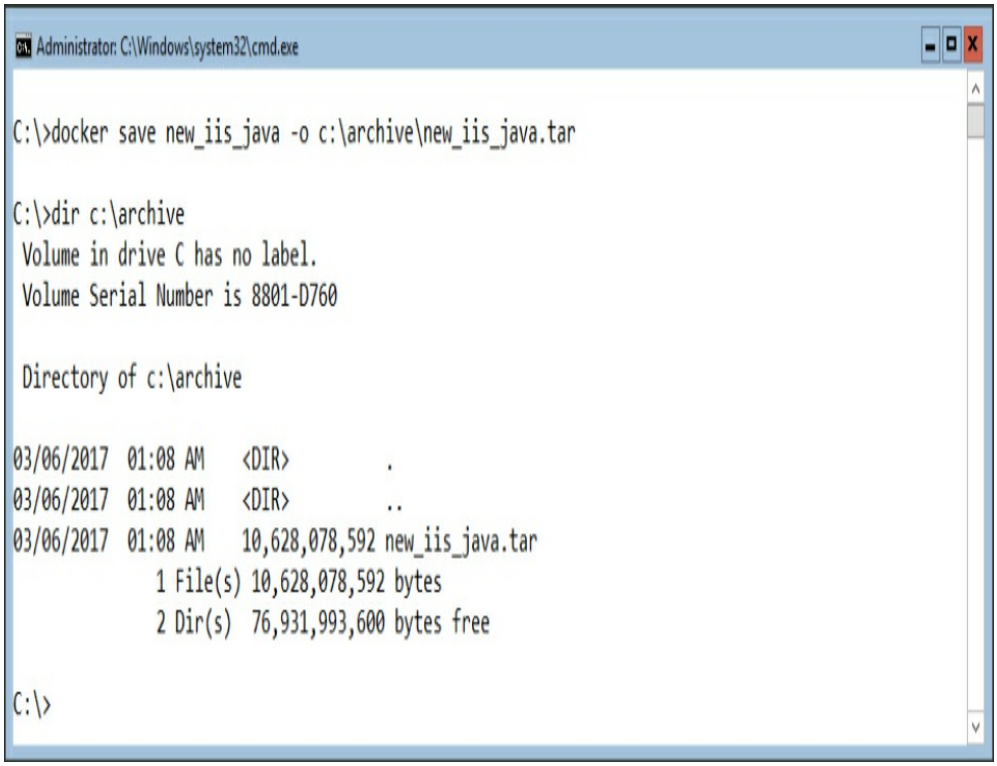


Figure 10-12 *Docker save* command

You can load a Docker image from a saved image using the *docker load* command. For example, Figure 10-13 shows loading a container image from the *c:\archive\new_iis_java.tar* file created earlier.

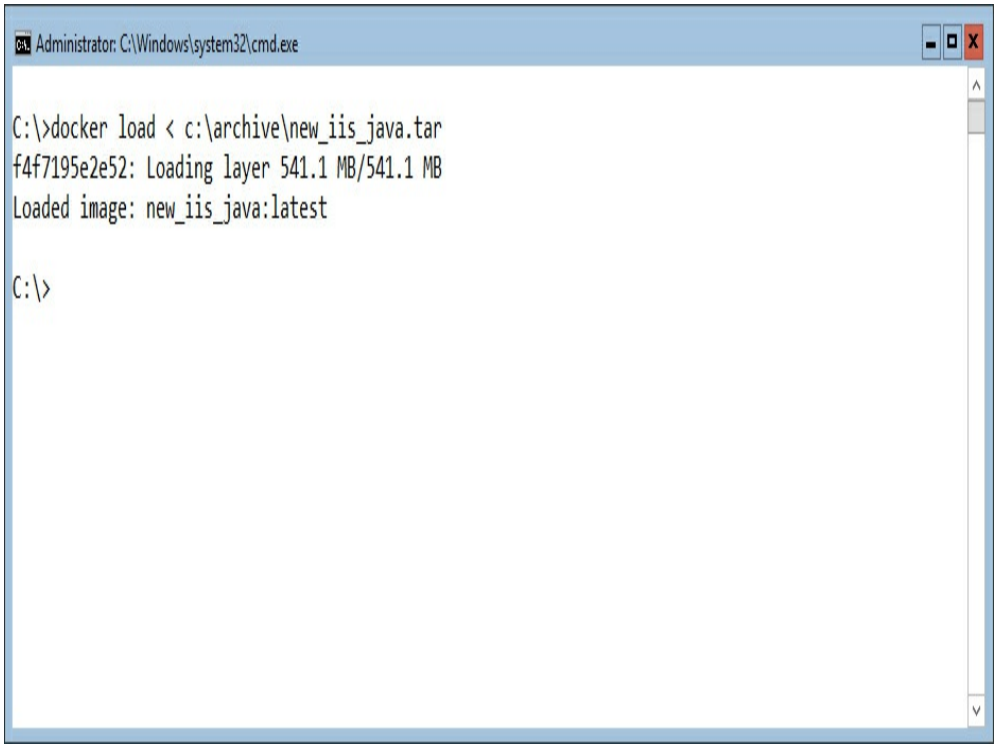


Figure 10-13 Load archived image

When you have multiple container images that have the same name, you can remove an image by using the image ID. You can determine the image ID by running the *docker images* command, as shown in Figure 10-14. You can also use Windows Admin Center to view this information. You can't remove a container image until the last container created from that image has been deleted either directly or indirectly.

Technet24

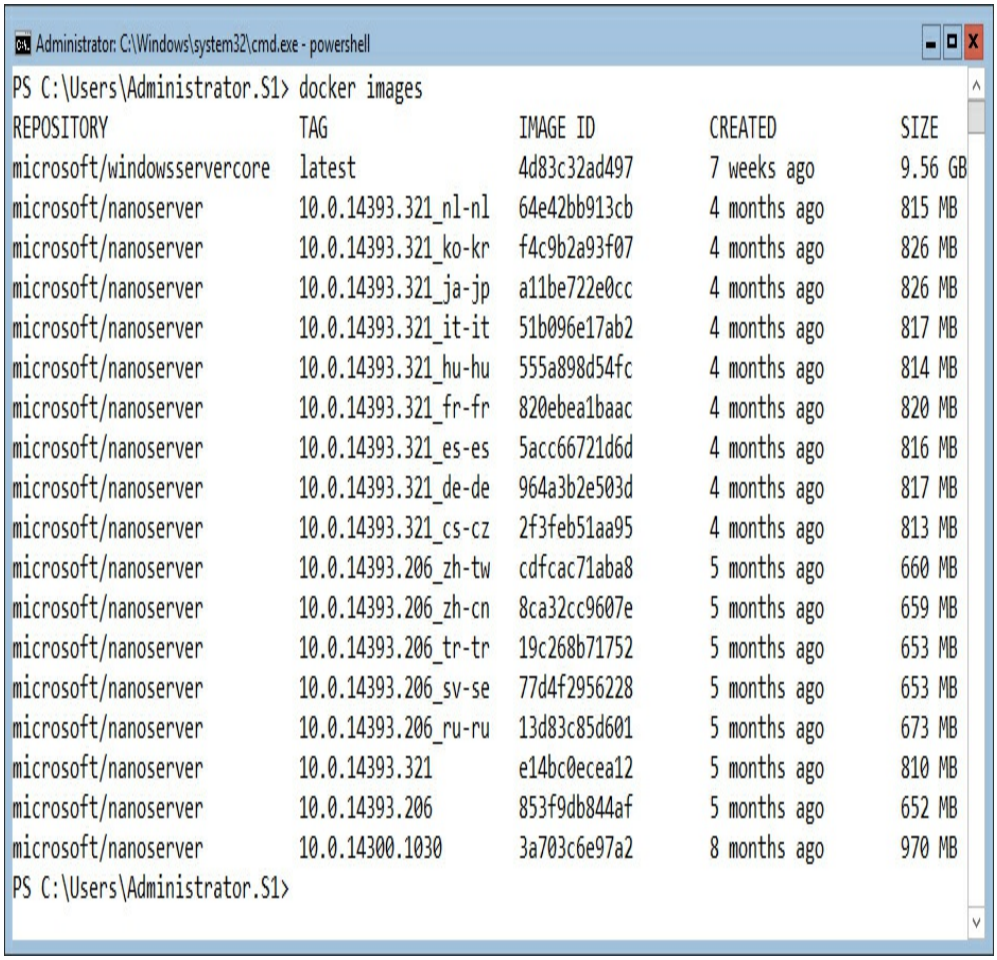


Figure 10-14 Image list

You then remove the image by using the *docker rmi* command with the image ID. For example, Figure 10-15 shows the removal of the image with the *ID a896e5590871* using this command:

```
docker rmi a896e5590871
```

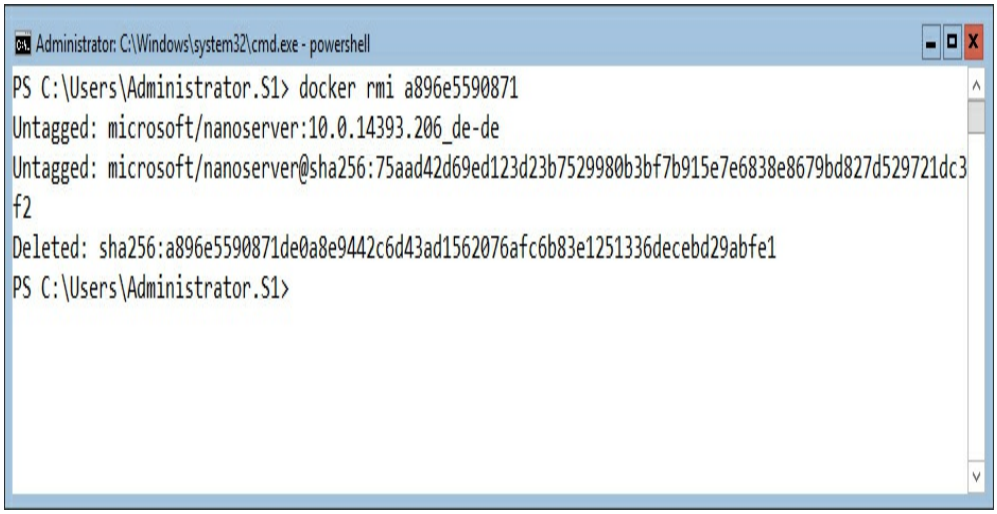


Figure 10-15 Remove image

You can also remove a container image by selecting the image in the list of containers in Windows Admin Center and clicking Delete.

Service accounts for Windows containers

Although containers based on the Server Core and Nano Server operating systems have most of the same characteristics as a virtual machine or a bare-metal deployment of the Server Core or Nano Server versions of Windows Server, one thing that you can't do with containers is to join them to a domain. This is because containers are supposed to be temporary, rather than permanent, and domain-joining them would clog up Active Directory with unnecessary computer accounts.

While containers can't be domain-joined, it is possible to use a group-managed service account (gMSA) to

Technet24

provide one or more containers with a domain identity similar to that used by a device that is realm-joined. Performing this task requires downloading the Windows Server Container Tools and ensuring that the container host is a member of the domain that hosts the gMSA. When you perform this procedure, the container's LocalSystem and Network Service accounts use the gMSA. This gives the container the identity represented by the gMSA.

To configure gMSA association with a container, perform the following steps:

1. Ensure that the Windows Server container host is domain-joined.
2. Add the container host to a specially created domain security group. This domain security group can have any name.
3. Create a gMSA and grant gMSA permission to the specially created domain security group of which the container host is a member.
4. Install the gMSA on the container host.
5. Use the *New-CredentialSpec* cmdlet on the container host to generate the gMSA credentials in a file in JSON format. This cmdlet is located in the *CredentialSpec* PowerShell module, which is available as a part of the Windows Server Container tools. For example, if you created a gMSA named *MelbourneAlpha*, you would run the following command:

   Click here to view code image

   ```
   New-CredentialSpace -Name MelbourneAlpha -
   AccountName MelbourneAlpha
   ```

6. You can verify that the credentials have been saved in JSON format by running the *Get-CredentialSpec* cmdlet. By default, credentials are stored in the *C:\ProgramData\Docker\CredentialSpecs* folder.

7. Start the container using the *--security-opt "credentialspec="* option and specify the JSON file containing the credentials associated with the gMSA. For example, run the following command if the credentials are stored in the file *twt_webapp01.json*:

   Click here to view code image

```
docker run --security-opt
"credentialspec=file://twt_webapp01.json" --
hostname
webapp01 -it
mcr.microsoft.com/windows/servercore:ltsc2019
powershell
```

Once you've configured the container to indirectly use the gMSA for its Local System and Network Service accounts, you can provide the container with permissions to access domain resources by providing access to the gMSA. For example, if you wanted to provide the container with access to the contents of a file share hosted on a domain member, you can configure permissions so that the gMSA has access to the file share.

More Info

Active Directory Container Service Accounts

You can learn more about Active Directory Container Service Accounts at https://docs.microsoft.com/en-us/virtualization/windowscontainers/manage-containers/manage-serviceaccounts.

APPLYING UPDATES

One of the concepts that many IT operations personnel find challenging is that you don't update a container that is deployed in production. Instead, you create a fresh container from the original container image, update that container, and then save that updated container as a new container image. You then remove the container in production and deploy a new container to production that is based on the newly updated image.

For example, you have a container that hosts a web app deployed from a container image named *WebApp1* that is deployed in production. The developers in your organization release an update to *WebApp1* that involves changing some existing settings. Rather than modifying the container in production, you start another container from the *WebApp1* image, modify the settings, and then create a new container image named *WebApp2*. You then deploy a new container into production from the *WebApp2* container image and remove the original un-updated container.

While you can manually update your container base OS images by applying software updates, Microsoft releases updated versions of the container base images each time a new software update is released. Once a new container OS base image is released, you or your organization's developers should update existing images that are dependent on the container OS base image. Regularly updated container base OS images provide an excellent

reason for eventually moving toward using Dockerfiles to automate the process of building containers. If you have a Dockerfile configured for each container image used in your organization, updating your Container base OS images when a new Container base OS image is released is a quick, painless, and automated process.

> **Inside OUT**
>
> **Keep old OS images**
>
> You'll need to keep your old container OS base images around until all the container images that are dependent on that old container OS base image have been replaced by newly updated container images.

CONTAINER NETWORKING

Each container has a virtual network adapter. This virtual network adapter connects to a virtual switch, through which inbound and outbound traffic is sent. Networking modes determine how network adapters function in terms of IP addressing—meaning whether they use NAT or are connected to the same network as the container host.

Windows Containers support the following networking modes:

- **NAT.** Each container is assigned an IP address from the private 172.16.0.0 /12 address range. When using NAT, you can configure port forwarding from the container host to the container endpoint. If you create a container without specifying a network, the container will use the default NAT network. The Docker service

creates its own default NAT network. When the container host reboots, the NAT network will not be created until the Docker service has restarted. Any container that was attached to the existing NAT network and that is configured to persist after reboot (for example, because it uses the *–restart always* option) will reattach to the NAT network that is newly created when the Docker service restarts.

- **Transparent.** Each container endpoint connects to the physical network. The containers can have IP addresses assigned statically or through DHCP.
- **Overlay.** Use this mode when you have configured the Docker Engine to run in Swarm mode. Overlay mode allows container endpoints to be connected across multiple container hosts.
- **L2 Bridge.** Container endpoints are on the same IP subnet used by the container host. IP addresses must be assigned statically. All containers on the container host have the same MAC address.
- **L2 Tunnel.** This mode is only used when you deploy containers in Azure.

You can list available networks using the following command, as shown in Figure 10-16:

```
docker network ls
```

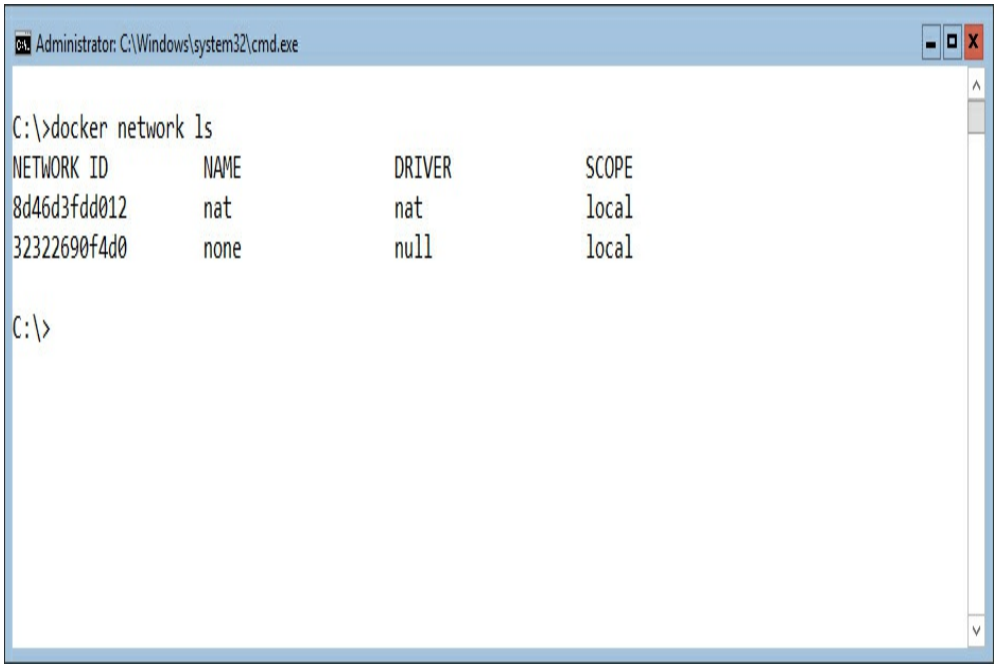


Figure 10-16 Network list

You can also view a list of container networks in Windows Admin Center, as shown in Figure 10-17.

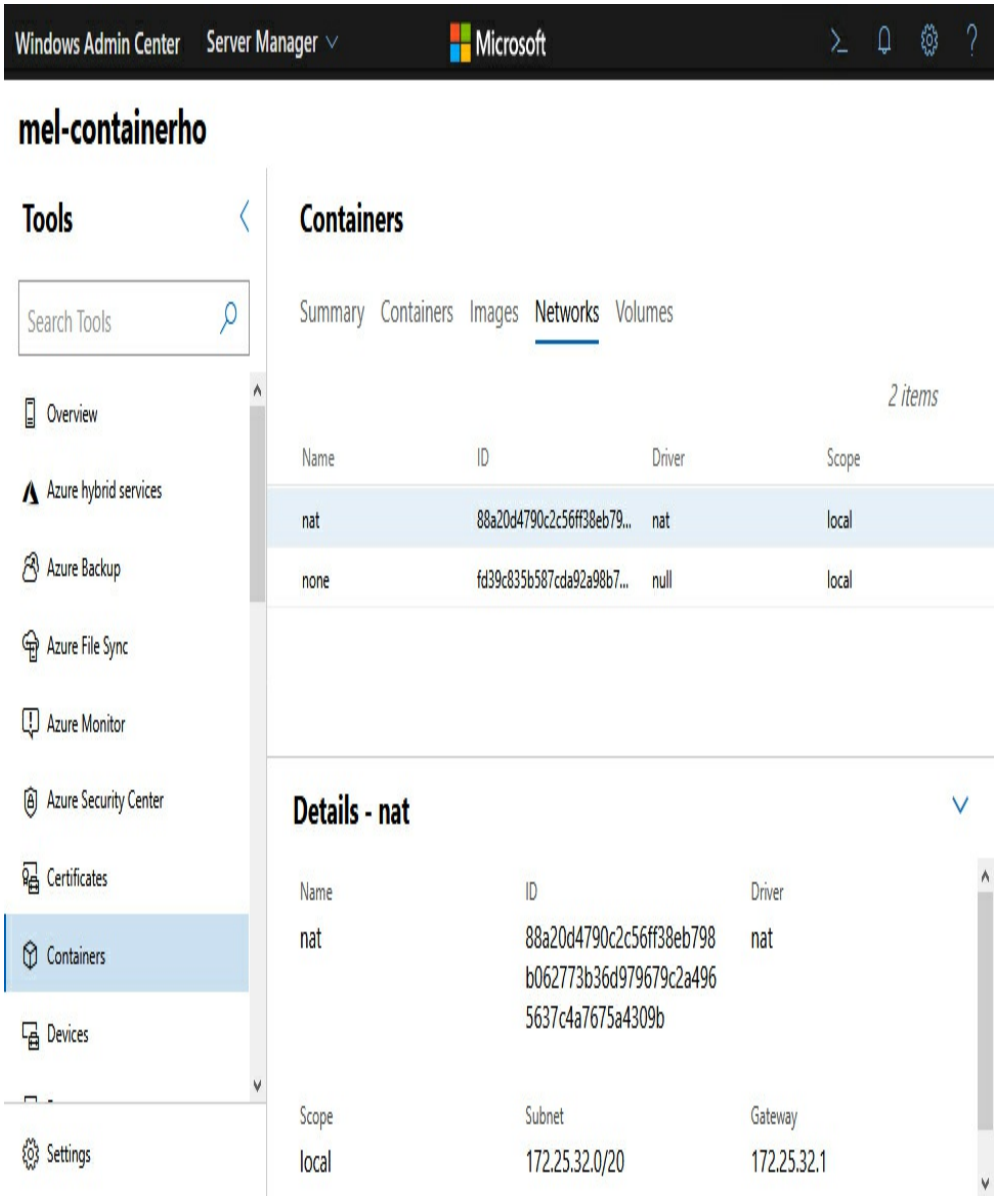


Figure 10-17 Windows Admin Center network list

To view which containers are connected to a specific network, run this command:

Click here to view code image

```
Docker network inspect <network name>
```

You can create multiple container networks on a container host, but you need to keep in mind the following limitations:

- If you are creating multiple networks of the transparent or L2 bridge type, you need to ensure that each network has a separate network adapter.
- If you create multiple NAT networks on a container host, additional NAT network prefixes will be partitioned from the container host's NAT network's address space. (By default, this is 172.16.0.0/12.) For example, these will be 172.16.1.0/24, 172.16.2.0/24, and so on.

NAT

Network Address Translation (NAT) allows each container to be assigned an address in a private address space, while connecting to the container host's network uses the container host's IP address. The default NAT address range for containers is 172.16.0.0 /16. If the container host's IP address is in the 172.16.0.0 /12 range, you will need to alter the NAT IP prefix. You can do this by performing the following steps:

1. Stop the Docker service.
2. Remove any existing NAT networks using the following command:

   Click here to view code image

   ```
   Get-ContainerNetwork | Remove-ContainerNetwork
   ```

3. Perform one of the following actions:

   1. Edit the *daemon.json* file and configure the *"fixed-cidr":"< IP Prefix > /Mask"* option to the desired network address prefix.

2. Edit the *daemon.json* file and set the *"bridge": "none"* option, and then use the *docker network create -d* command to create a network. For example, use the following command to create a network that uses the 192.168.15.0/24 range, the default gateway of 192.168.15.1, and is named *CustomNat*:

   Click here to view code image

   ```
   Docker network create -d nat --
   subnet=192.168.15.0/24 --
   gateway=192.168.15.1
   CustomNat
   ```

4. Start the Docker service.

You can also allow connections to custom NAT networks when a container is run by allowing the use of the --*Network* parameter and specifying the custom NAT network name. To do this, you need to have the *"bridge: none"* option specified in the *daemon.json* file. Use the following command to run a container and join it to the *CustomNat* network created earlier:

Click here to view code image

```
Docker run -it --network=CustomNat <ContainerImage> <cmd>
```

Port mappings allow ports on the container host to be mapped to ports on the container. For example, use the following command to create the container to map port 8080 on the container host to port 80 on a new container created from the *windowsservercore* image and to run *powershell.exe* interactively on the container:

Click here to view code image

```
Docker run -it -p 8080:80 microsoft/windowsservercore
powershell.exe
```

Port mappings must be specified when the container is created or when it is in a stopped state. You can specify them using the *-p* parameter or the *expose* command in a Dockerfile when using the *-p* parameter. A random port will be assigned if you do not specify a port on the container host, but you do specify one on the container itself. For example, run this command:

Click here to view code image

```
Docker run -itd -p 80
mcr.microsoft.com/windows/servercore/iis:windowsservercore-
ltsc2019 powershell.exe
```

A random port on the container host can be mapped through to port 80 on the container. You can determine which port is randomly assigned using the *docker ps* command. Figure 10-18 shows a container named *youthful_hopper* that has had the port 54742 on the container host assigned to port 80 on the container.

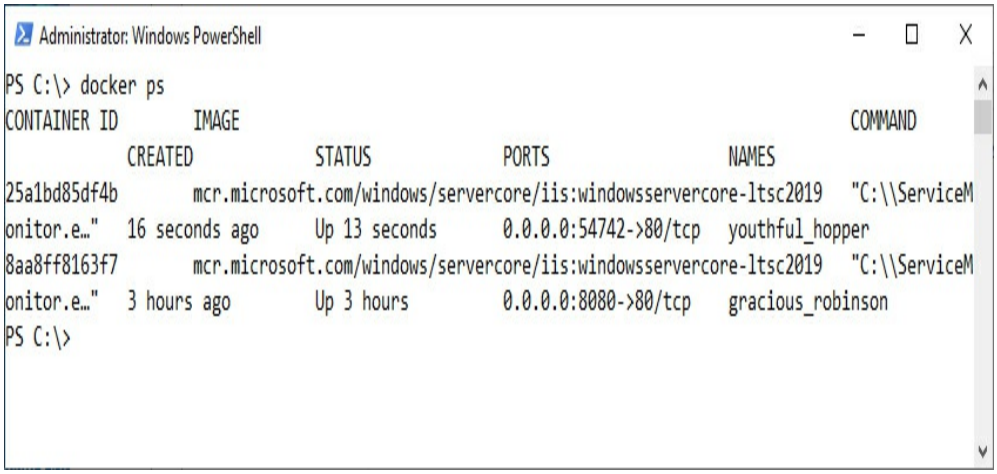


Figure 10-18 *docker ps* command

When you configure port mapping, firewall rules on the container host will be created automatically that will allow traffic through.

Transparent

Transparent networks allow each container endpoint to connect to the same network as the container host. You can use the Transparent Networking mode by creating a container network that has the driver name *transparent*. The driver name is specified with the *-d* option. Use this command:

Click here to view code image

```
docker network create -d transparent TransparentNetworkName
```

If the container host is a virtual machine running on a Hyper-V host, and you want to use DHCP for IP address assignment, it's necessary to enable

MACAddressSpoofing on the VM network adapter, as shown in Figure 10-19. The Transparent Network mode supports IPv6. If you are using the Transparent Network mode, you can use a DHCPv6 server to assign IPv6 addresses to containers.

Technet24

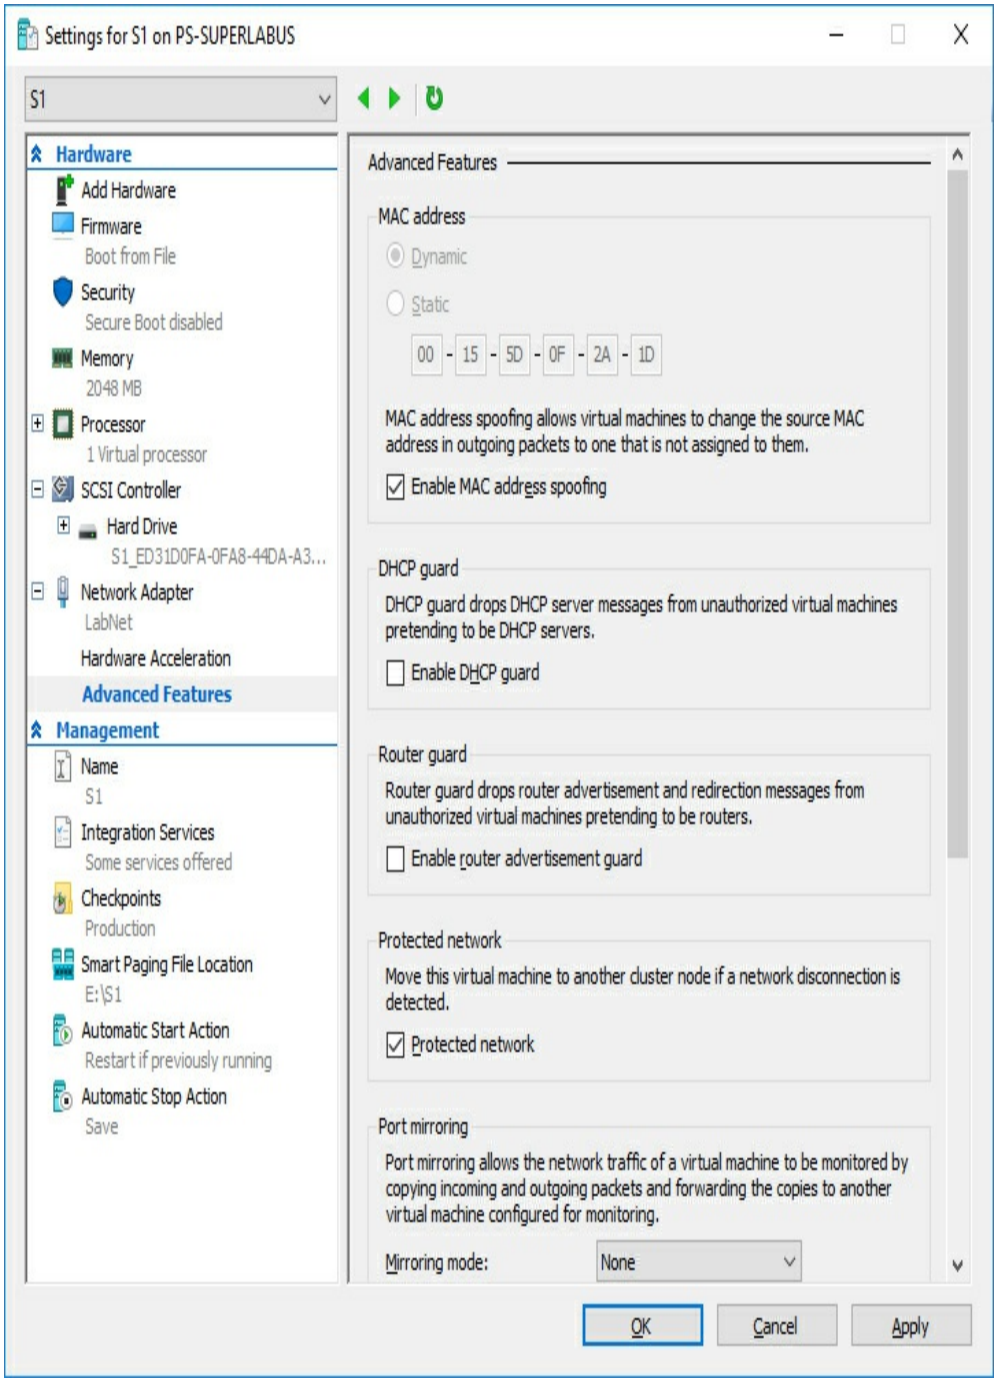


Figure 10-19 Enabling MAC address spoofing

If you want to manually assign IP addresses to

containers, you must specify the subnet and gateway parameters when you create the transparent network. These network properties need to match the network settings of the network to which the container host is connected.

For example, let's say your container host is connected to a network that uses the 192.168.30.0/24 network and uses 192.168.30.1 as the default gateway. Run this command to create a transparent network that will allow static address assignment for containers on this network called *TransNet*:

Click here to view code image

```
Docker network create -d transparent --subnet=192.168.30.0/24 --
gateway=192.168.30.1
TransNet
```

Once the transparent network is created with the appropriate settings, you can specify an IP address for a container using the *--ip* option. For example, to start a new container from the *microsoft/windowsservercore* image, enter the command prompt within the container. Run this command to assign it the IP address *192.168.30.101* on the *TransNet* network:

Click here to view code image

```
Docker run -it --network:TransNet --ip 192.168.30.101
microsoft/windowsservercore cmd.exe
```

Just like when you use a transparent network,

containers are connected directly to the container host's network; you don't need to configure port mapping into the container.

Overlay

You can only use Overlay Networking mode if the Docker host is running in Swarm mode as a Manager node. Each overlay network you create on a swarm cluster has its own IP subnet defined by an IP address prefix in the private address space. You create an overlay network by specifying *overlay* as the driver. For example, you can create an overlay network for the subnet 192.168.50.0/24 with the name *OverlayNet* by running the following command from a Swarm Manager node:

Click here to view code image

```
docker network create --driver=overlay -subnet=192.168.50.0/24
OverlayNet
```

Layer 2 Bridge

Layer 2 Bridge (L2 Bridge) networks are similar to transparent networks in that they allow containers to have IP addresses on the same subnets as the container host. Layer 2 Bridge networks differ in that IP addresses must be assigned statically; this is because all containers on the container host that use an L2 Bridge network have the same MAC address.

When creating an L2 Bridge network, you must specify the network type as *l2bridge*. You must also specify subnet and default gateway settings that matches the subnet and default gateway settings of the container host. For example, use the following command to create an L2 Bridge network named *L2BridgeNet* for the IP address range 192.168.88.0/24 and with the default gateway address 192.168.88.1:

Click here to view code image

```
Docker network create -d l2bridge –subnet=192.168.88.0/24 --
gateway=192.168.88.1
L2BridgeNet
```

More Info

Container networking

You can learn more about networking of containers at https://docs.microsoft.com/en-us/virtualization/windowscontainers/container-networking/architecture.

LINUX CONTAINERS ON WINDOWS

If you configure Docker to run in Experimental mode, you can run Linux containers on a Windows Server computer that supports the Hyper-V Isolation mode concurrently with containers based on Windows Server images.

To configure Docker to run in the Experimental mode, edit the *daemon.json* in the

c:\programdata\docker\config folder, and add this line:

Click here to view code image

```
{
    "experimental":true
}
```

You'll then need to run the following command from an elevated command prompt to set the appropriate environment variable:

Click here to view code image

```
[Environment]::SetEnvironmentVariable("LCOW_SUPPORTED", "1",
"Machine")
```

You'll also need to go to the following GitHub repository at *https://github.com/linuxkit/lcow/releases* and download the files named *initrd.img* and *kernel* and place them in the *C:\Program Files\Linux Containers* folder, as shown in Figure 10-20.

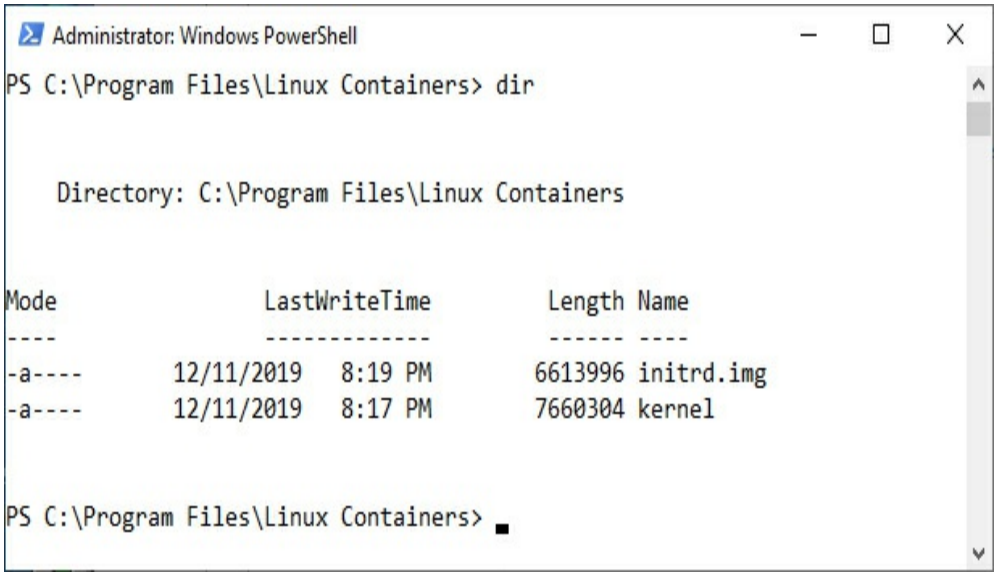


Figure 10-20 Files required for Linux Containers on Windows

Once you've performed these steps, restart the Docker service by running this command:

Click here to view code image

```
Restart-service *docker*
```

Once you've taken these steps, you'll be able to execute the following command to pull and run a Linux image, such as *busybox*, as shown in Figure 10-21.

Click here to view code image

```
docker run --platform linux busybox uname -a
```

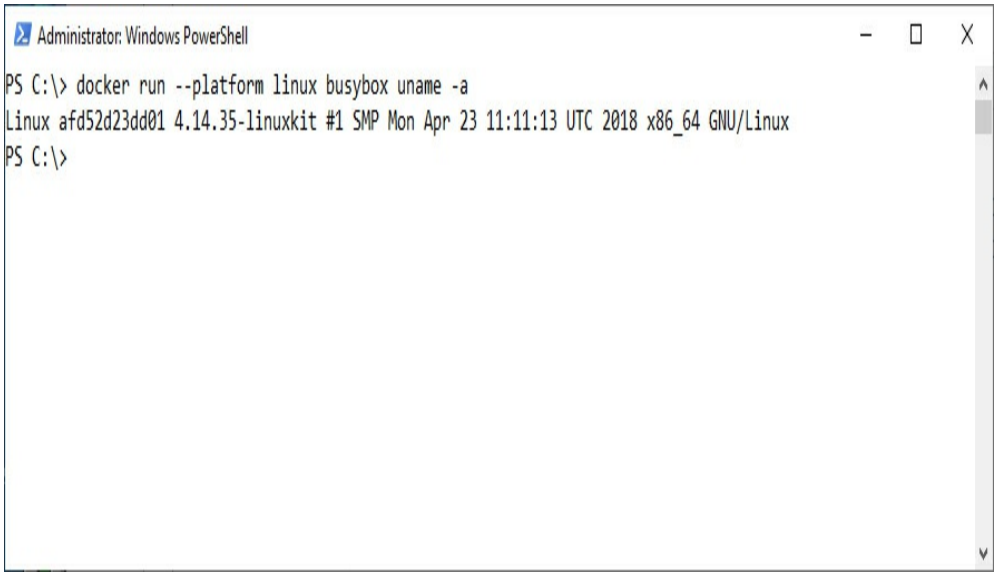


Figure 10-21 A Linux container running on Windows Server

Technet24

You only need to use the *--platform* option the first time you run a container that needs to be retrieved from a repository.

CONTAINER ORCHESTRATION

Container Orchestration allows you to perform the following tasks with containers:

- **Scheduling.** Allows container instances to be started on specific container hosts that participate in the orchestration based on the container host load.
- **Affinity.** Allows you to ensure that container instances that need to be near each other are kept in proximity. Can also be configured using anti-affinity to ensure that container instances that need to be highly available are kept on separate container hosts.
- **Health monitoring.** Determines which containers may have become unresponsive and replaces them with functional container instances.
- **Failover.** Tracks the container hosts. If a container host fails, the container instances that were running on that host are started on another host.
- **Scaling.** Allows new instances of containers to be started or stopped to meet the demand on the application. As load increases, the orchestrator starts containers. As load decreases, the orchestrator stops containers.
- **Networking.** Creates a special overlay network that allows containers to communicate when hosted on separate container hosts.
- **Service discovery.** Allows containers to track other containers as those containers are shifted between container hosts and as they change IP addresses.

- **Application upgrades.** Orchestration can ensure that containers are upgraded in such a way that application down time is minimized. Orchestration can also ensure that rollback occurs in the event that an upgrade causes problems.

Docker on Windows Server supports two forms of orchestration: Kubernetes and Docker Swarm.

More Info

Container orchestration

You can learn more about container orchestration at: https://docs.microsoft.com/en-us/virtualization/windowscontainers/about/overview-container-orchestrators.

Kubernetes

A Windows Server container host can participate in a Linux-based Kubernetes orchestration cluster. This allows Windows Server–based container images to take advantage of Kubernetes orchestration. In Windows Server 2019, a Kubernetes cluster must be managed by a Linux server, and you cannot deploy a Windows Server–only Kubernetes cluster. Both Linux worker container hosts and Windows Server container hosts can participate in the same Kubernetes cluster.

More Info

Kubernetes and Windows

You can learn more about Kubernetes and Windows at https://docs.microsoft.com/en-us/virtualization/windowscontainers/kubernetes/getting-started-kubernetes-windows.

Docker Swarm

Docker Swarm mode provides container orchestration. This includes native clustering of container hosts and scheduling of container workloads. You can configure a collection of container hosts to form a swarm cluster. To do this, you need to configure the Docker Engine on the container hosts to run together in Swarm mode.

Swarms include two types of container hosts:

- **Manager nodes.** Swarms are initialized from Manager nodes. All Docker commands for the control and management of the swarm must be run from a Manager node. A swarm must have at least one Manager node, but it's possible for a swarm to have multiple Manager nodes. Manager nodes ensure that the state of the swarm matches the intended state. They keep track of services running on the swarm.
- **Worker nodes.** Manager nodes assign Worker nodes tasks that the Worker nodes execute. Worker nodes use a *join* token, which are generated during swarm initialization to join a swarm.

Swarm mode has the following requirements:

- Docker Engine v1.13.0 or later
- TCP port 2377 for cluster management communication open on each node
- TCP and UDP port 7946 for intra-node communication
- TCP and UDP port 4789 for overlay network traffic

Creating swarm clusters

Run the following command on the container host that will function as a Manager node and specify the IP

address of the container host:

Click here to view code image

```
Docker swarm init --advertise-addr <container_host_ip> --listen-
addr
<container_host_ip>:2377
```

The container host will now function as the first Manager node in a new Docker swarm. Running the *Docker swarm init* command will also generate a *join* token that you will use to add other nodes to the swarm. You can retrieve the Worker node *join* token at any point in the future from the Manager node by running the following command:

Click here to view code image

```
Docker swarm join-token worker -q
```

Joining additional Manager nodes requires the Manager node *join* token. You can retrieve the Manager node *join* token at any point by running the following command on the Manager node.

Click here to view code image

```
Docker swarm join-token manager -q
```

You can add container hosts to the swarm as Worker nodes by running the following command:

Click here to view code image

```
Docker swarm join --token <worker-join-token>
<manager_container_host_ip>
```

You can add container hosts as additional Manager nodes to the swarm by running the following command:

Click here to view code image

```
Docker swarm join --token <manager-join-token>
<manager_container_host_ip>
```

You can view which nodes are members of a swarm by running the command:

```
Docker node ls
```

Creating overlay networks

Overlay mode allows container endpoints to be connected across multiple container hosts. In a multi-node environment, VXLAN (Virtual Extensible LAN) encapsulation occurs in VFP (Virtual Filtering Platform) forwarding extensions included in the Hyper-V virtual switch. Inter-host communications directly reference IP endpoints. In single-node environments, intra-host communication occurs through a bridged connection on the Hyper-V virtual switch.

You can create an overlay network by running the following command from a Manager node, where 10.0.10.0/24 defines the overlay network address space:

Click here to view code image

```
Docker network create --driver=overlay --subnet=10.0.10.0/24
<OverlayNetworkName>
```

Deploying and scaling swarm services

A swarm service consists of a specific container image that you want to run in multiple instances across nodes in the swarm cluster. Docker Swarm allows you to scale the service as required. For example, if you had a container image that allowed you to deploy a front-end web server for a multi-tier application, you can deploy that container image as a swarm service. As requirements dictate, you can increase the number of container instances by scaling the swarm service. Traffic can be automatically load balanced across the swarm service using DNS Round-Robin. This is substantially simpler than starting new container instances across separate container hosts and manually configuring load balancing to take account of the new container instances.

To create a swarm service, you use the *Docker service create* command to specify the service name, network name, and container image name, and to set the Endpoint mode to DNS Round-Robin. For example, to create a service named *SWARMFRONT* using the *WEBFRONTEND* container image and the *OverNet* overlay network, issue this command:

Click here to view code image

```
Docker service create –name=SWARMFRONT --endpoint-mode=dnsrr --
network=OverNet
WEBFRONTEND
```

Once the service is created, you can scale it using the *Docker service scale* command and specifying the number of container instances. For example, to scale the *SWARMFRONT* service to 10 instances, run this command:

Click here to view code image

```
Docker service scale SWARMFRONT=10
```

You can have more than one service present on a Docker Swarm. You can view which services are present by running the following command on a Manager node:

```
Docker service ls
```

You can view which container instances make up a service, including which swarm nodes individual containers are running on, by running the following command:

Click here to view code image

```
Docker service ps <servicename>
```

More Info

Swarm Mode

You can learn more about Swarm mode at https://docs.microsoft.com/en-us/virtualization/windowscontainers/manage-containers/swarm-mode.

Chapter 11

Clustering and high availability

Failover clustering

Network Load Balancing

The primary high-availability technology available in Windows Server 2019 is failover clustering. Clustering allows you to ensure that your workloads remain available if server hardware or even a site fails. Cluster sets are a new feature available in Windows Server 2019, and they allow clusters to be loosely federated. Windows Server 2019 also includes network load-balancing functionality, which allows you to scale out non-stateful workloads, such as web server front ends.

FAILOVER CLUSTERING

Failover clustering is a stateful, high-availability solution that allows an application or service to remain

available to clients if a host server fails. You can use failover clustering to provide high availability to applications such as SQL Server and scale out file servers and virtual machines.

Failover clustering is supported in both the Standard and Datacenter editions of Windows Server 2019. In some earlier versions of the Windows Server operating system, you only gained access to failover clustering if you used the Enterprise edition. Windows Server 2019 supports up to 64 nodes in a failover cluster.

Generally, all servers in a cluster should run either a similar hardware configuration or should be similarly provisioned virtual machines. You should also use the same edition and installation option. For example, you should aim to have cluster nodes that either run the full GUI or Server Core versions of Windows Server, but you should avoid having cluster nodes that have a mix of computers running Server Core and the full GUI version. Avoiding this mix ensures that you use a similar update routine. A similar update routine is more difficult to maintain when you use different versions of Windows Server.

You should use the Datacenter edition of Windows Server when building clusters that host Hyper-V virtual machines because the virtual machine–licensing scheme available with this edition provides the most

VM licenses.

To be fully supported by Microsoft, cluster hardware should meet the Certified for Windows Server logo requirement. An easy way of accomplishing this is to purchase and deploy Azure Stack HCI, a pre-built hyper-converged Windows Server installation available from select vendors. Even though it is called Azure Stack HCI and sounds as though it is far more of a cloud-based solution, it's primarily just an optimized Windows Server deployment on a certified configuration with all the relevant clustering and Software-Defined Datacenter features lit up.

You can use serial-attached SCSI (SAS), iSCSI, Fibre Channel, or Fibre Channel over Ethernet (FcoE) to host shared storage for a Windows Server failover cluster. Failover clustering only supports IPv4- and IPv6-based protocols. You install failover clustering by installing the Failover Clustering feature and the Failover Clustering Remote Server Administration Tools (RSAT). RSAT allows you to manage other failover clusters in your environment.

Inside OUT

The future of clustering

In the past, you likely clustered workloads such as Exchange and SQL to ensure that they were highly available. Today, both products come with technologies that allow you to have a highly available solution without deploying on a failover cluster. These technologies are Always On Availability Groups and Database Availability Groups. Today, you're more likely to deploy these workloads on virtual machines

and to ensure that the virtual machine host is configured in a failover cluster than you are than to directly make the workloads highly available by installing them on a failover cluster.

Cluster quorum modes

A cluster quorum mode determines how many nodes and witnesses must fail before the cluster is in a failed state. Nodes are servers that participate in the cluster. Witnesses can be stored on shared storage, on file shares, in Windows Server 2019, and even on a USB drive attached to a network switch; shared storage is the preferred method. This recommendation might eventually change, however, as continuously available file shares are more widely adopted.

Inside OUT

No DFS Witness

For unknown reasons, some people use DFS shares as file share witnesses when setting up their failover clusters. To stop this type of shenanigan from occurring in the future, Microsoft has configured Windows Server 2019 failover clustering so that it explicitly blocks the use of DFS namespaces when configuring a file share witness.

Microsoft recommends that you configure a cluster so that an odd number of total votes is spread across member nodes and the witness. This limits the chance of a tie during a quorum vote.

There are four cluster quorum modes:

- **Node Majority.** This cluster quorum mode is recommended for

clusters that have an odd number of nodes. When this quorum type is set, the cluster retains quorum when the number of available nodes exceeds the number of failed nodes. For example, if a cluster has five nodes and three are available, quorum is retained.

- **Node And Disk Majority.** This cluster quorum mode is recommended when the cluster has an even number of nodes. A disk witness hosted on a shared storage disk, for example iSCSI or Fibre Channel, that is accessible to cluster nodes has a vote when determining quorum, as do the quorum nodes. The cluster retains quorum as long as the majority of voting entities remain online. For example, if you have a four-node cluster and a witness disk, a combination of three of those entities needs to remain online for the cluster to retain quorum. The cluster retains quorum if three nodes are online or if two nodes and the witness disk are online.
- **Node And File Share Majority.** This configuration is similar to the Node And Disk Majority configuration, but the quorum is stored on a network share rather than on a shared-storage disk. It is suitable for similar configurations to Node And Disk Majority. This method is not as reliable as Node And Disk Majority because file shares generally do not have the redundancy features of shared storage.
- **No Majority: Disk Only.** This model can be used with clusters that have an odd number of nodes. It is only recommended for testing environments because the disk hosting the witness functions as a single point of failure. When you choose this model, as long as the disk hosting the witness and one node remain available, the cluster retains quorum. If the disk hosting the witness fails, quorum is lost, even if all the other nodes are available.

When you create a cluster, the cluster quorum is automatically configured for you. You might want to alter the quorum mode, however, if you change the number of nodes in your cluster. For example, you might want to alter the quorum mode if you change

from a four-node to a five-node cluster. When you change the cluster quorum configuration, the Failover Cluster Manager provides you with a recommended configuration, but you can choose to override that configuration if you want.

You can also perform advanced quorum configuration to specify what nodes can participate in the quorum vote, which you can set on the Select Voting Configuration page of the Configure Cluster Quorum Wizard. When you do this, only the selected nodes' votes are used to calculate quorum. Also, it's possible that fewer nodes would need to fail to cause a cluster to fail than would otherwise be the case if all nodes participated in the quorum vote. This can be useful when configuring how multisite clusters calculate quorum when the connection between sites fails.

Cluster storage and cluster shared volumes

Almost all failover cluster scenarios require access to some form of shared storage. Windows Server failover clusters can use SAS, iSCSI, or Fibre Channel for Shared Storage. With the inclusion of the iSCSI Target software in Windows Server, iSCSI is one of the simplest and cheapest of these technologies to implement.

You should configure disks used for failover clustering as follows:

- Volumes should be formatted using NTFS or ReFS.
- Use Master Boot Record (MBR) or GUID Partition Table (GPT).
- Avoid allowing different clusters access to the same storage device. This can be accomplished through LUN masking or zoning.
- Any multipath solution must be based on Microsoft Multipath I/O (MPIO).

Cluster Shared Volumes (CSV) is a technology that was introduced in Windows Server 2008 R2 that allows multiple cluster nodes to have concurrent access to a single Logical Unit Number (LUN). Prior to Windows Server 2008 R2, only the active node had access to shared storage. CSV allows you to have virtual machines on the same LUN run on different cluster nodes. CSV also has the following benefits:

- Support for scale-out file servers
- Support for BitLocker volume encryption
- SMB 3.0 support
- Integration with Storage Spaces
- Online volume scan and repair

You can enable CSV only after you create a failover cluster and you have provided the storage that will be available to the CSV.

Cluster networks

While you can create failover clusters with nodes that have a single network adapter, best practice is to have

separate networks and network adapters for the following:

- A connection for cluster nodes to shared storage
- A private network for internal cluster communication
- A public network that clients use to access services hosted on the cluster

In scenarios where high availability is critical, you might have multiple redundant networks connected through several separate switches. If you have a cluster where everything is connected through one piece of network hardware, you can almost guarantee that piece of network hardware is the first thing that fails.

You can either use IPv4 or IPv6 addresses that are dynamically or statically assigned, but you should not use a mix of dynamically and statically assigned IP addresses for nodes that are members of the same cluster. If you use a mixture of dynamically and statically assigned IP addresses, the Validate A Configuration Wizard generates an error.

Inside OUT

Default Gateways

Even if the cluster validation wizard only gives you warnings when you perform the test, you cannot create a failover cluster unless each node is configured with a default gateway. The default gateway doesn't have to be a host that exists, but if you're having trouble in your virtual machine lab with creating a failover cluster, go back and check whether you've configured a default gateway for each node.

MPIO

If there are multiple paths to physical storage on a server, you will need to enable multipath I/O (MPIO). Enabling MPIO aggregates multiple paths to a physical disk into a single logical path for data access. Enabling MPIO improves resiliency to failure. You should enable MPIO on all nodes that will participate in the cluster. You can enable MPIO with the following PowerShell command:

Click here to view code image

```
Install-WindowsFeature -ComputerName Node1 -Name MultiPath-IO
```

Cluster Aware Updating

Cluster Aware Updating (CAU) is a feature included with Windows Server that allows you to automate the process of applying software updates to failover clusters. Cluster Aware Updating integrates with Windows Update, Windows Server Update Services (WSUS), System Center Configuration Manager, and other software update-management applications.

CAU uses the following process:

1. Obtains the update files from the source location.
2. Puts the first node into maintenance mode.
3. Moves any cluster roles off the node to other nodes in the cluster.
4. Installs software updates.
5. Restarts the cluster node, if necessary.

6. Checks for additional updates. If updates are found, it performs Steps 4 through 6 until all updates are applied.
7. Brings the node out of maintenance mode.
8. Reacquires clustered roles that were moved to other nodes.
9. Puts the next node into maintenance mode and repeats the cycle starting at Step 3 until all nodes have been updated.

The main benefit of CAU is that it updates a process that previously had to be performed manually. You can configure CAU to automatically apply updates to cluster nodes when they are approved through WSUS or System Center Configuration Manager, or you can trigger CAU manually as needed.

You can configure the CAU options so that the updates are rolled back across the cluster if an update fails to install on a node after a specified number of attempts. You can also configure Advanced Options, such as requiring scripts to run before and after the update process has occurred and whether services hosted on the cluster require special shutdown or startup configurations.

Failover and preference settings

Cluster preference settings allow you to configure the preferred owner for a specific cluster role, and you can also configure different preferred owners for different cluster roles. Where possible, the role is hosted on the preferred owner. You can configure a list of preferred

owners, so that if the most preferred owner isn't available, the next preferred owner hosts the role. You configure a role-preferred owner on the role's Properties dialog box.

You configure whether the clustered role fails back to the preferred owner on the Failover tab of the cluster role's Properties dialog box. When configuring failback, you need to:

- Determine whether you want to prevent failback
- Determine whether you want to have failback occur automatically as soon as the preferred owner is in a healthy state
- Configure failback to occur within a certain number of hours of the preferred owner returning to a healthy state

Node quarantine settings allow you to configure a node so that it is unable to rejoin the cluster if the node fails a certain number of times within a certain period. Configuring node quarantine blocks workloads from being placed back on a quarantined node until the reason for the repeated failure can be dealt with by a server administrator.

A fault domain is a collection of hardware components that share a common point of failure. Windows Server 2019 allows you to configure fault domains at the site, rack, chassis, and node levels. You can use Storage Spaces and Storage Spaces Direct with fault domains to ensure that copies of data are kept in separate fault

domains. For example, if you defined fault domains at the site level, you can configure storage spaces so that redundant data resides in separate sites rather than residing only on separate cluster nodes.

Multisite clusters

Failover clusters can span multiple sites. When configuring a cluster that spans two sites, you should consider the following:

- Ensure that there are an equal number of nodes in each site.
- Allow each node to have a vote.
- Enable dynamic quorum. Dynamic quorum allows quorum to be recalculated when individual nodes leave the cluster one at a time. Dynamic quorum is enabled by default on Windows Server 2019 failover clusters.
- Use a file share witness. You should host the file share witness on a third site that has separate connectivity to the two sites that host the cluster nodes, as shown in Figure 11-1. When configured in this manner, the cluster retains quorum if one of the sites is lost. An alternative to a file share witness is a cloud witness.

Technet24

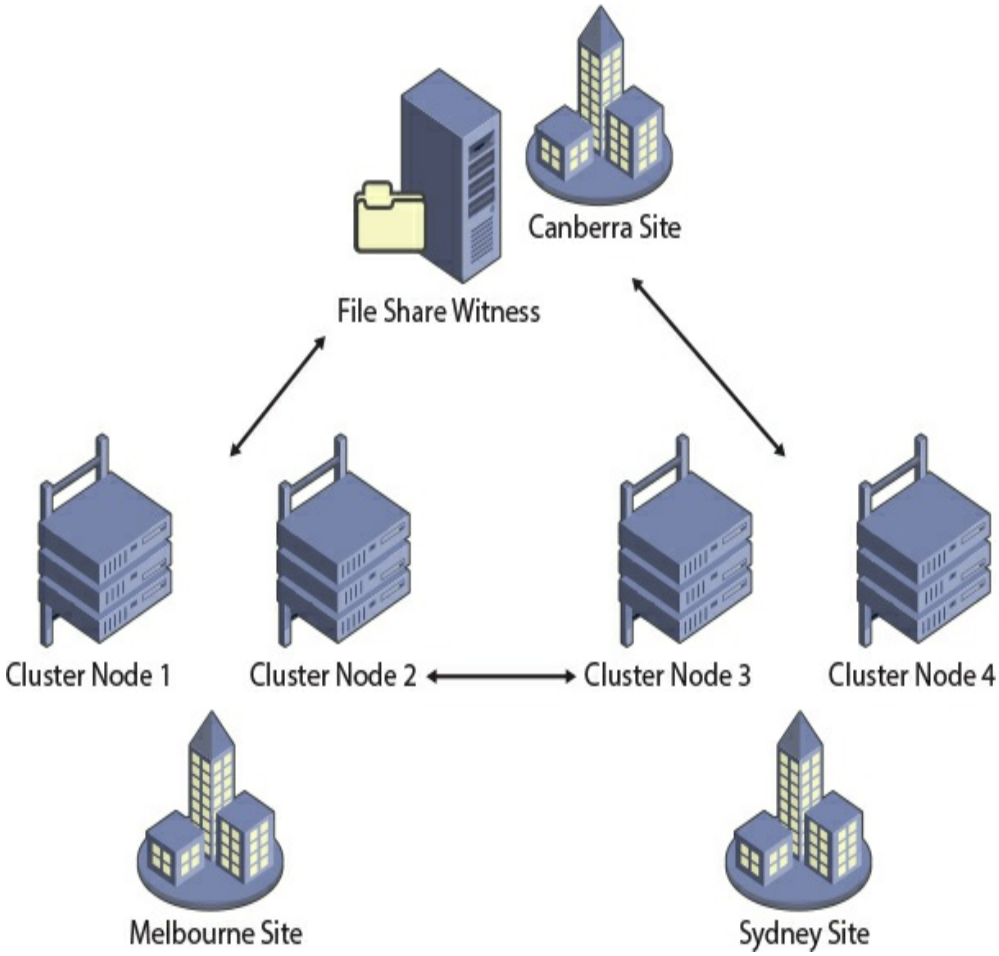


Figure 11-1 Multisite cluster

If you have only two sites and are unable to place a file share witness in an independent third site, you can manually edit the cluster configuration to reassign votes so that the cluster recalculates quorum.

Manually reassigning votes is also useful to avoid split-brain scenarios. Split-brain scenarios occur when a failure occurs in a multisite cluster and when both sides of the cluster believe they have quorum. Split-brain

scenarios cause challenges when connectivity is restored and make it necessary to restart servers on one side of the multisite cluster to resolve the issue. You can manually reassign votes so that one side always retains quorum if intersite connectivity is lost. For example, by setting the Melbourne site with two votes and the Sydney site with one vote, the Melbourne site always retains quorum if intersite connectivity is lost.

Cloud witness

Cloud witness has the cluster witness role hosted in Azure rather than at a third site. Cloud witness is suitable for multisite clusters. To configure a cloud witness, create a storage account in Azure, copy the storage access keys, note the endpoint URL links, and then use these links with the Configure Cluster Quorum Wizard and specify a cloud witness as shown in Figure 11-2.

Technet24

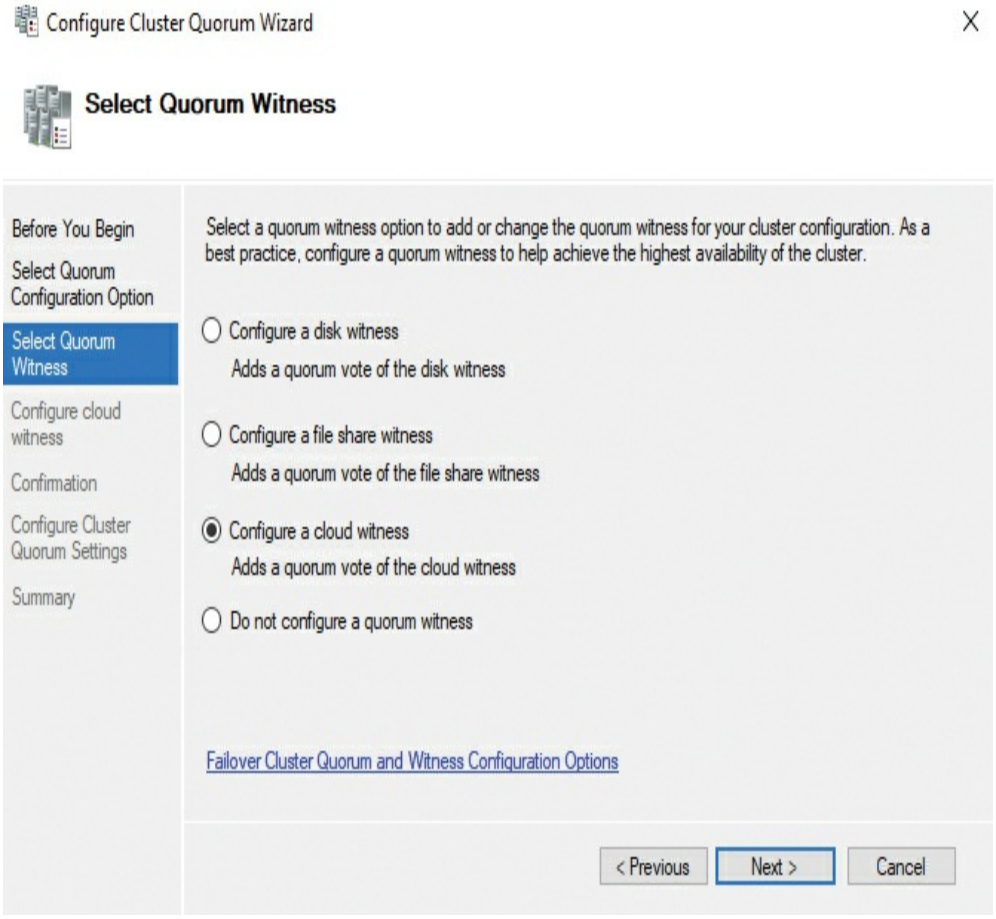


Figure 11-2 Configure a cloud witness

More Info

Cloud Witness

You can learn more about cloud witness at at https://docs.microsoft.com/en-us/windows-server/failover-clustering/deploy-cloud-witness.

Virtual machine failover clustering

One of the most common uses for failover clusters is hosting virtual machines. By deploying a workload such as SQL Server or Exchange on a highly available virtual

machine, you can achieve high availability without the need for the application to be aware that it is now highly available. The virtual machine functions normally, provides services to clients on the network, and can switch between cluster nodes as necessary in the event that the individual cluster node hosting it requires maintenance or experiences some sort of failure. Building a Hyper-V failover cluster first involves creating a failover cluster and then adding the Hyper-V role to each node of the cluster.

Virtual machine storage

You should use cluster shared volumes to store virtual machines on a Hyper-V cluster because cluster shared volumes allow multiple cluster nodes to manage a single shared storage device. This allows you to put multiple virtual machines on the same shared storage device but have those virtual machines hosted by different nodes in the failover cluster. Cluster-shared volumes are mapped under the *C:\ClusterStorage* folder on cluster nodes.

When creating a new virtual machine on a failover cluster, first select which cluster node hosts the virtual machine.

When creating a highly available virtual machine, specify the cluster-shared volume path as the location to store the virtual machine. If you have an existing

machine that you want to make highly available, you can move the virtual machine to this path. As an alternative, you also have the option to specify a SMB 3.0 file share as the storage location for the highly available virtual machine. Whether to select a cluster shared volume or an SMB 3.0 file share depends on your organization's storage configuration.

After the virtual machine is created, you can control it by using the Failover Cluster Manager console. The move option in the Actions pane allows you to select the cluster node to which you want to move the virtual machine.

In production environments, you should ensure that each Hyper-V host has an identical hardware configuration. However, in development environments, this is not always possible. If different processor types are used—for example, an Intel processor on one node and an AMD processor on another—you might have to perform a quick migration. Quick migration allows migration between nodes but does cause a disruption in client connectivity. You can allow migration between Hyper-V nodes with different processor types or versions by enabling the processor compatibility setting on the virtual machine.

VM load balancing

VM load balancing is a feature of Windows Server 2019

that identifies nodes in a Hyper-V cluster where CPU and memory resources are under pressure; VM load balancing redistributes workloads to other nodes that have lower CPU and/or memory utilization. By default, VM load balancing is enabled, and it automatically moves VMs to a new node after you add that node to an existing Windows Server 2019 Hyper-V cluster.

You can configure VM load balancing so that workload redistribution happens automatically when a specific threshold is reached, or you can configure load balancing to occur on a periodic basis. VMs are moved between Hyper-V cluster nodes using live migration without incurring downtime. You configure VM load balancing aggressiveness by using the following Windows PowerShell command:

Click here to view code image

```
(Get-Cluster).AutoBalancerLevel = <value>
```

The values listed in Table 11-1 determine automatic load balancer thresholds.

Table 11-1 Auto VM load balancer aggressiveness

AutoBalancerLevel	Aggressiveness	Action
1 (default)	Low	Move workloads when the node reaches 80 percent memory or processor utilization

2	Medium	Move workloads when the node reaches 70 percent memory or processor utilization
3	High	Move workloads when the node reaches 60 percent memory or processor utilization

You configure how the balancer functions by using the following Windows PowerShell command:

Click here to view code image

```
(Get-Cluster).AutoBalancerMode = <value>
```

The values listed in Table 11-2 determine automatic load balancer mode.

Table 11-2 Auto VM load balancer mode

AutoBalancer mode	Aggressiveness
0	Disables VM load balancing
1	Load balances each time a new node joins the Hyper-V cluster
2 (default)	Load balances each time a new node joins the Hyper-V cluster and every 30 minutes

Rolling upgrades

Cluster OS Rolling Upgrade allows you to upgrade a cluster hosting Hyper-V or scale out file server

workloads without taking those workloads offline. You can perform a Cluster OS Rolling Upgrade from Windows Server 2012 R2 to Windows Server 2016, or you can perform a rolling upgrade from Windows Server 2016 to Windows Server 2019. Other cluster hosted workloads, such as SQL Server, are only offline during failover from an original cluster node to a new or upgraded cluster node.

Upgrading a cluster involves either adding new nodes running Windows Server 2016 to an existing cluster running Windows Server 2012 R2 or upgrading existing nodes running Windows Server 2012 R2 to Windows Server 2016. The cluster can function normally with both Windows Server 2012 R2 and Windows Server 2016 nodes.

The nodes that you can add to a cluster depend on the cluster functional level. A cluster that only has nodes running Windows Server 2012 R2 is running at cluster functional level 8. Cluster functional level 8 supports nodes that run both Windows Server 2012 R2 and Windows Server 2016. You can also introduce cluster nodes that run either Windows Server 2012 or 2016 to the cluster as long as the cluster is still configured to use cluster functional level 8. It’s even possible to roll back to Windows Server 2012 R2 when all nodes are running Windows Server 2016 as long as you haven’t upgraded the cluster functional level from level 8.

A cluster running at functional level 9 supports nodes running Windows Server 2016 and Windows Server 2019. A cluster running at functional level 10 only supports nodes running Windows Server 2019.

You should only upgrade the cluster functional level when you are absolutely certain that you don't want to roll back the cluster upgrade and that all nodes have been upgraded. You can determine which functional level a cluster is running at with the following Windows PowerShell command:

Click here to view code image

```
Get-Cluster | Select ClusterFunctionalLevel
```

You can also use the *Get-ClusterNodeSupportedVersion* cmdlet on a computer running Windows Server 2019 to determine which nodes in the cluster are running at a specific functional level and cluster upgrade version.

You can upgrade the cluster functional level by using the *Update-ClusterFunctionalLevel* cmdlet.

Cluster OS Rolling Upgrade has the following conditions:

- You can only upgrade a Hyper-V cluster if the CPUs support Second-Level Addressing Table (SLAT). This is because Windows Server 2016 and Windows Server 2019 Hyper-V is only supported on processors that support SLAT.
- When performing the rolling cluster upgrade, always use the most recent version of Windows Server to add nodes to the cluster.

- All cluster management tasks should be performed from the more recent operating system node than the previous version node.
- Cluster rolling upgrade cannot be used to upgrade VM guest clusters that use a shared virtual hard disk as shared storage. VM guest clusters are clusters in which each node is a virtual machine. A VM guest cluster that uses a shared virtual hard disk for shared storage cannot be upgraded, but a VM guest cluster that uses another form of shared storage, such as iSCSI, can be upgraded by using a rolling cluster upgrade.
- Before upgrading a node, drain the node of workloads and evict the node from the cluster. After you've upgraded the node's operating system, you can join it to the cluster again.
- Prior to commencing Cluster OS Rolling Upgrade, ensure that all elements of the cluster and its workloads are backed up.
- You can't lower the cluster functional level after it has been raised.
- Cluster OS Rolling Upgrade only works when the original cluster is running Windows Server 2012 R2 or later. Cluster OS Rolling Upgrade does not work with Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008.

More Info

Cluster OS Rolling Upgrade

You can learn more about rolling upgrades at https://docs.microsoft.com/en-us/windows-server/failover-clustering/cluster-operating-system-rolling-upgrade.

Workgroup clusters

Workgroup clusters are a special type of cluster in which cluster nodes are not members of an Active Directory domain. Workgroup clusters are also known as Active Directory Detached Clusters. The following workloads

are supported for workgroup clusters:

- **SQL Server.** When deploying SQL Server on a workgroup cluster, it's recommended that you use SQL Server Authentication for databases and SQL Server Always On Availability Groups.
- **File Server.** A supported but not recommended configuration as Kerberos will not be available as an authentication protocol for SMB traffic.
- **Hyper-V.** A supported but not recommended configuration. Hyper-V live migration is not supported, through it is possible to perform quick migration.

When creating a workgroup cluster, you first need to create a special account on all nodes that will participate in the cluster that has the following properties:

- The special account must have the same username and password on all cluster nodes.
- The special account must be added to the local Administrators group on each cluster node.
- The primary DNS suffix on each cluster node must be configured with the same value.
- When creating the cluster, ensure that the *AdministrativeAccessPoint* parameter when using the *New-Cluster* cmdlet is set to *DNS*. Ensure that the cluster name is present in the appropriate DNS zone, which depends on the primary DNS suffix, when running this command.
- You will need to run the following PowerShell command on each node to configure the *LocalAccountTokenFilterPolicy* registry setting to *1*.

Click here to view code image

```
new-itemproperty -path
HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System
```

```
-Name LocalAccountTokenFilterPolicy -Value 1
```

To create a workgroup cluster, use the *New-Cluster* cmdlet with the *name* parameter listing the cluster name, the node parameters listing the nodes that you wish to join to the cluster where the nodes have been configured according to the prerequisites, and the *administrativeaccesspoint* parameter configured for *DNS*. For example, to create a new workgroup cluster named *workgrpclst* with member nodes *node1* and *node2*, run the following command on one of the nodes:

Click here to view code image

```
New-Cluster -name workgrpclst -node node1,node2 -
AdministrativeAccessPoint DNS
```

Cluster sets

Cluster sets are a Windows Server 2019 technology that provide a method of grouping multiple failover clusters. When you implement cluster sets, you can combine smaller clusters into larger virtual clusters. These virtual clusters support virtual machine fluidity and a unified storage namespace, meaning that virtual machines can be moved between clusters in a cluster set as easily as they are moved between nodes in a traditional failover cluster. While virtual machines can be live migrated between clusters, Windows Server 2019 cluster sets do not allow virtual machines to fail over between clusters in a cluster set.

Only clusters running at the Windows Server 2019 cluster functional level 10.X or higher can participate in cluster sets. You cannot join clusters running Windows Server 2016 or Windows Server 2012 R2 to a cluster set. All clusters in a cluster set must be members of the same Active Directory forest.

Cluster sets improve the failover cluster lifecycle. Rather than adding and removing nodes from a single cluster when the node hardware is to be retired, you can add a new cluster to the cluster set, migrate workloads across from the original cluster to the new cluster, and then retire that original cluster. Cluster sets are currently supported by Microsoft for up to 64 cluster nodes, which is the same number of nodes supported in an individual failover cluster. That being said, there is no specific limit to the number of nodes that may exist within a cluster set, so going beyond 64 nodes in a cluster set is possible should you want to try it.

Cluster sets consist of a management cluster and member clusters. The management cluster is the cluster set that holds the management role for the cluster set and also hosts the unified storage namespace for the cluster set Scale-Out File Server. The management cluster does not host workloads, such as virtual machines, and its role is to mange the relationship between other clusters in the cluster set and to host the storage namespace. The new role that the management

cluster hosts is termed the “Cluster Set Master” (*CS-Master*). Member clusters hold the “Cluster Set Worker” or *CS-Worker* role.

> **More Info**
>
> **Cluster Sets**
>
> You can learn more about rolling upgrades at https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/cluster-sets.

Managing failover clustering with PowerShell

You can use cmdlets in the *FailoverClusters* module, as described in Table 11-3, to manage failover clusters on Windows Server 2016.

Table 11-3 Failover cluster PowerShell cmdlets

Noun	Verbs	Function
Cluster	Get, Test, Stop, New, Start, Remove	Manage failover clusters
ClusterAccess	Block, Get, Remove, Grant	Manage permissions to control access to a failover cluster
ClusterAvailableDisk	Get	View disks that support failover clustering that are visible to all nodes but are not yet clustered disks
ClusterChe	Add, Remove, Get	Manage cryptographic key or

Technet24

ckpoint		registry checkpoints for a cluster resource
ClusterDiagnosticInfo	Get	Get cluster diagnostic information
ClusterDisk	Add	Add a disk to a cluster
ClusterDiskReservation	Clear	Clears persistent reservation on a disk in a failover cluster
ClusterFaultDomain	Get, Set, Remove, New	Manage cluster fault domains
ClusterFaultDomainXML	Set, Get	Configure cluster fault domains using XML
ClusterFileServerRole	Add	Add the file server role to a cluster
ClusterFunctionalLevel	Update	Upgrade the cluster functional level
ClusterGenericApplicationRole	Add	Configure high availability for an application that was not originally designed to run in a failover cluster
ClusterGenericScriptRole	Add	Configure an application controlled by a script that runs in Windows Script host on a failover cluster

ClusterGenericServiceRole	Add	Configure high availability for a service that was not originally designed to run in a failover cluster
ClusterGroup	Remove, Start, Stop, Add, Get, Move	Manage cluster roles/resource groups in a failover cluster
ClusterGroupFromSet	Remove	Remove a group from a cluster group set
ClusterGroupSet	Remove, Get, New, Set	Manage cluster group sets
ClusterGroupSetDependency	Add, Get, Remove	Manage cluster group set dependencies
ClusterGroupToSet	Add	Add a group to a cluster group set
ClusterIPResource	Update	Renew or release DHCP lease for IP address resource used by a failover cluster
ClusteriSCSITargetServerRole	Add	Create a highly available iSCSI Target server
ClusterLog	Get, Set	Manage cluster log settings
ClusterNa	New	Create a cluster name account in

Technet24

meAccount		Active Directory
ClusterNetwork	Get	View cluster network properties
ClusterNetworkInterface	Get	View cluster network interface properties
ClusterNetworkNameResource	Update	Register cluster resource names in DNS
ClusterNode	Get, Resume, Stop, Start, Remove, Add, Clear, Suspend	Manage specific cluster nodes
ClusterOwnerNode	Set, Get	Configure which nodes can own a resource
ClusterParameter	Set, Get	View and configure information about an object in a failover cluster
ClusterQuorum	Get, Set	Manage cluster quorum
ClusterResource	Remove, Move, Start, Suspend, Add, Get, Stop, Resume	Manage cluster resources
ClusterRes	Get, Set, Add,	Manage cluster resource

ourceDependency	Remove	dependencies
ClusterResourceDependencyReport	Get	Generate a cluster resource dependency report
ClusterResourceFailure	Test	Test failure of a cluster resource
ClusterResourceType	Add, Get, Remove	Manage cluster resource type
ClusterScaleOutFileServerRole	Add	Add a cluster scale-out file server role
ClusterServerRole	Add	Add a cluster server role
ClusterSharedVolume	Get, Add, Remove, Move	Manage cluster shared volumes
ClusterSharedVolumeState	Get	View the state of cluster shared volumes
ClusterStorageSpacesDirect	Repair, Set, Disable, Enable, Get	Manage cluster storage spaces direct

ClusterStorageSpacesDirectDisk	Set	Configure cluster storage spaces direct disks
ClusterVirtualMachineConfiguration	Update	Update the configuration of a clustered virtual machine hosted on a failover cluster
ClusterVirtualMachineRole	Add, Move	Manage clustered virtual machine live migration
ClusterVMMonitoredItem	Add, Remove, Get	Manage clustered virtual machine–monitored items
ClusterVMMonitoredState	Reset	Reset clustered virtual machine monitoring state

NETWORK LOAD BALANCING

Network Load Balancing (NLB) distributes traffic across multiple hosts in a balanced manner. NLB directs new traffic to cluster nodes under the least load. NLB works as a high-availability solution because it detects node failures and automatically redistributes traffic to available nodes. NLB is also scalable, which means that you can start with a two-node NLB cluster and keep

adding nodes until you reach a maximum of 32 nodes. A node is a computer running the Windows Server operating system that participates in the NLB cluster. A computer can only be a member of one Windows NLB cluster.

NLB functions through the creation of a virtual IP and a virtual network adapter with an associated Media Access Control (MAC) address. Traffic to this address is distributed to the cluster nodes. In the default configuration, traffic is distributed to the least-utilized node. You can also configure NLB so that specific nodes are preferred and process more traffic than other nodes in the cluster.

NLB is a high-availability solution that is suitable for stateless applications. Imagine that you have a two-tier web application where you have four servers running IIS as the web tier and two servers running SQL Server as the database tier. In this scenario, you add Web Server 1, Web Server 2, Web Server 3, and Web Server 4 to an NLB cluster as web servers host stateless applications. To make the database servers highly available, you would configure a failover cluster or Always On Availability Group, but you wouldn't configure NLB because SQL Server is a stateful application.

NLB is failure aware as long as the failure occurs on the

node. For example, if you have a two-node NLB cluster that you use with a web application, and the network card fails on one of the nodes, NLB is aware of the failure and stops directing traffic to the node. If, instead of a network card failure, the web application fails, but everything else remains operational, NLB remains unaware of the failure and continues to direct requests to the node that hosts the failed web application. In a real-world environment, you'd set up more sophisticated monitoring that allows you to detect the failure of the application, not just the failure of the node. Some hardware-based NLB solutions are sophisticated enough to detect application failures rather than just host failures.

Network Load Balancing prerequisites

In terms of setup, NLB is straightforward and only requires you to install the feature. Although you can install the feature on a node, you must ensure that it meets some prerequisites before you can add it to a cluster. Before you add a host to an NLB cluster, you must ensure the following:

- All nodes in the NLB cluster must reside on the same subnet. While you can configure a subnet to span multiple geographic locations, the cluster is unlikely to converge successfully if the latency between nodes exceeds 250 milliseconds.
- All network adapters must either be configured as unicast or multicast. You can't use Windows NLB to create an NLB cluster

that has a mixture of unicast and multicast addresses assigned to network adapters.

- If you choose to use the unicast cluster configuration mode, the network adapter needs to support changing its MAC address.
- The network adapter must be configured with a static IP address.
- Only TCP/IP is supported on the network adapter that is used for NLB. You cannot bind other protocols, such as IPX, Token Ring, or Banyan Vines, on the adapter.

There is no restriction on the number of network adapters in each host. You can have one host with three network adapters in a team and another host with five separate adapters participate in the same NLB cluster. You can run nodes on different editions of Windows Server, although considering that the only difference between the Standard and Datacenter edition is virtual machine licensing, you are unlikely to be running NLB on the Datacenter edition. While it is possible to have working NLB clusters if you are running different versions of Windows Server, such as Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019, you should upgrade all nodes as soon as possible to the same operating system. You should also attempt to ensure that nodes have similar hardware capacity or are running on similarly provisioned virtual machines so that a client doesn't get a radically different experience when connecting to different nodes.

NLB cluster operation modes

When configuring an NLB cluster, you can choose between one of three separate cluster-operation modes. The mode that you select depends on the configuration of the cluster nodes and the type of network switch to which the cluster nodes are connected. You can configure the cluster-operation mode when creating the cluster. The operation modes are as follows:

- **Unicast Mode.** If you configure an NLB cluster to use the unicast cluster-operation mode, all nodes in the cluster will use the same unicast MAC address. Outbound traffic from node members uses a modified MAC address that is determined by the cluster host's priority settings. The drawback to using unicast mode with nodes that have a single adapter is that you can only perform management tasks from computers located on the same TCP/IP subnet as the cluster nodes. This restriction doesn't apply, however, when the cluster nodes have multiple network adapters that are not NIC team members. When you use unicast with multiple network adapters, one adapter is dedicated to cluster communication, and you can connect to any other adapters to perform management tasks. You can improve cluster operation by placing the adapters that you use for cluster communication and the adapters you use for management tasks on separate VLANs.
- **Multicast Mode.** This mode is suitable when each node only has one network adapter. When you configure multicast mode, each cluster host keeps its original network adapter's MAC address, but it is also assigned a multicast MAC address. All nodes in the cluster use the same multicast MAC address. You can only use multicast mode if your organization's switches and routers support it. Unless you've got bargain-basement network equipment, it's likely that your current network hardware supports multicast mode for NLB.
- **IGMP Multicast.** Internet Group Management Protocol (IGMP) multicast mode is an advanced form of multicast mode that

reduces the chances of the network switch being flooded with traffic. It works by only forwarding traffic through the ports that are connected to hosts that participate in the NLB cluster, but it also requires a switch that supports this functionality.

Managing cluster hosts

You manage NLB clusters through the Network Load Balancing Manager console. The NLB console is available as part of the RSAT tools, meaning that you can manage one or more NLB clusters from a computer that is not a member of the cluster. You do, however, need to manage clusters by using an account that is a member of the local Administrators group on each cluster node. You can use this console to create and manage clusters and cluster nodes.

After the cluster is set up and functioning, most maintenance operations, such as applying software updates, are performed on cluster nodes. When performing maintenance tasks, you should perform the task on one node at a time so the application that the cluster provides to users continues to remain available. Prior to performing maintenance tasks on a node, you should first stop incoming connections and allow existing connections to be completed naturally. When there are no connections to the node, you can then stop the node, apply the updates, and restart the node. You have the following options for managing cluster nodes:

- **Drainstop.** Blocks new connections to the cluster node but doesn't

terminate existing connections. Use this prior to planned maintenance to gracefully evacuate the cluster of connections.

- **Stop.** Stops the cluster node. All connections to the cluster node from clients are stopped. Use Stop after you use Drainstop so that you can then perform maintenance tasks, such as applying updates.
- **Start.** Starts a cluster node that is in a stopped state.
- **Suspend.** Pauses the cluster node until you issue the Resume command. Using Suspend does not shut down the cluster service, but it does terminate current connections as well as block new connections.
- **Resume.** Resumes a suspended cluster node.

Port rules

An NLB port rule allows you to configure how the NLB cluster responds to incoming traffic on a specific port and protocol, such as TCP port 80 (HTTP) or UDP port 69 (Trivial FTP). Each host in an NLB cluster must use the same port rules. While it is possible to configure port rules on a per-node basis, it's safer to configure them at the cluster level to ensure that they are consistent. The default port rule redirects traffic on any TCP or UDP port to each node in the cluster in a balanced way. If you want to create specific rules, delete the default rule.

Filtering and affinity

When you create a port rule, you also choose a filtering mode, which might require you to configure affinity. Filtering allows you to determine whether incoming

traffic is handled by one, some, or all nodes in the NLB cluster. If you choose to allow multiple nodes to accept incoming traffic, you need to configure affinity. The affinity setting determines whether one node handles all subsequent requests from the same client, or if subsequent requests from the same client are redistributed across other nodes. Affinity is often important with web applications where a consistent session must exist between the client and one web server after the session is established. You can configure the following filtering modes:

- **Multiple Host.** This is the default filtering mode. This mode allows traffic to be directed to any node in the NLB cluster. You also need to specify one of the following affinity settings:
  - **Single.** When you configure this option, the incoming request is distributed to one of the cluster nodes. All subsequent traffic from the originating host is directed to that same node for the duration of the session. This is the default multiple host affinity. Don't confuse Multiple Host, Single Affinity with Single Host filtering.
  - **None.** Incoming traffic is distributed to the cluster node under the least load even if there is an existing session. This means that multiple nodes may handle traffic from a single session.
  - **Network.** This option directs clients to cluster nodes based on the requesting client's network address.
- **Single Host.** A single node handles all traffic sent to the cluster on this port rule. In this case, no load balancing occurs. For example, if you have a four-node cluster, but you want node three to be the only one that handles traffic on TCP port 25, you would configure a port rule with the filtering mode set to Single Host.

- **Disable The Port Range.** When you configure this setting, the NLB cluster drops traffic sent to the cluster that matches the port rule.

Managing NLB with PowerShell

You can use the Windows PowerShell cmdlets listed in Table 11-4 to manage NLB clusters on Windows Server 2019.

Table 11-4 NLB PowerShell cmdlets

Noun	Verbs	Function
NlbClusterNode	Add, Get, Remove, Resume, Set, Start, Stop, Suspend	Configure and manage a cluster node
NlbClusterNodeDip	Add, Get, Remove, Set	Configure the cluster node's dedicated management IP address
NlbClusterPortRule	Add, Disable, Enable, Get, Remove, Set	Create and manage port rules
NlbClusterVip	Add, Get, Remove, Set	Configure the cluster's virtual IP address
NlbCluster	Get, New, Remove, Resume, Set, Start, Stop, Suspend	Configure and manage the cluster
NlbClusterDriver	Get	Provides information

Info		about the cluster driver
NlbClusterNodeNetworkInterface	Get	Provides information about the node's network interface driver
NlbClusterIpv6Address	New	Configure the cluster's IPv6 address
NlbClusterPortRuleNodeHandlingPriority	Set	Manage priority on a per-port rule basis
NlbClusterPortRuleNodeWeight	Set	Configure the node weight on a per-port rule basis

Technet24

Chapter 12

Active Directory Certificate Services

CA types

Certificate revocation lists

Certificate Services role services

Certificate templates

Certificate autoenrollment and renewal

CA management

Active Directory Certificate Services is the infrastructure service that most administrators don't pay much attention to until they need to perform a task, such as to generate certificates for an internal application or to support a specific type of server deployment.

You can install Certificate Services on both Server Core and Server with a GUI system, with Server Core being

the more secure option because of its reduced attack surface. Because Certificate Services is an important security role, you should consider deploying it on a computer separate from other roles. You should also limit access to the server hosting Certificate Services so that only those directly responsible for performing administration tasks have access.

CA TYPES

Windows Server Certificate Authorities (CAs) come in four basic types: enterprise root, enterprise subordinate, standalone root, and standalone subordinate. The type that you deploy depends on your certificate needs. Unless you are performing a new deployment, the type of deployment also likely depends on the particular CA hierarchy that you want to use.

Before you can understand what type of CA you need to deploy, you need to understand how CA hierarchies work. The CA at the top of a hierarchy is known as the Root CA. A Root CA can have any number of subordinate CAs. Subordinate CAs can, in turn, function as the parent CAs of child CAs, which are also called subordinate CAs.

Each CA has a special certificate known as the CA certificate. The Root CA signs its own certificate. A subordinate CA has its CA certificate signed by a parent

CA. For example:

- A Root CA named *APEX_CA* signs its own CA certificate.
- *APEX_CA* signs the CA certificate of a subordinate CA named *CHILD_CA*.
- *CHILD_CA* signs the CA certificate of a subordinate CA named *GRANDCHILD_CA*.

The way that certificate trust works is that if a client trusts the Root CA's CA certificate, it trusts all certificates issued by CAs that are directly or indirectly subordinates of that CA. For example, if a client trusts the *ROOT_CA* CA certificate, it automatically trusts the certificates issued by *GRANDCHILD_CA*.

You can use the pkiview console (*pkiview.msc*) to view your current organization's Certificate Services configuration. Figure 12-1 shows an enterprise PKI that has a root CA and a subordinate CA.

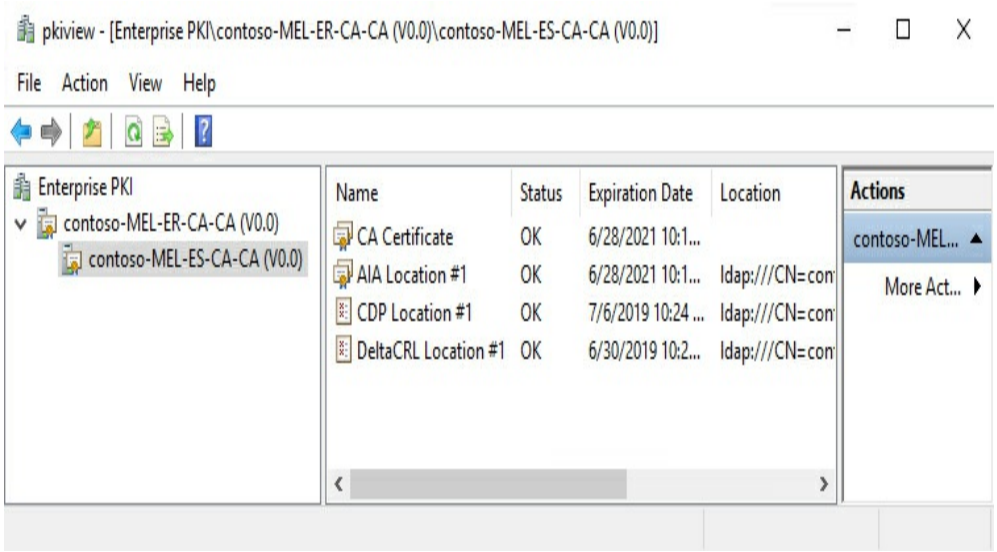


Figure 12-1 PKIview

Another concept that it is important to understand in the context of CA hierarchies is certificate revocation. A parent CA can revoke a child CA's CA certificate. When this happens, all certificates issued by the child CA are also invalidated. For example, if the administrator of *APEX_CA* from the earlier example revoked the CA certificate of *CHILD_CA*, all the certificates issued by CHILD_CA would no longer function. This would include the CA certificate of *GRANDCHILD_CA*, which, in turn, would invalidate all the certificates issued by *GRANDCHILD_CA*. Needless to say, you only revoke the CA certificate of a subordinate CA in dire circumstances, such as when the child CA has been compromised by malware or hackers. If a child CA has been compromised, it might have issued certificates that you don't know about, and the most secure response is to revoke all certificates issued from the CA and to start over from scratch.

Some other things to keep in mind before we delve more into the types of CAs that you can configure with Windows Server 2019 are:

- You can't convert a root CA to a subordinate CA.
- You can't convert a subordinate CA to function as a root CA; however, a subordinate CA can function as the parent CA of another subordinate CA.

Technet24

- A subordinate CA's CA certificate cannot have a lifetime that extends beyond the current lifetime of the parent CA's certificate.

> **Inside OUT**
>
> **Remember to renew CA certificates**
>
> Remember, a CA's server certificate needs to be renewed. A CA can't issue certificates that have a lifespan beyond that of its CA certificate. If a CA's certificate has expired, so have all the certificates issued by the CA.

Enterprise CA

An Enterprise CA is fully integrated with Active Directory. This has the following benefits:

- The Enterprise CA is automatically trusted by any computer that is a member of the same Active Directory forest.
- All certificates published by the CA are automatically trusted by members of the same forest.
- Certificate requests can be performed automatically and don't require the use of a request file.
- Certificates requests can be approved automatically, and certificates are automatically distributed and installed on domain clients.
- Certificate revocation information is published to Active Directory.
- Enterprise CAs can use certificate templates.
- Enterprise CAs can issue certificates to devices that have domain membership and also can issue certificates to devices that are not members of an Active Directory domain.
- An Enterprise subordinate CA can have a standalone CA as a parent CA.
- A standalone subordinate CA can have an Enterprise CA as a parent CA.

One of the primary drawbacks of an Enterprise CA is that it always needs to be online. If an Enterprise CA goes offline, certificate services in the domain don't function properly. Temporary shutdowns due to maintenance are fine, but shutting down or having an enterprise CA offline for a couple of weeks is likely to lead to headaches. Enterprise CAs must also always be installed on computers that are members of an Active Directory domain.

Enterprise Root

An Enterprise Root CA is an enterprise CA that signs its own signing certificate. As mentioned earlier, an Enterprise Root CA can only be installed on a computer that is a domain member, and that computer should always be online. Because the Enterprise Root CA publishes certificate information to all domains in a forest, you need to configure it using an account that is a member of the Enterprise Admins group.

> **Inside OUT**
>
> **Single CA Deployments**
>
> Many smaller organizations, such as those with fewer than 300 users, have a single Enterprise Root CA as their only CA. While it's not the most secure way you can do things, a single Enterprise Root CA is a pragmatic solution for organizations that have basic certificate services requirements. The best practice is to use an Offline Root CA, which is discussed later in this chapter.

After you've installed the binary files for a CA, you are

then able to configure it. You do this by running the AD CS Configuration Wizard. You can run this wizard using the signed-in user's credentials, or by using separate credentials, which you specify on the Credentials page of the wizard shown in Figure 12-2. In the next few pages, we'll go through the process of installing an Enterprise Root CA. The process of installing an Enterprise Root CA is similar to the process of installing other CA types.

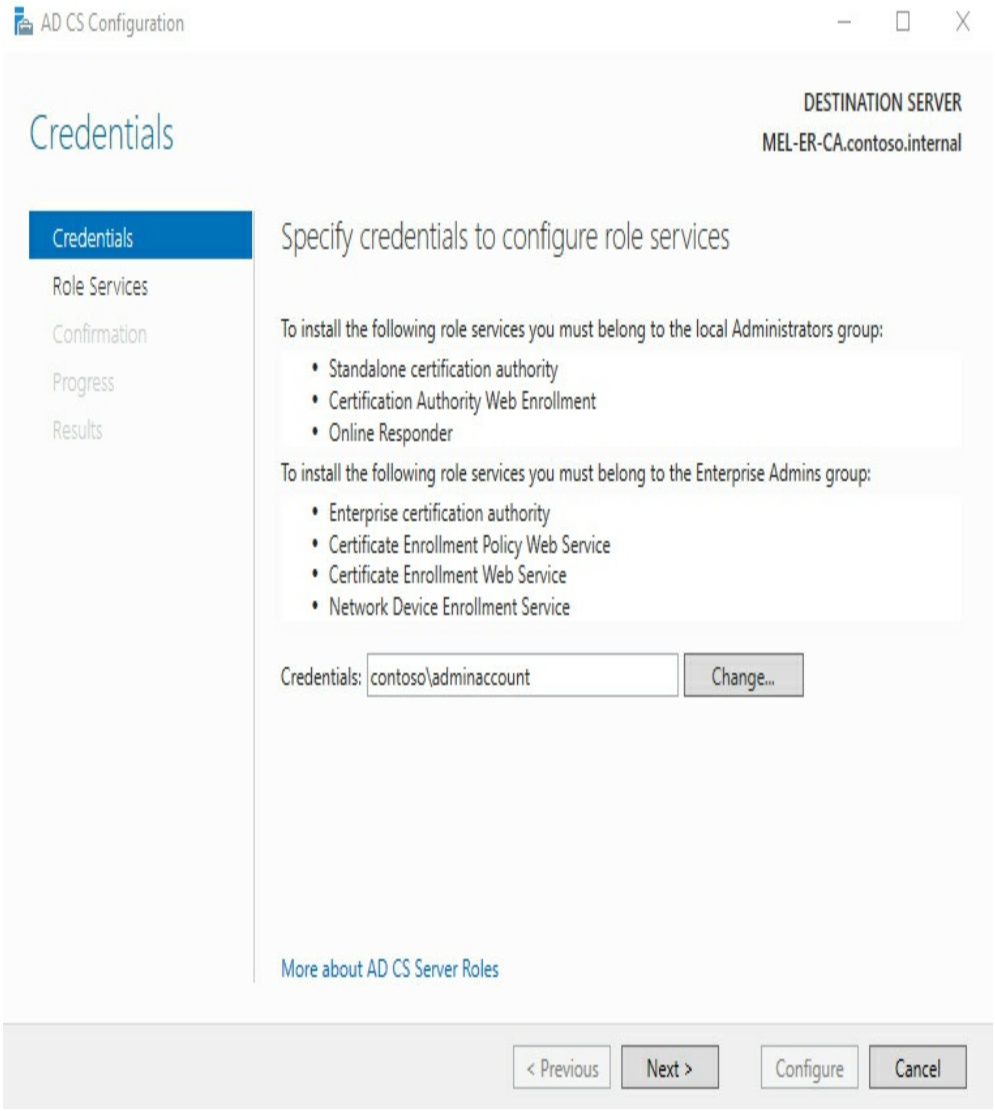


Figure 12-2 Credentials configuration

The next step in the wizard involves specifying which roles to configure. While it is possible to configure multiple roles at once, it's usually better to get the Certification Authority role functioning properly before adding and configuring additional role services, such as Certification Authority Web Enrollment. The other role

services are discussed later in this chapter. Figure 12-3 shows the Certification Authority role selected.

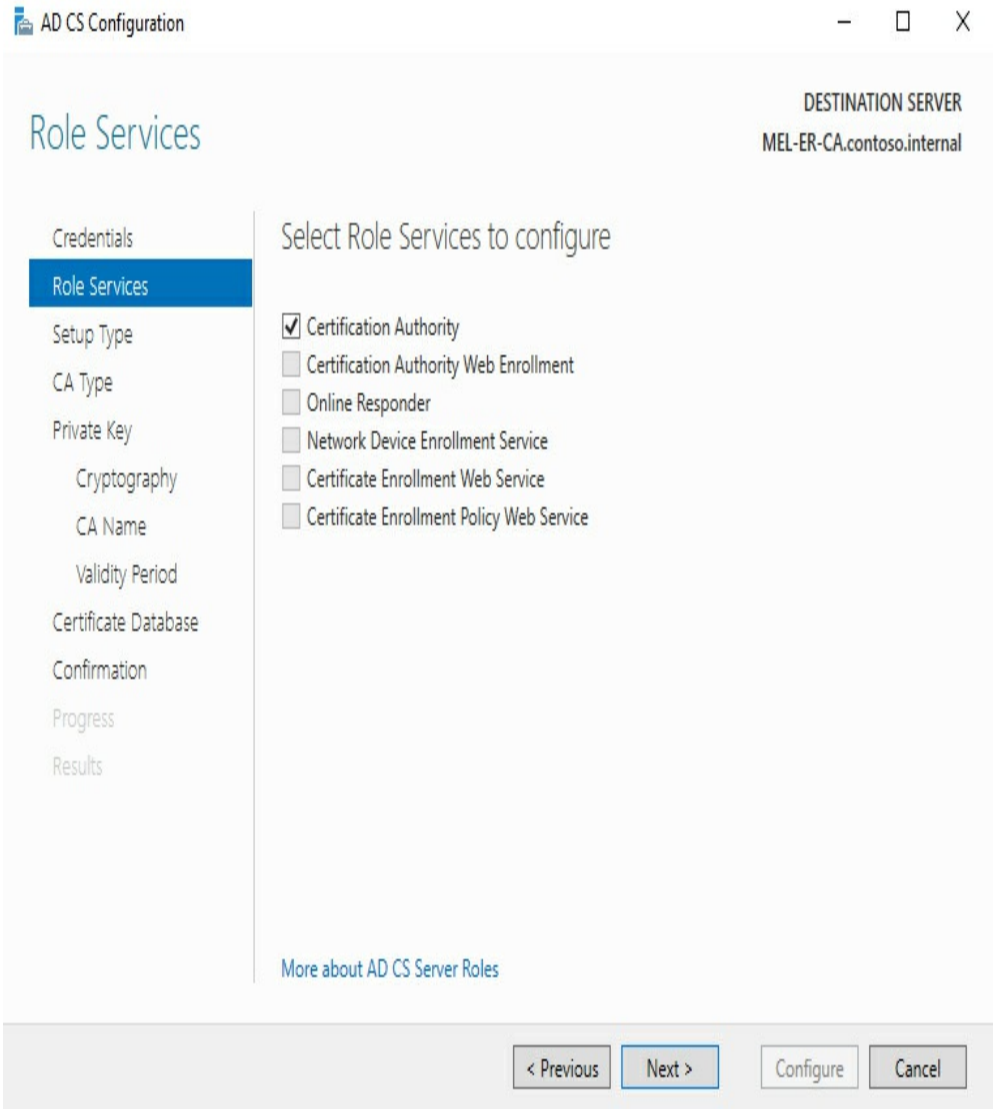


Figure 12-3 Certification Authority role

If the computer is a member of an AD DS domain, you get the option to install an Enterprise CA or a Standalone CA, as shown in Figure 12-4. If the computer

is not a member of an AD DS domain, you can't select the Enterprise CA option. If you're sure that the computer is a member of the domain, but you can't select the Enterprise CA option, make sure that you are using an account that has Enterprise Admin privileges.

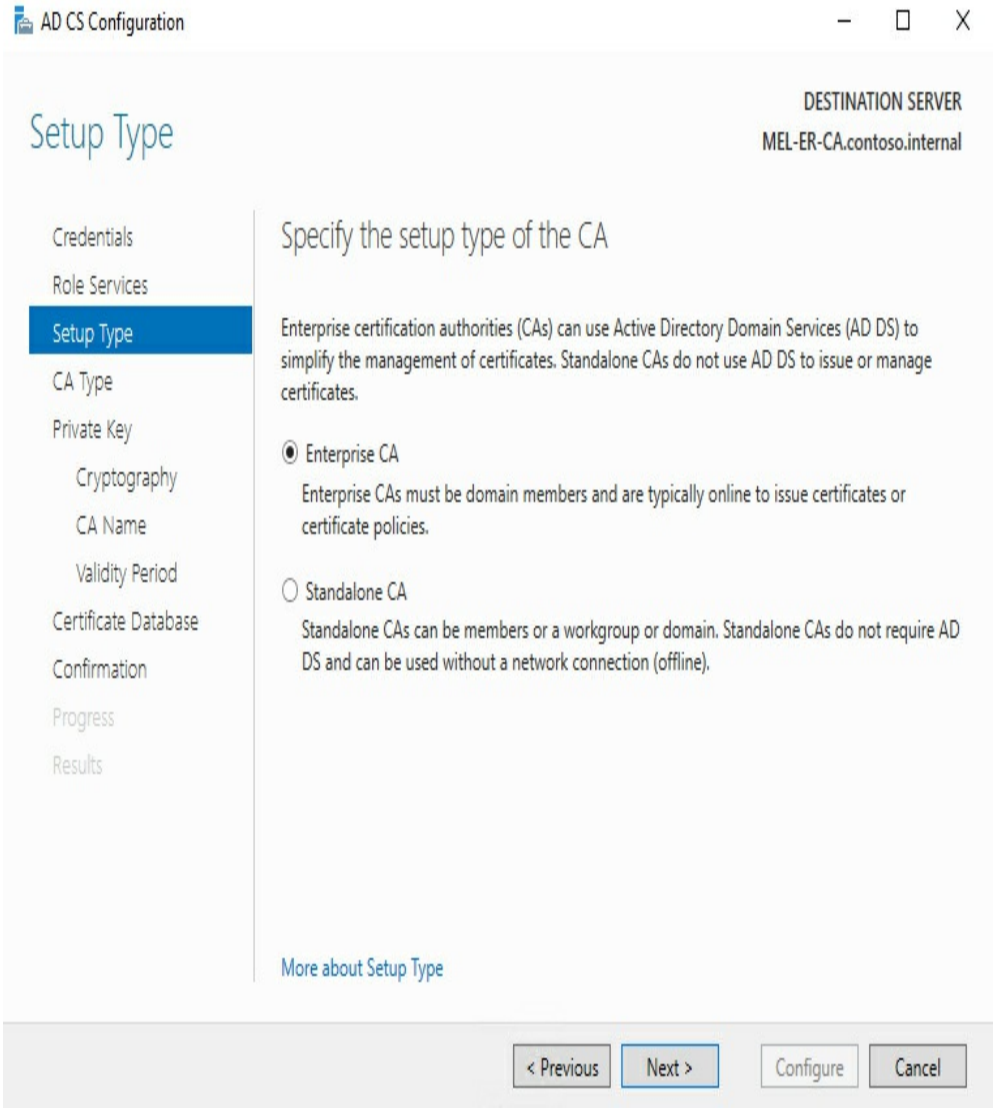


Figure 12-4 Enterprise CA

After you've chosen between Enterprise CA and Standalone CA, you need to choose between Root CA and Subordinate CA, as shown in Figure 12-5. Although it is possible to do, Microsoft doesn't recommend having multiple enterprise root CAs in a forest. Unless you are deploying a single CA, the recommendation is to have no enterprise root CAs and to use a combination of an offline standalone root CA and enterprise subordinate CAs as this provides the most secure deployment type.

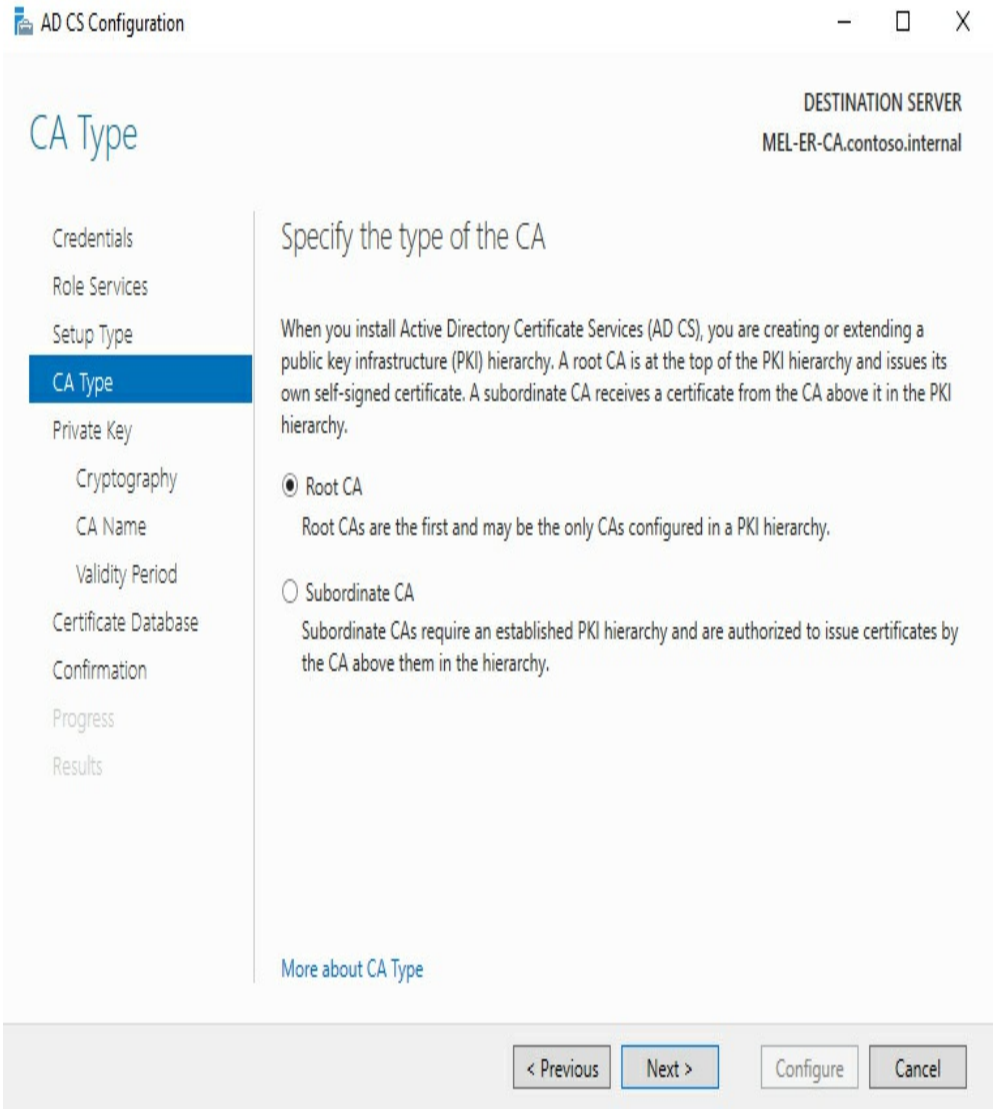


Figure 12-5 CA Type

On the Private Key page, shown in Figure 12-6, specify whether you want to create a new private key or use an existing private key. If you're deploying a new CA, create a new private key. If you are replacing a failed CA or migrating from an earlier version of Windows Server to a new Windows Server 2019 CA, use an existing private

key. The trick with CA migration is to ensure that the destination server has the same name as the source server.

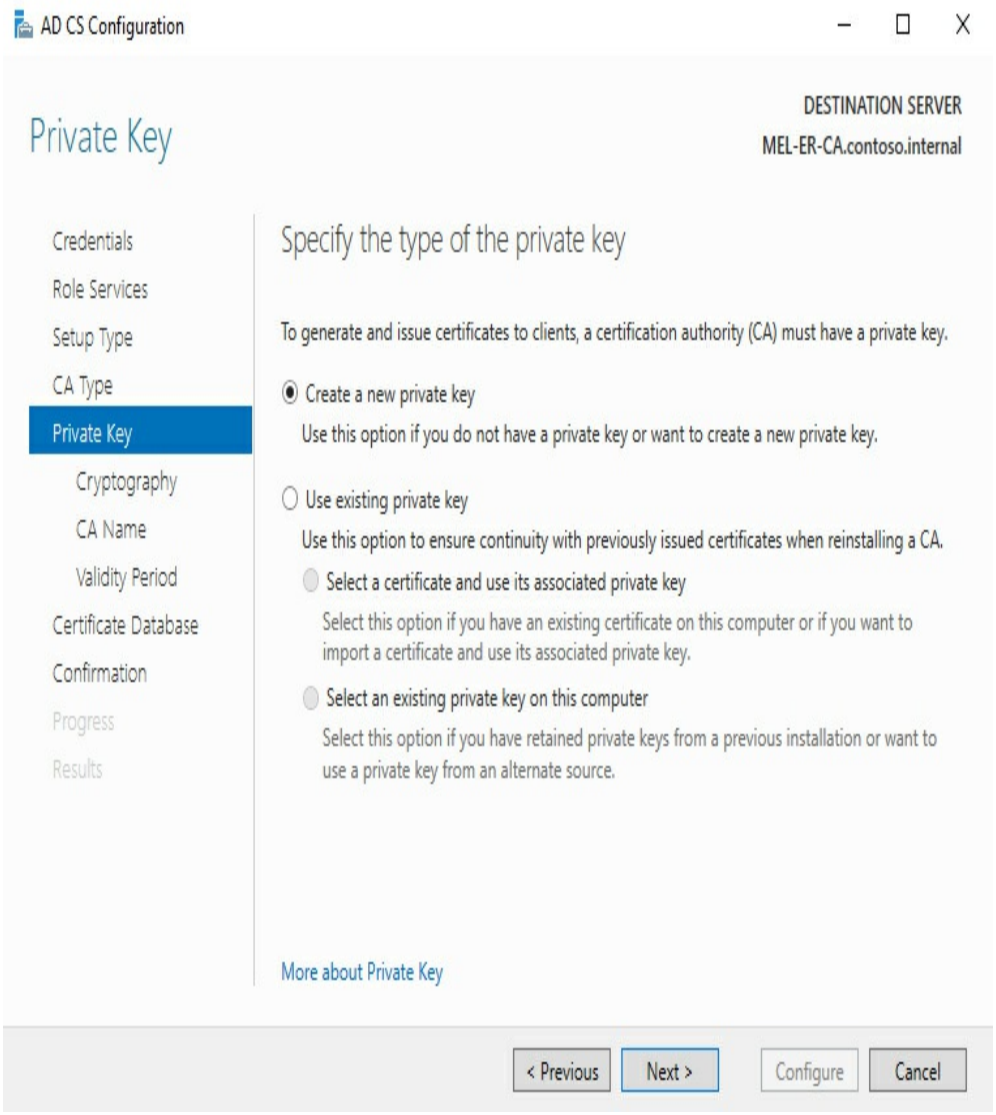


Figure 12-6 The Private Key page

More Info
CA migration

You can find out more about CAmigration at https://techcommunity.microsoft.com/t5/ITOps-Talk-Blog/Step-By-Step-Migrating-The-Active-Directory-Certificate-Service/ba-p/697674.

If you choose to create a new private key, you need to specify the cryptographic options for that key. In most circumstances, the default options, shown in Figure 12-7, are fine. In some cases, however, you might want to configure stronger cryptographic options. The important thing to remember is that the more advanced and stronger the cryptography you use, the more you limit yourself as to which clients can interact with the key. Table 12-1 provides a list of available cryptographic options.

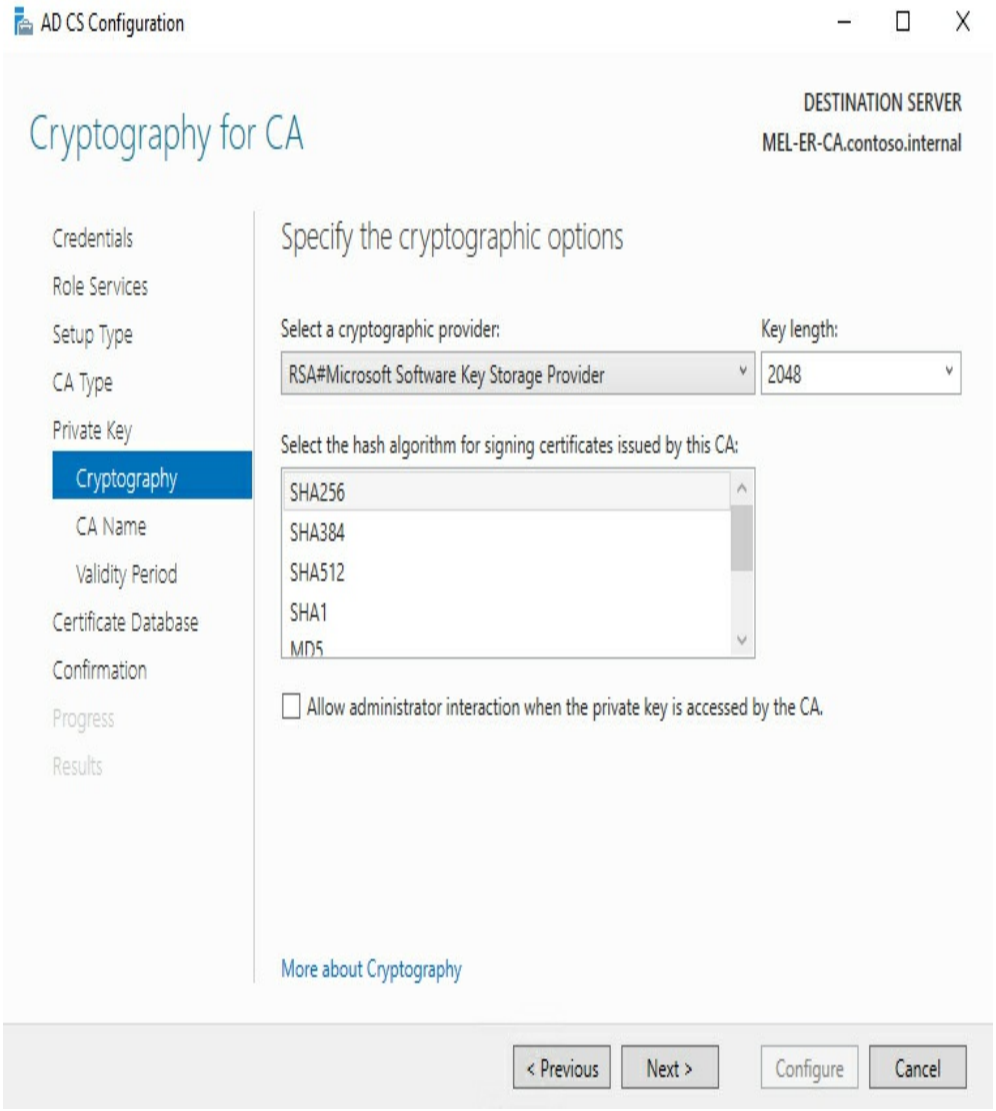


Figure 12-7 Cryptography For CA

Table 12-1 CA Private Key cryptography options

Cryptographic providers	Available hash algorithms	Available key lengths
Microsoft Base Smart Card Crypto Provider	SHA1	1024

	MD2	2048
	MD4	4096
	MD5	
Microsoft Enhanced Cryptographic Provider v1.0	SHA1	512
	MD2	1024
	MD4	2048
	MD5	4096
ECDSA_P256#Microsoft Smart Card Key Storage Provider	SHA256	256
	SHA384	
	SHA512	
	SHA1	
ECDSA_P521#Microsoft Smart Card Key Storage Provider	SHA256	521
	SHA384	
	SHA512	
	SHA1	

RSA#Microsoft Software Key Storage Provider (Default)	SHA256	512
	SHA384	1024
	SHA512	2048
	SHA1	4096
	MD5	
	MD4	
	MD2	
Microsoft Base Cryptographic Provider v1.0	Sha1	512
	Md2	1024
	Md4	2048
	Md5	4096
ECDSA_P521#Microsoft Software Key Storage Provider	SHA256	521
	SHA384	
	SHA512	
	SHA1	

ECDSA_P256#Microsoft Software Key Storage Provider	SHA256 SHA384 SHA512 SHA1	256
Microsoft Strong Cryptographic Provider	SHA1 MD2 MD4 MD5	512 1024 2048 4096
ECDSA_P384#Microsoft Software Key Storage Provider	SHA256 SHA384 SHA512 SHA1	384
Microsoft Base DSS Cryptographic Provider	SHA1	512 1024

RSA#Microsoft Smart Card Key Storage Provider	SHA256	1024
	SHA384	2048
	SHA512	4096
	SHA1	
	MD5	
	MD4	
	MD2	
DSA#Microsoft Software Key Storage Provider	SHA1	512
		1024
		2048
ECDSA_P384#Microsoft Smart Card Key Storage Provider	SHA256	384
	SHA384	
	SHA512	
	SHA1	

The CA name window, shown in Figure 12-8, is

automatically populated based on the computer's identity information. This name is important because it is present on each certificate issued by the CA. The distinguished name suffix is automatically generated. Unless you are issuing certificates to outside third parties, the default names provided here are likely to suffice.

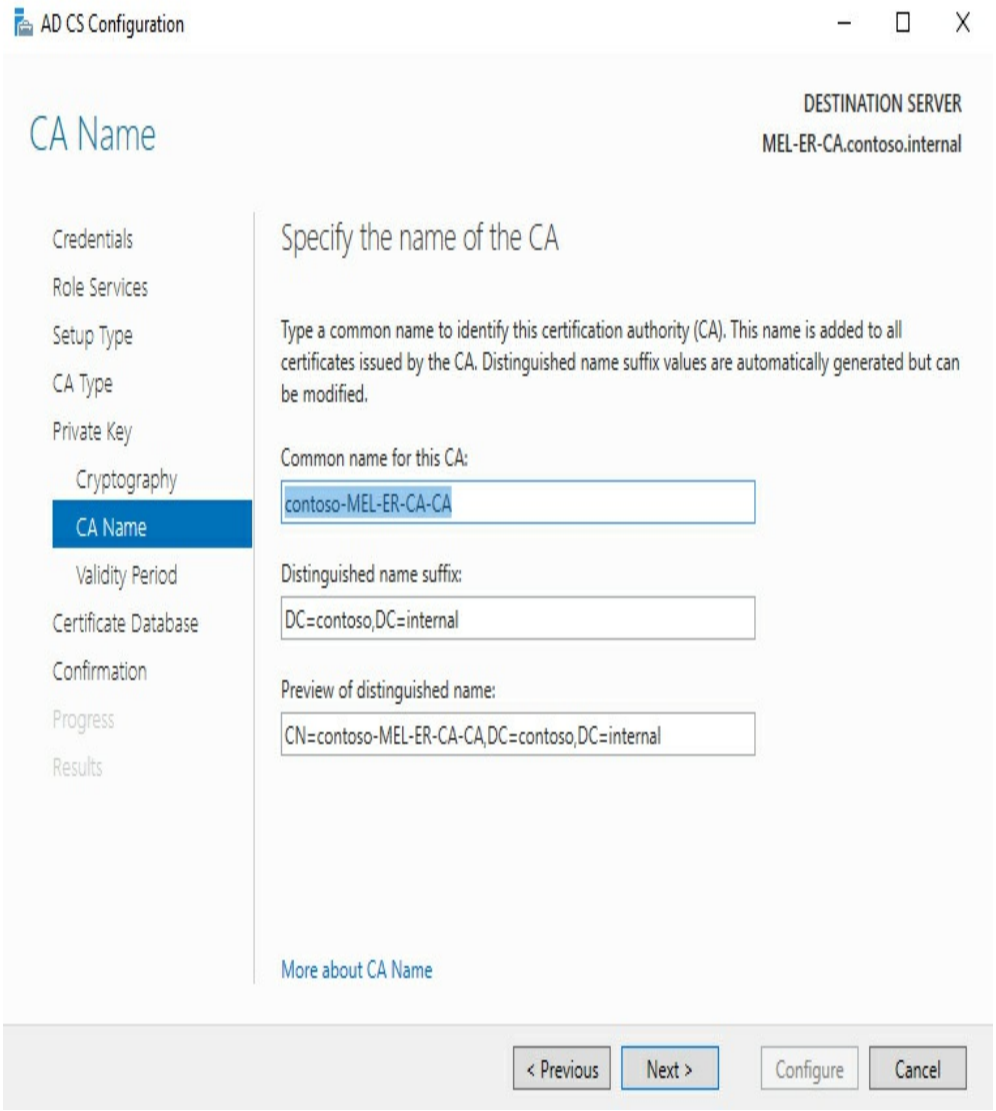


Figure 12-8 CA Name

On the Validity Period page, shown in Figure 12-9, specify the validity of the CA's certificate. When configuring this for a root CA, remember that no certificate issued by the CA or a subordinate CA can have a lifetime that exceeds the value of the CA's certificate. This is an important consideration when it comes to subordinate CAs. For example, let's say you are installing a subordinate CA and you request a five-year CA certificate from a root CA. If the root CA issuing the subordinate CA certificate has only four years of validity left on its CA certificate, the expiration date of the subordinate CA certificate is four years, not five years.

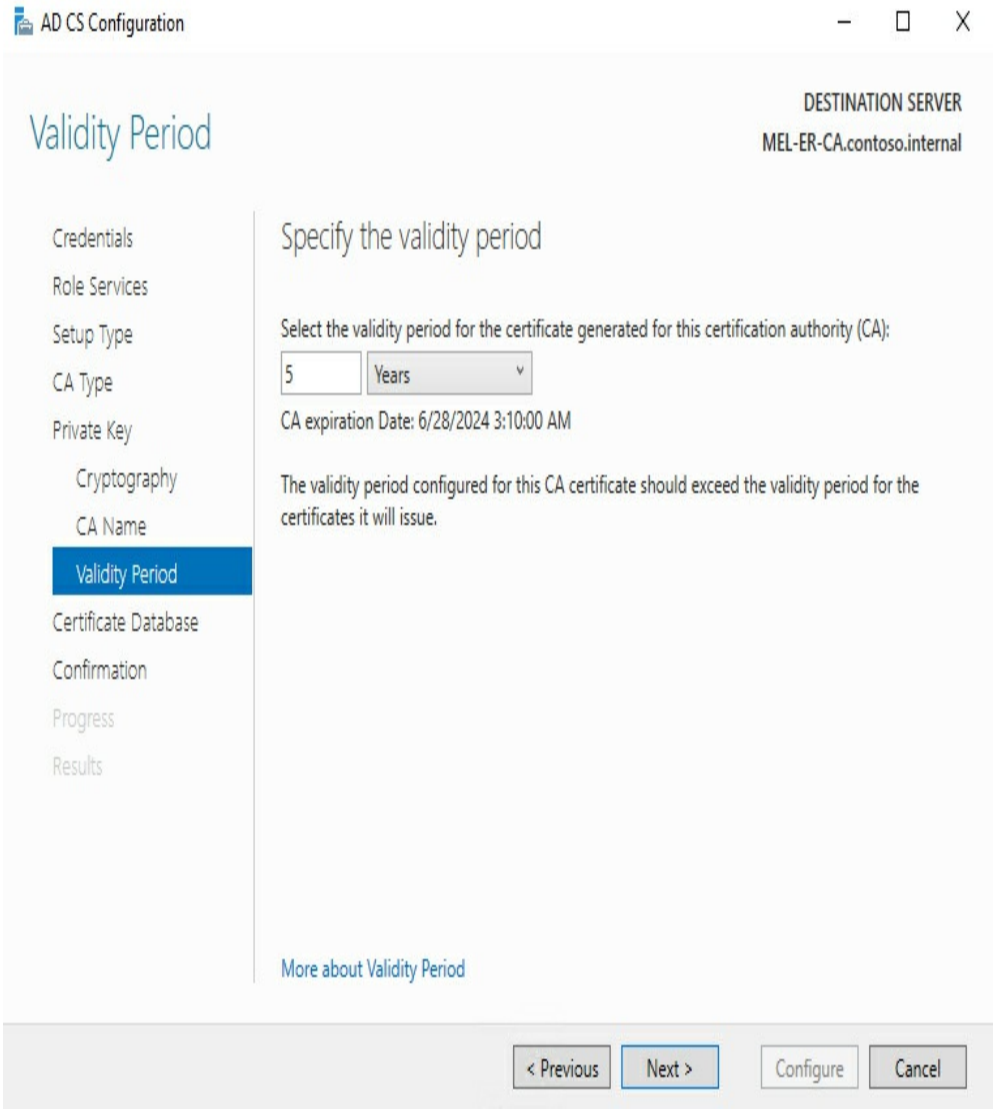


Figure 12-9 Validity Period

The final configuration step for an Enterprise Root CA is the location of the certificate database and the database log, as shown in Figure 12-10. These locations are important primarily from the perspective of backing up and recovering the database. If you enable key archiving, which is covered later in this chapter, the private keys of

Technet24

certificates issued by the CA are also written to this database in encrypted form.

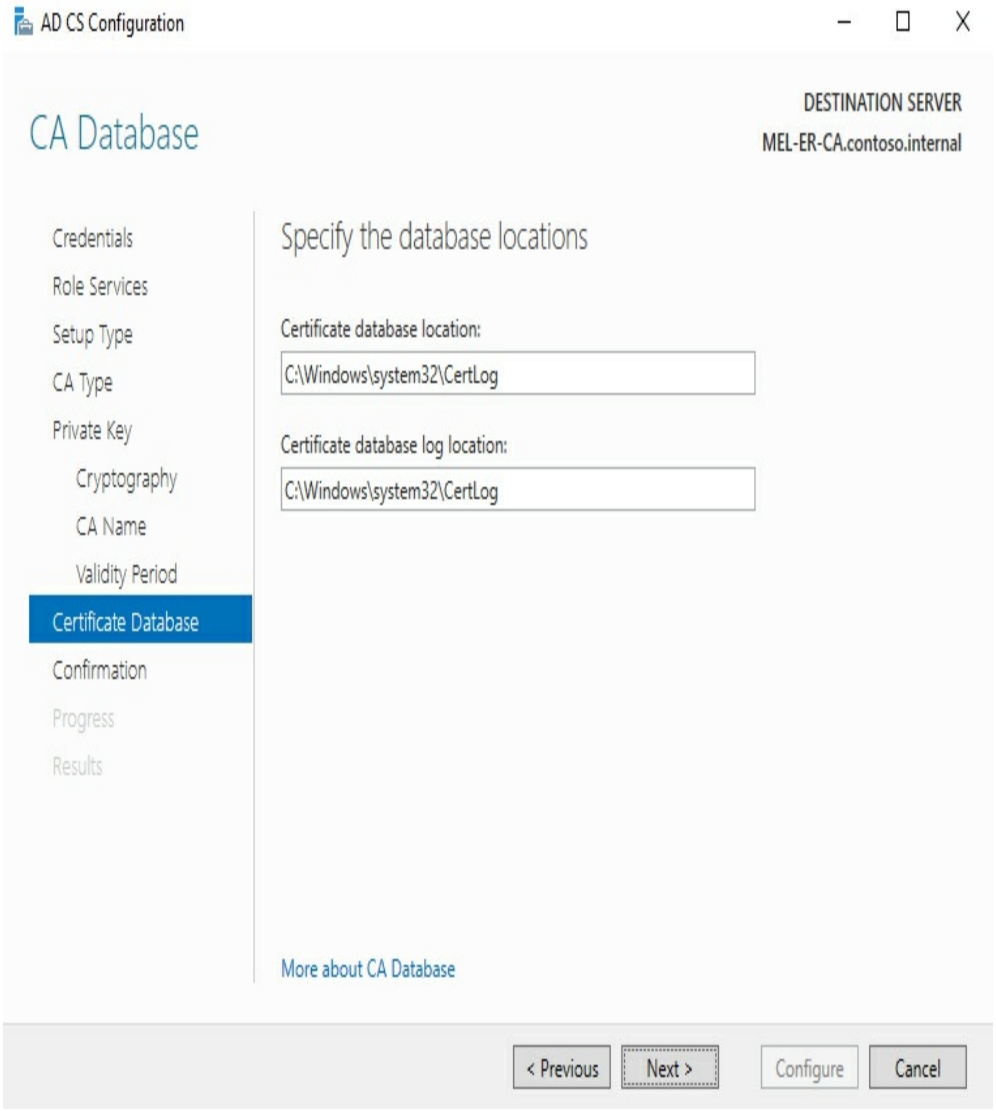


Figure 12-10 CA Database

Enterprise subordinate

An enterprise subordinate CA is integrated into Active Directory but has another CA as the apex of the

certificate trust chain. Similar to enterprise root CAs, enterprise subordinate CAs must remain online and must be a member of an Active Directory domain. Domain membership means computers that are members of the same forest as the enterprise subordinate automatically trust certificates issued by the CA. As is the case with enterprise root CAs, certificate revocation data is also published to Active Directory. Enterprise subordinate CAs can use certificate templates independently of where their CA certificate comes from.

You can acquire a signing certificate for an enterprise subordinate CA from an enterprise CA or from a standalone CA. Microsoft recommends using enterprise subordinate CAs with an offline root CA. You'll learn about offline root CAs later in this chapter.

If you are using a standalone CA, the installation wizard automatically generates a certificate request file and places it in the root directory on volume *C:* when the wizard completes. You can then submit this file to the CA to generate the subordinate CA's certificate. Until you install the subordinate CA certificate, the CA is not operational.

To install the CA certificate issued by the parent CA, perform the following steps:

1. Open the Certification Authority console.

Technet24

2. Select the CA. In the Action menu, click All Tasks and then click Install CA Certificate.
3. In the Install CA Certificate dialog box, locate and open the certificate issued by the parent CA.
4. Ensure the CA is selected. In the Action menu, click All Tasks and then click Start Service.

Standalone CAs

Standalone CAs differ from enterprise CAs primarily in that enterprise CAs are highly integrated into Active Directory and the other is not. Additional differences between standalone CAs and enterprise CAs include:

- An enterprise CA must be installed on a domain member. You can install a standalone CA on a computer that isn't domain joined or on a computer that is domain joined.
- Standalone CA certificates are not automatically published to Active Directory.
- You can't use certificate templates with standalone CAs. The properties of the certificate are dependent on the certificate request.
- Standalone CAs cannot authenticate or use the identity of Active Directory users as part of the certificate request process even when installed on computers that are members of the domain.
- By default, all certificate requests to a standalone CA must be approved before a certificate is issued. With an enterprise CA, you can configure whether a certificate is automatically approved by configuring template properties.

While you can use standalone CAs to issue certificates to domain clients, they are not directly integrated into Active Directory in the way that an enterprise CA is. This means that domain members don't automatically trust

certificates issued by a standalone CA and aren't automatically able to determine certificate revocation information by performing a query against Active Directory.

While you can deploy a standalone CA on a computer that is a domain member, standalone CAs are almost always deployed on computers that are not members of an Active Directory domain. Benefits of standalone CAs include:

- Because the CA is not a member of the domain, it can be deployed on a perimeter network or in Azure.
- The CA can be shut down for extended periods.
- Members of Active Directory security groups, such as members of the Domain Admins group, don't have access to the CA if it is deployed on a computer that does not have domain membership.

You can configure a client to trust a CA by importing the CA's certificate into the client's Trusted Root Certification Authority store. To do this, perform the following steps:

1. Open the Certificates Snap-In of the Microsoft Management Console and set the snap-in to use the Computer Account, as shown in Figure 12-11.

Technet24

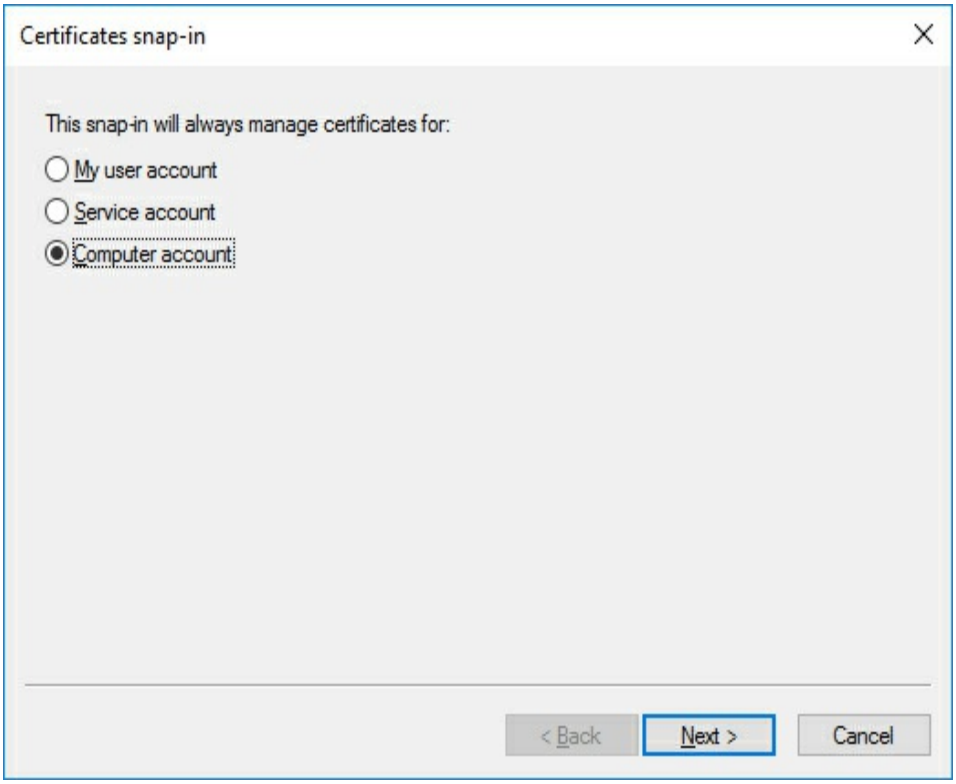


Figure 12-11 Selecting the Computer Account snap-in

2. Select the Certificates node under the Trusted Root Certification Authorities node, as shown in Figure 12-12. On the Action menu, click All Tasks and then click Import.

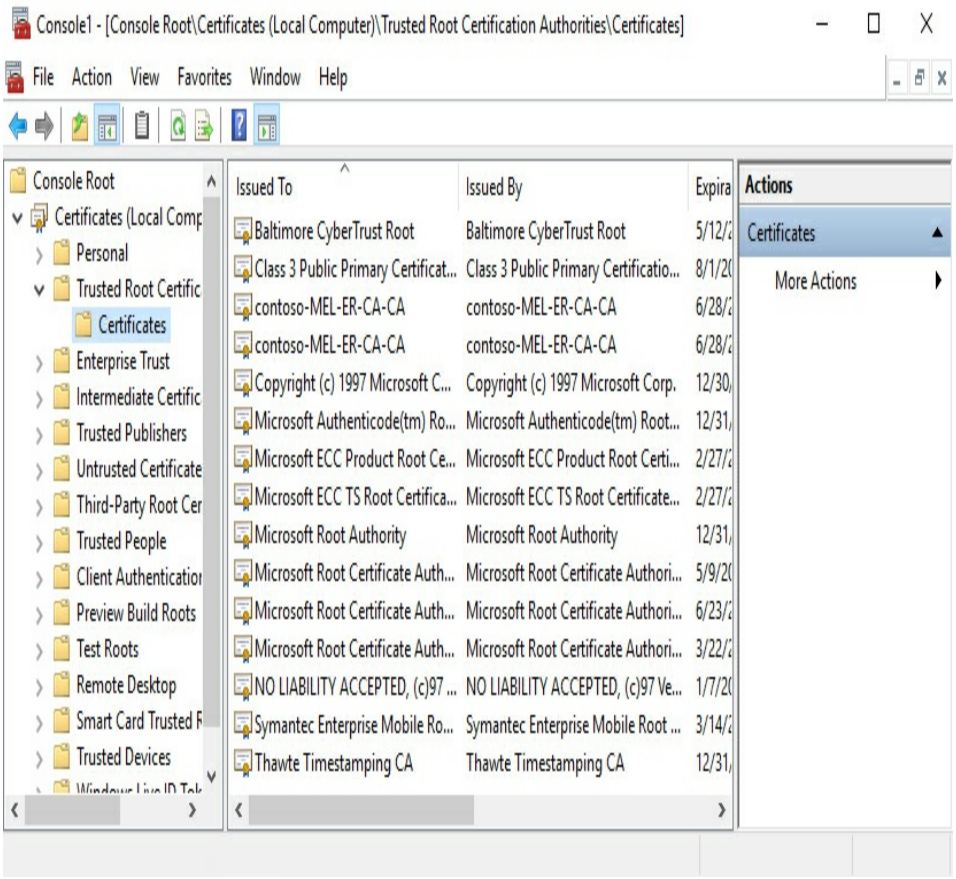


Figure 12-12 Trusted Root CAs

3. After you import the CA certificate, the client computer trusts all certificates issued by that CA. It also trusts all certificates issued by CAs that are subordinate to that CA in the trust chain.

If you want to have all clients in a domain trust the certificates issued by a standalone root CA, you can use the *certutil.exe* command to publish the root CA's certificate to Active Directory. To do this, you need to run the *certutil.exe* command. Do this using an account that has Domain Admin privileges and substitute the root CA certificate file name for the file name *SA_ROOT_CA.crt* used in this example command:

Technet24

Click here to view code image

```
Certutil.exe -dspublish -f SA_ROOT-CA.crt
```

Inside OUT

Deploy root certificates

You can use Microsoft Intune or System Center Configuration Manager to deploy root certificates to domain-joined clients, nondomain–joined clients, and managed mobile devices.

Standalone root

Standalone root CAs sign their own certificates, but they aren't directly integrated into AD DS. A standalone root CA can be offline, but this causes problems if an online certificate revocation list distribution point is not present. Standalone root CAs often have the Certification Authority Web Enrollment service installed as a method of allowing people to submit certificate requests to the CA. After a request is submitted to the CA, you can approve or deny the request using the Certification Authority console. If the request is in the form of a *.req* file, you can use the *certreq.exe* utility to manage the request. You'll learn more about managing certificate requests later in this chapter.

Standalone subordinate

A standalone subordinate CA can be the child CA of an Enterprise CA or a standalone CA. Standalone

subordinates are often deployed on perimeter networks when it is necessary to issue certificates to people outside the organization. Don't use an Enterprise CA in this scenario because enterprise CAs can only be installed on domain–joined computers, and you should avoid deploying domain–joined computers on perimeter networks. As is the case with standalone root CAs, standalone subordinate CAs often have the Certification Authority Web Enrollment service installed as a method of allowing people to submit certificate requests to the CA and as a method of retrieving certificates after they are issued.

Offline root

An offline root CA is a CA that is secured by being turned off and is only brought online when specific tasks occur, such as revoking an existing signing certificate, issuing a new signing certificate, or publishing a new certificate revocation list (CRL). Because the CA is offline, it is much less likely to be compromised.

The trick with deploying an offline root CA is that you need to configure the Certificate Revocation List and Authority Information Access distribution points, so they are hosted on a computer that remains online when the offline root CA is offline. You also need to export the CA certificate and place it in a location

accessible to clients performing revocation checks. You should renew the CA certificate for the offline root CA after updating the CRL and Authority Information Access (AIA) information because it is included with the offline root CA CA certificate. You'll learn more about configuring CRLs and AIA distribution points later in this chapter.

More Info

Offline Root CAs

You can find out more about offline root CAs at http://social.technet.microsoft.com/wiki/contents/articles/2900.offline-root-certification-authority-ca.aspx.

CERTIFICATE REVOCATION LISTS

A certificate revocation list (CRL) is a list of certificates a CA has issued that are no longer considered valid. A delta CRL is a list of all certificates that have been revoked since the publication of the last CRL. By default, a CA publishes a new CRL every week, and a delta CRL is published every day. Both the CRL and the delta CRL are stored on a CRL distribution point, also known as a CDP.

CRL distribution points

A CA can have multiple CDPs. CRL and CDP information is included with a certificate when the certificate is issued. Clients perform a check of

certificate validity each time they encounter a new certificate and perform additional checks based on the CRL and delta CRL publication intervals. Clients also check CDP locations in the order that they are configured on the CA.

You configure the CRL distribution point on the Extensions tab of the CA's properties, as shown in Figure 12-13. If you are publishing certificates that need to be used by people outside your organization, ensure that you have a CRL hosted in a publicly available location, such as on a web server on your perimeter network or in Microsoft Azure. Remember that any changes you make to CDP settings are only applied to certificates issued after the changes are made. If you make a change and then remove existing CDPs, clients encountering existing certificates can't perform revocation checks, and the certificates are deemed invalid.

Technet24

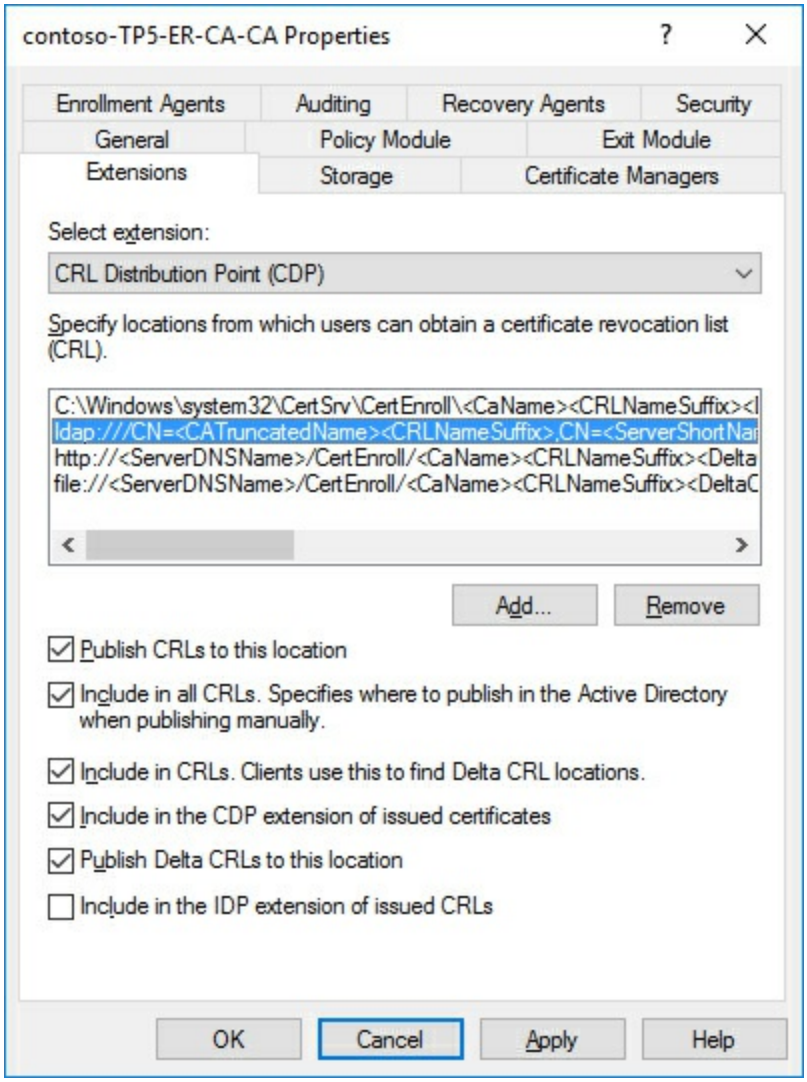


Figure 12-13 CRL DPs

Authority Information Access

The Authority Information Access (AIA) point settings specify where a client can retrieve the CA certificate of the CA that issued the certificate that the client has encountered. As with the CRL DP information, changes to the AIA information are only present in certificates

issued subsequent to the change. This means that you need to keep existing AIA points in place until currently issued certificates either expire or are renewed. You configure AIA extensions on the Extensions tab of the CA properties, as shown in Figure 12-14. The AIA extensions are visible when you set the Select Extension drop-down menu to Authority Information Access (AIA).

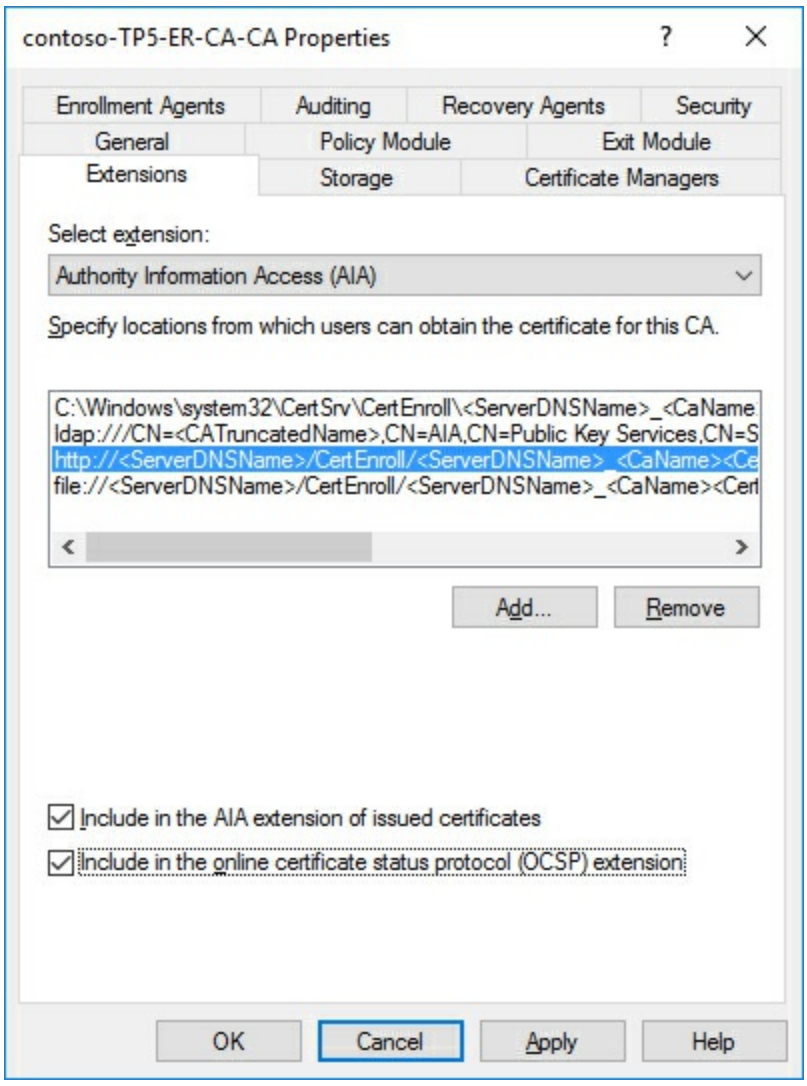


Figure 12-14 Selecting the Authority Information Access extension

Revoking a certificate

Revoking a certificate makes it invalid, and doing so also involves several steps. The first step is to revoke the certificate itself, something you do from the CA Management console. The second step is to update the CRL or delta CRL and then publish the updates to the appropriate location. Revoking a certificate requires that you know the identity of the certificate you are revoking.

To revoke a certificate, perform the following steps:

1. Locate the certificate in the list of issued certificates on the CA from which you issued the certificate. Certificate serial numbers are in hexadecimal but are listed sequentially in the console. Figure 12-15 shows a list of issued certificates.

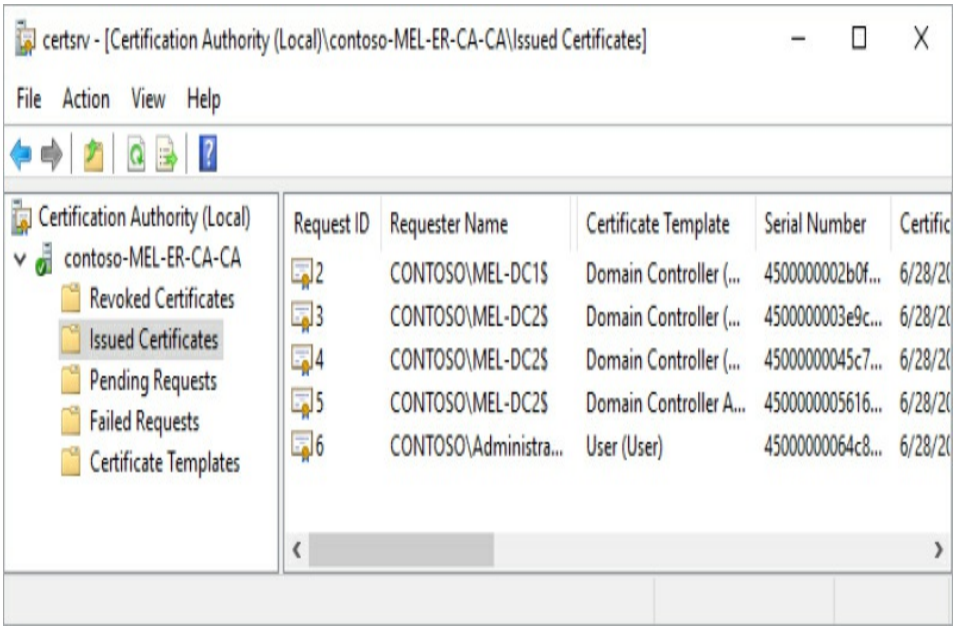


Figure 12-15 Issued certificates

2. On the Action menu, click All Tasks and then click Revoke Certificate.
3. In the Certificate Revocation dialog box, shown in Figure 12-16, select the Reason Code for the certificate revocation when you want the certificate revocation to take place (with the default being immediately), and click Yes. You are asked to confirm the revocation. The revoked certificate is then listed in the Revoked Certificates node in the Certification Authority console.

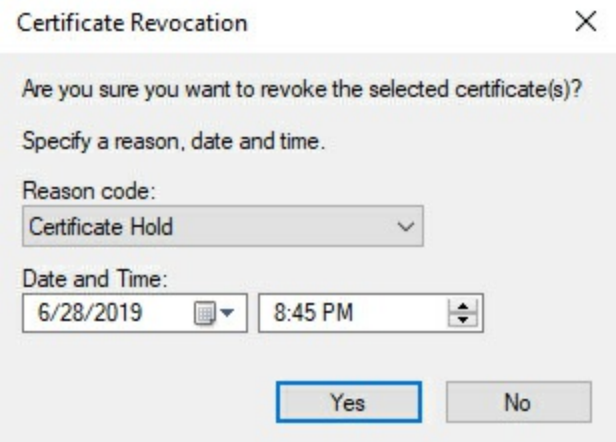


Figure 12-16 The Certificate Revocation dialog box

When you revoke a certificate, you need to specify a Reason Code. The following reason codes are available:

- Unspecified
- Key Compromise
- CA Compromise
- Change Of Affiliation
- Superseded
- Cease Of Operation
- Certificate Hold

Technet24

The reason you specify is important if your organization is required to document why a specific certificate has been revoked. All reasons except Certificate Hold are irreversible. If you choose the Certificate Hold reason, you can undo the certificate revocation at a later point in time. Depending on the reason for performing the revocation and its severity, such as finding a subordinate CA compromised by remote access Trojan malware, you might want to publish a CRL or delta CRL.

Publishing CRLs and delta CRLs

CRLs and delta CRLs are published automatically by CA. A new CRL is published by default once a week, and a new delta CRL is published once a day. You can configure the publication schedule using the Revoked Certificates Properties dialog box, shown in Figure 12-17. To access the dialog box, click the Revoked Certificates node of the Certification Authority console and then select Properties from the Action menu.

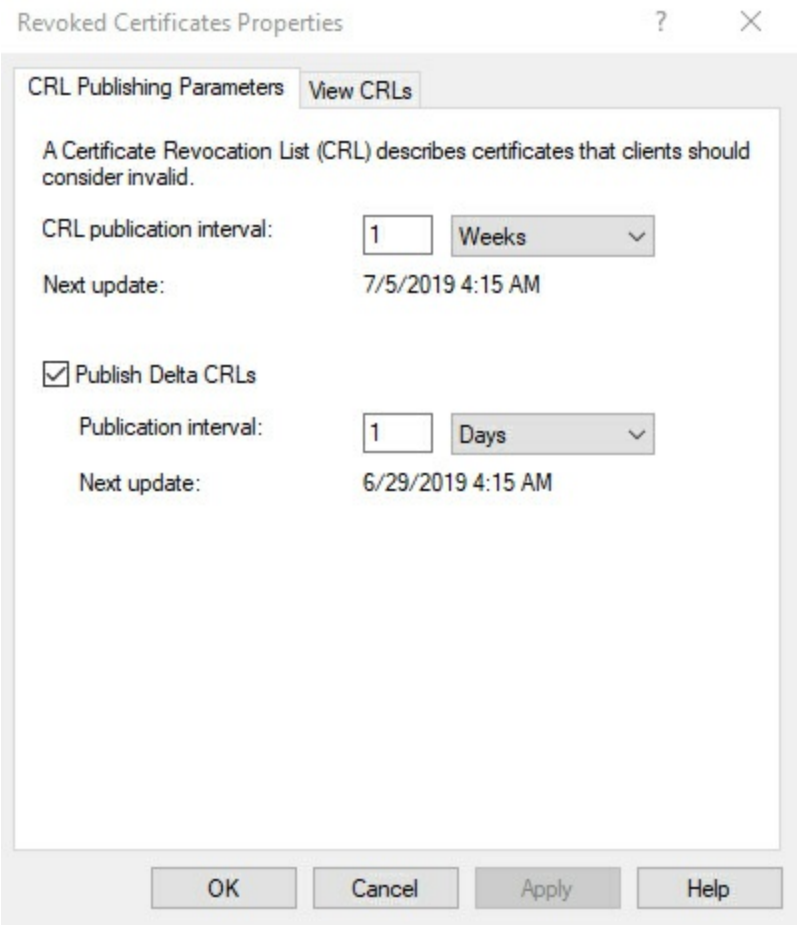


Figure 12-17 CRL Publishing Parameters

For most organizations, the default schedule is fine. If an organization issues certificates where revocation needs to propagate quickly, a shorter delta CRL publication interval would be more appropriate. This is because clients that have already accepted the certificate as valid have also learned about the CRL and delta CRL publication frequency. If you revoke a certificate, and you have a delta CRL publication frequency of once a week and a CRL publication frequency of once a month,

Technet24

it could take up to a week before all clients encountering that certificate recognize it as invalid.

You can view the last time a CRL was published on the View CRLs tab of the Revoked Certificates Properties dialog box, shown in Figure 12-18.

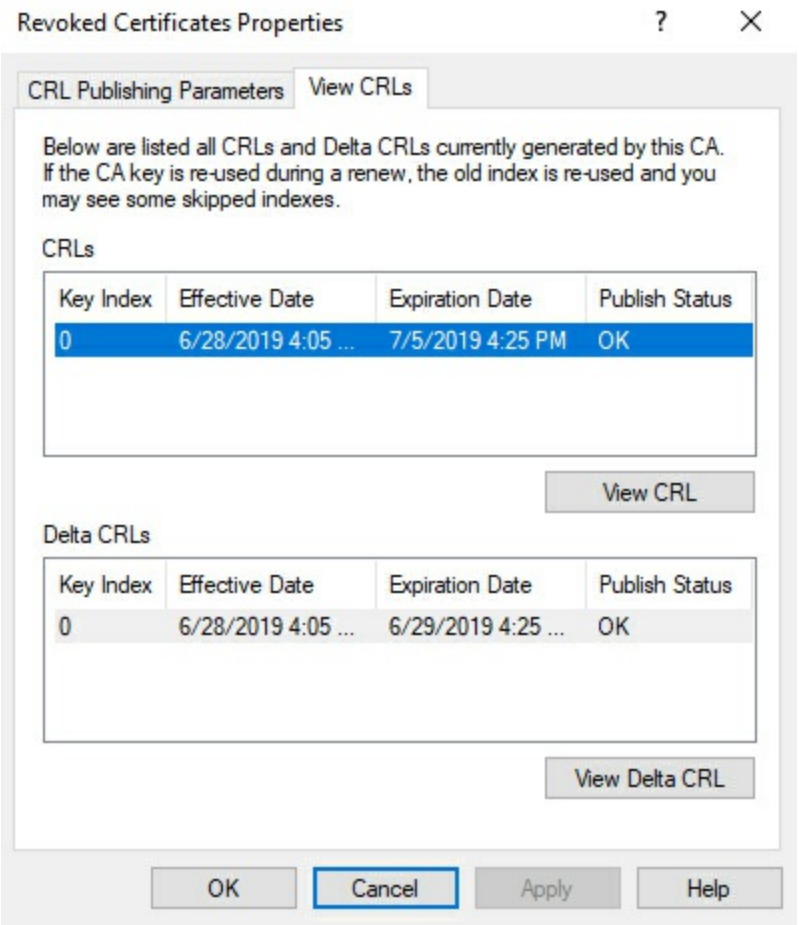


Figure 12-18 The View CRLs tab

You can manually trigger publication of a CRL or delta CRL by right-clicking the Revoked Certificates node, clicking All Tasks, and then clicking Publish. Select

whether you want to publish a New CRL or a Delta CRL Only, as shown in Figure 12-19.

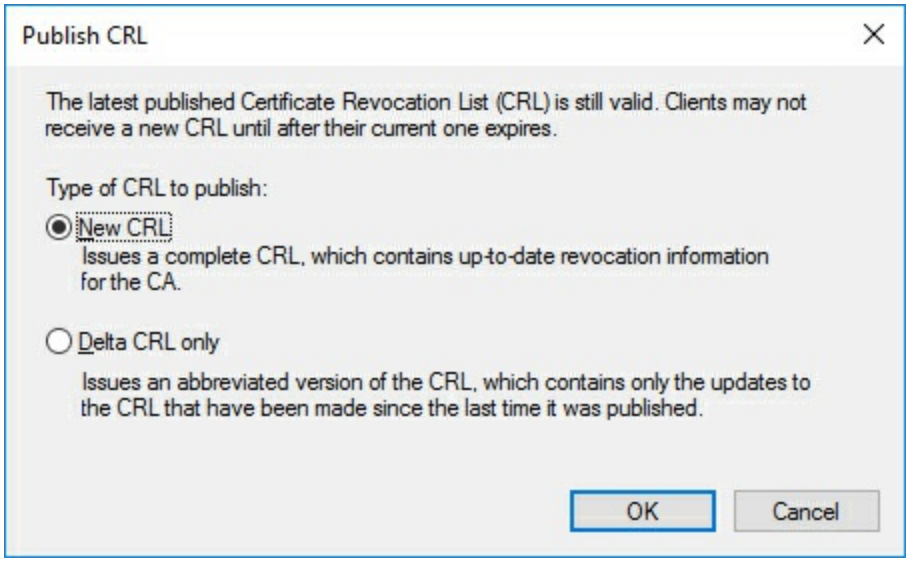


Figure 12-19 The Publish CRL window

CERTIFICATE SERVICES ROLE SERVICES

The primary function of AD CS is the CA. In addition to the CA role, there are several ancillary roles that can be used for certificate requests, enrollment, and to handle revocation traffic. These roles are as follows:

- **Certificate Enrollment Policy Web Service** Allows computers that are not members of the domain to obtain certificate enrollment policy information.
- **Certificate Enrollment Web Service** Enables users in the same forest as the CA to interact with the CA through a web browser to request certificates, renew certificates, and retrieve CRLs.
- **Certification Authority Web Enrollment** This role service

Technet24

allows people to navigate to a webpage and perform a certificate request. This can be either by filling out a web form with all the necessary details or by pasting the contents of a certificate request file into the form. When the certificate is issued, people can revisit this webpage to retrieve their issued certificate.

- **Network Device Enrollment Service** The Network Device Enrollment Service is a special service that you can configure to allow network devices like routers, switches, firewalls, and hardware-based VPN gateways to obtain certificates from the CA.
- **Online Responder** Online responders function in the same manner as CRL distribution points in that they provide a location where a client consuming a certificate can check whether a certificate is valid. The benefit of an online responder over a traditional CRL distribution point is that instead of downloading the entire revocation list, the client queries the online responder with a specific certificate ID.

CERTIFICATE TEMPLATES

Certificate templates allow you to configure the properties of different types of certificates. While it is possible to manually format a certificate request using templates, an enterprise CA can automatically populate a certificate request with relevant information stored in Active Directory.

By default, Certificate Authorities are only configured to issue a subset of all the pre-configured certificate templates stored in Active Directory. You can configure each enterprise CA in your environment to issue certificates based on specific certificate templates. For example, as a security precaution, you might want to

configure your PKI infrastructure so that only one CA can issue Data Recovery Agent certificates.

Deleting a template from the list of templates that can be issued in the Certification Authority console does not delete the certificate template from Active Directory. You can't remove any of the default certificate templates included with Windows Server 2019 from Active Directory using either the Certification Authority console or the Certificate Templates console. You can, however, remove custom certificate templates from Active Directory. Removing a certificate template just means that no new certificates can be issued from that template. It doesn't revoke certificates issued from the template that are already in circulation.

To configure a CA to issue a specific certificate template, perform the following steps:

1. Select the Certificate Templates node under the CA in the Certification Authority console, as shown in Figure 12-20. This shows you which certificate templates the CA is already configured to issue.

Technet24

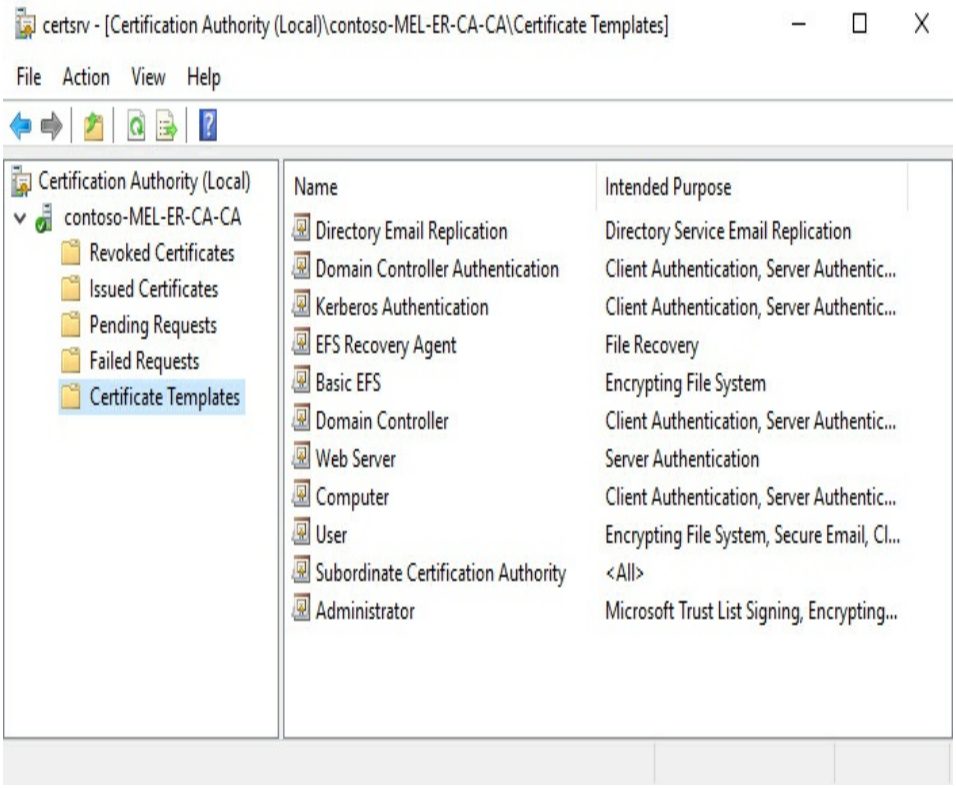


Figure 12-20 Certificate Templates in the Certification Authority console

2. If you want to have the CA issue certificates from a template not listed, you need to add the template. To do this, select the Certificate Templates node and click New from the Action menu. Next, click Certificate Template To Issue. This brings up the Enable Certificate Templates dialog box, shown in Figure 12-21. Select the certificate template that you want to issue and then click OK.

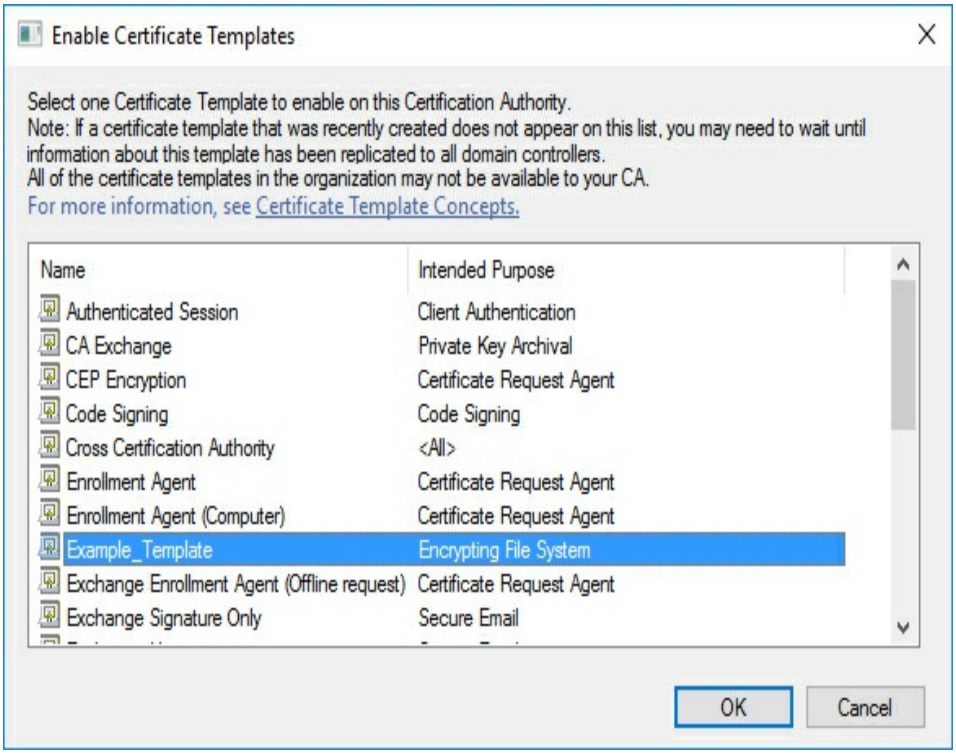


Figure 12-21 Enable Certificate Templates

Template properties

You manage certificate templates from the Certificate Templates console. Each certificate template has properties that you configure to determine how certificates are issued from the template function. The properties available in a template depend on the compatibility settings of the template. There are two compatibility settings: one for the CA and another for the certificate recipient. Figure 12-22 shows the CA setting configured for Windows Server 2016 and the certificate recipient set to Windows 10 / Windows Server 2016. Windows Server 2016 is the most recent

version that can be configured on a CA running the Windows Server 2019 operating system, and there are no specific Windows Server 2019 settings or features. When you specify Windows Server 2016, this template can only be used with a Windows Server 2016 or a Windows Server 2019 CA. Over the next few pages, you'll learn about some of the important areas of a certificate template's properties.

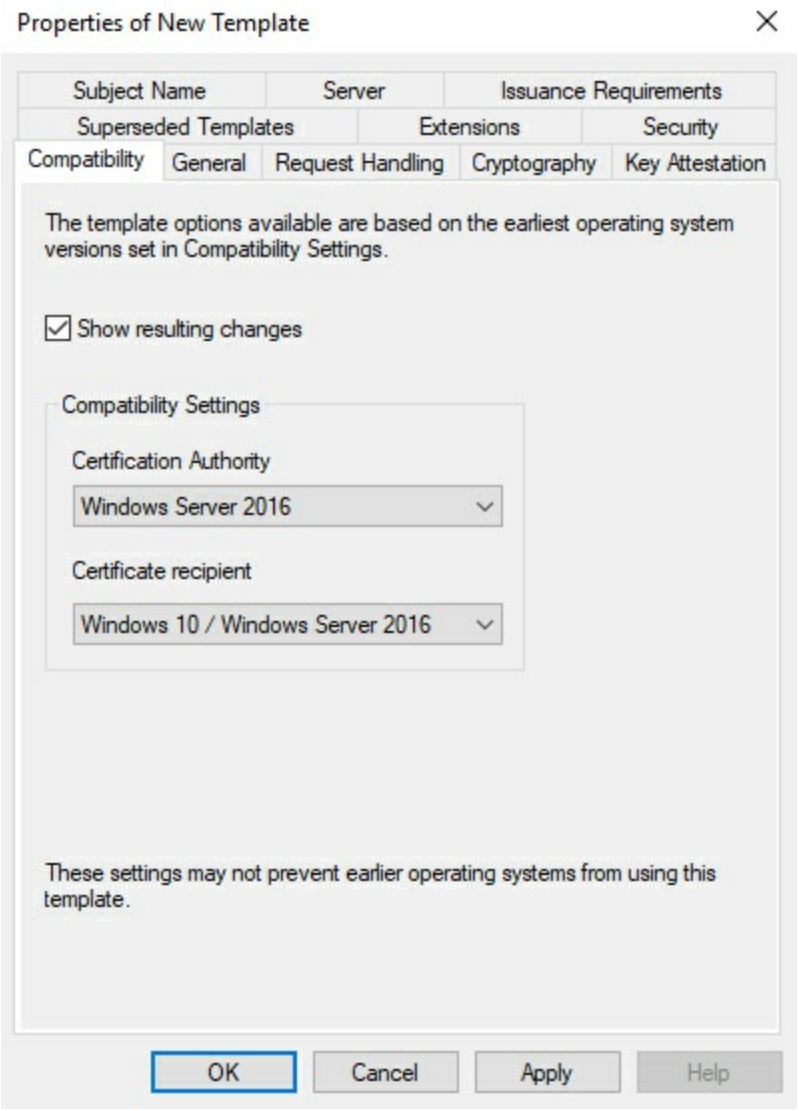


Figure 12-22 Compatibility tab

The General tab of a template, shown in Figure 12-23, allows you to configure both the Template Name and the Template Display Name. You also configure the certificate Validity Period and the Renewal Period in this tab. You can also choose whether the certificate is published in Active Directory. Publication to Active

Technet24

Directory is useful when other people in the organization might need to access a public certificate. If it's published in Active Directory, the public certificate doesn't have to be manually shared or retrieved. The private key is not published to Active Directory when the Publish Certificate in Active Directory option is enabled. In this example, Template Display Name is set to *Example_Template*. The Validity Period is set to 1 year, and the Renewal Period is set to 6 weeks. The template is configured to be published in Active Directory.

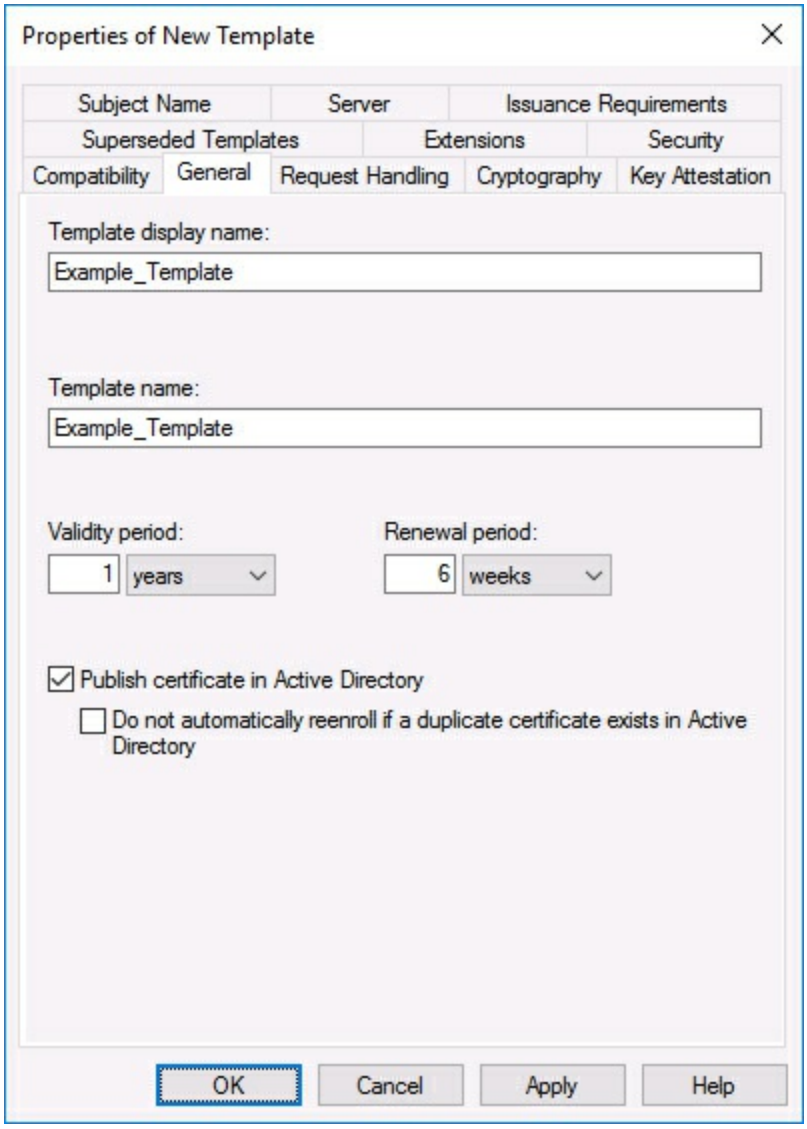


Figure 12-23 General tab

The Request Handling tab of the certificate template properties dialog box, shown in Figure 12-24, allows you to specify the certificate's purpose, which can be one of the following:

- Encryption

Technet24

- Signature
- Signature And Encryption
- Signature And Smartcard Logon

In Figure 12-24, the Purpose is set to Encryption. The Include Symmetric Algorithms Allowed By The Subject option is enabled. The Allow Private Key To Be Exported and the Enroll Subject Without Requiring Any User Input options are enabled. You can also use this tab to specify whether the subject's encryption private key is archived in the CA database, whether the private key can be exported, whether the user is prompted during enrollment, and when the private key is used.

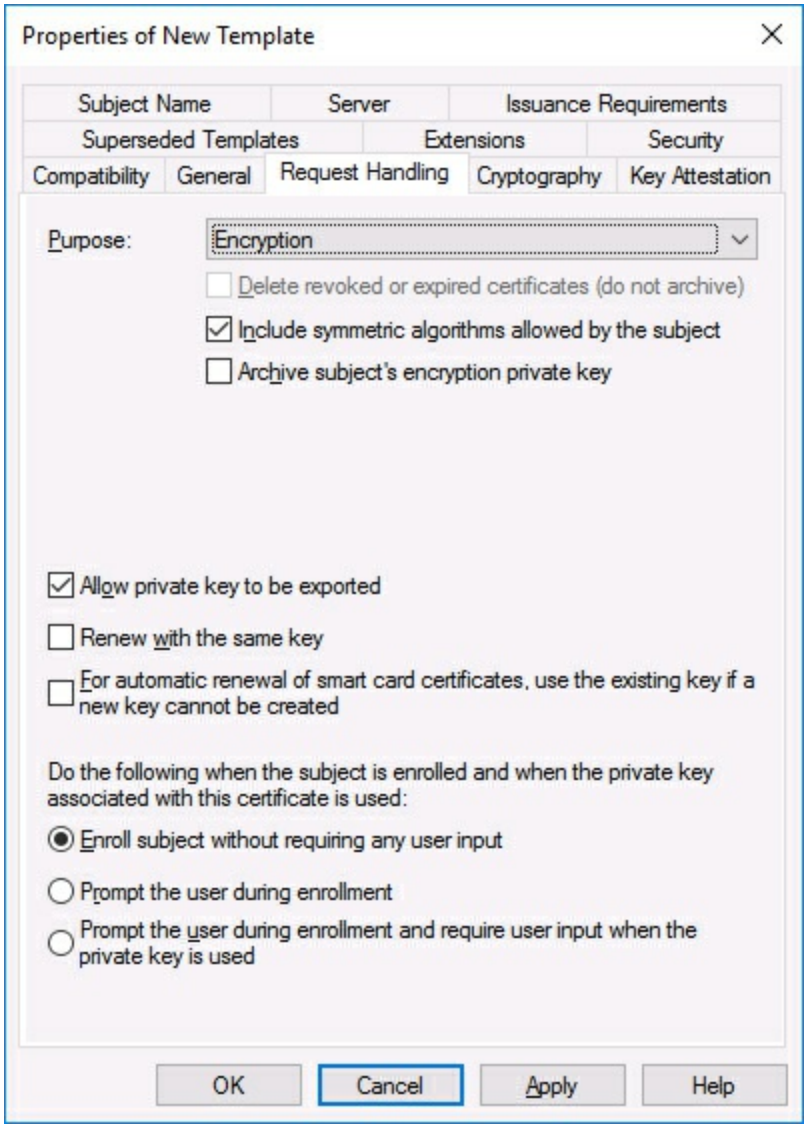


Figure 12-24 Request Handling tab

The Cryptography tab of a certificate's properties, shown in Figure 12-25, allows you to specify the cryptographic properties of the certificate issued from the template. This includes specifying a Minimum Key Size. You also use this tab to specify the properties of the cryptographic providers used for the certificate request.

Technet24

You do this to ensure that the certificate request meets a specific cryptographic benchmark. These settings are useful in environments that have very specific security requirements. In this example, Provider Category is set to Legacy Cryptographic Service Provider. Algorithm Name is set to Determined By CSP. Minimum Key Size is 2048. The Microsoft Enhanced Cryptographic Provider v 1.0 is selected.

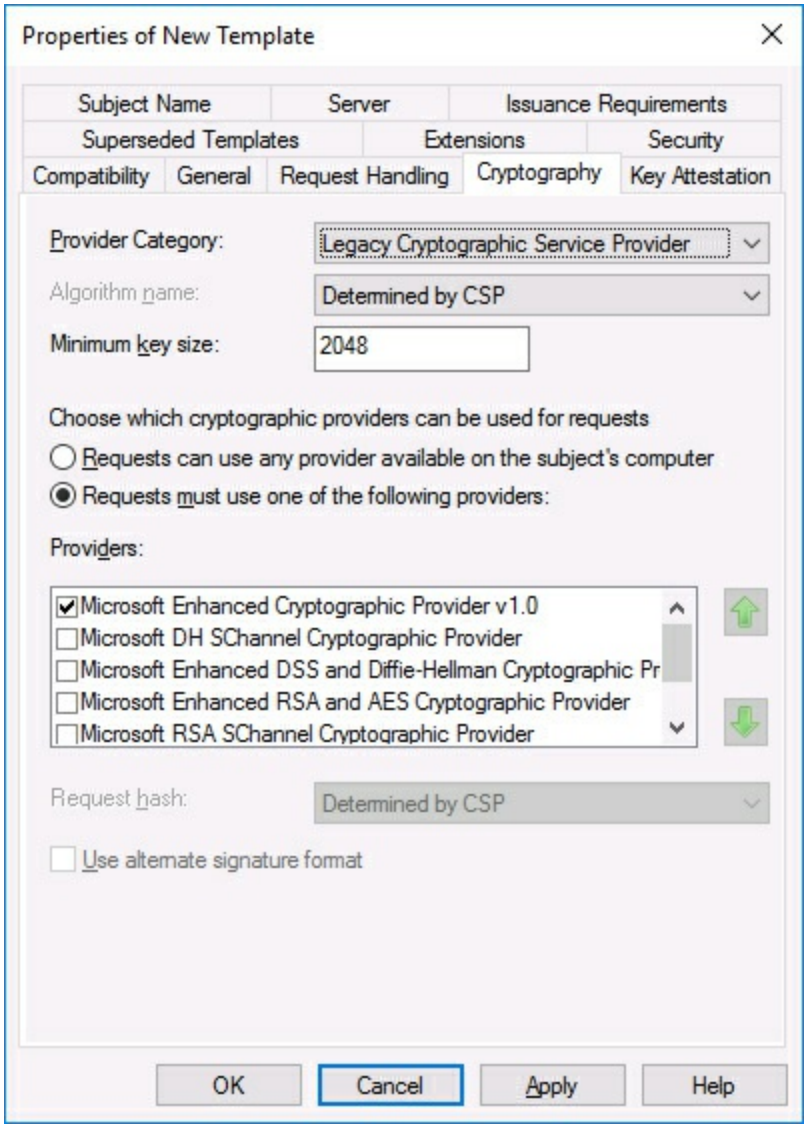


Figure 12-25 Cryptography tab

The Superseded Templates tab, shown in Figure 12-26, allows you to configure existing certificate templates that a newer template supersedes. For example, you might create a more advanced EFS certificate template that uses stronger cryptography and then specify the existing Basic EFS template as a template that the new

Technet24

template supersedes. When you configure autoenrollment correctly, a certificate from the new template automatically replaces the certificate from the superseded template.

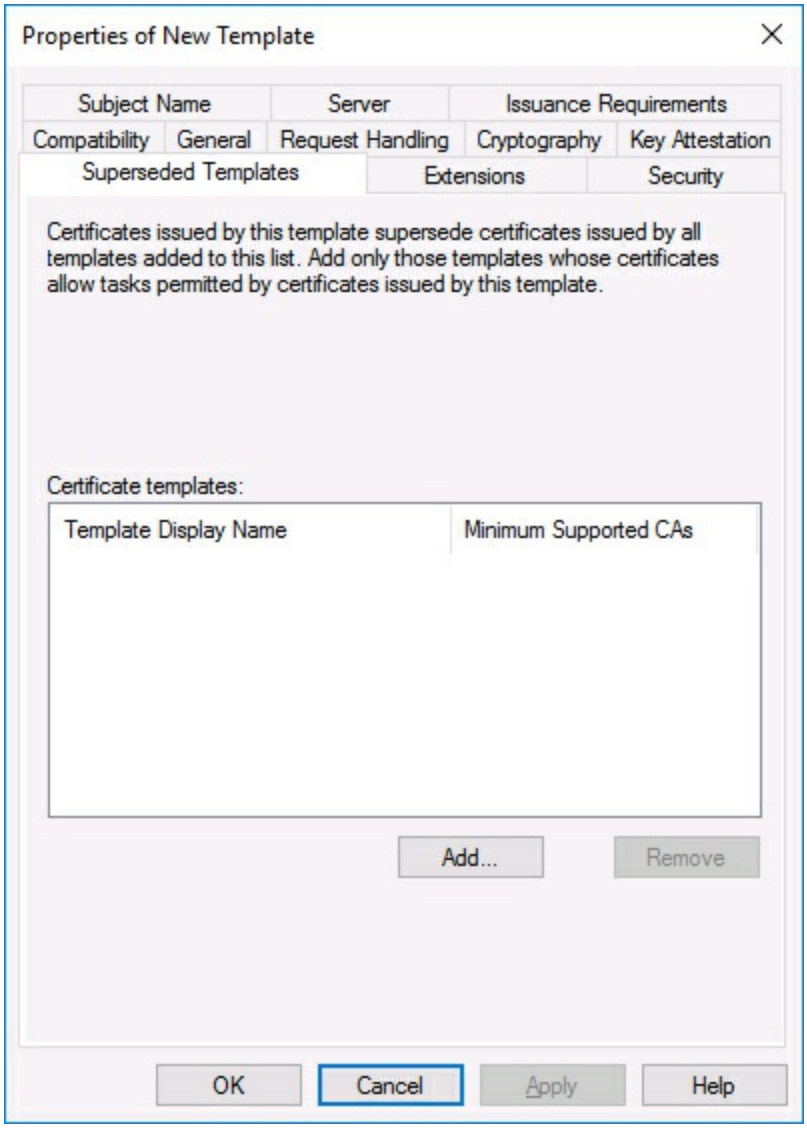


Figure 12-26 Superseded Templates tab

The Security tab, shown in Figure 12-27, allows you to

configure which users and groups have permission to enroll and autoenroll in the certificate. A user who has permission to enroll in the certificate can request the certificate. Autoenrollment means the certificate request is performed automatically. You configure autoenrollment through group policy. In Figure 12-27, the Admin Account is assigned the Read and Write (Allow) permissions. This process is described in more detail later in this chapter.

Technet24

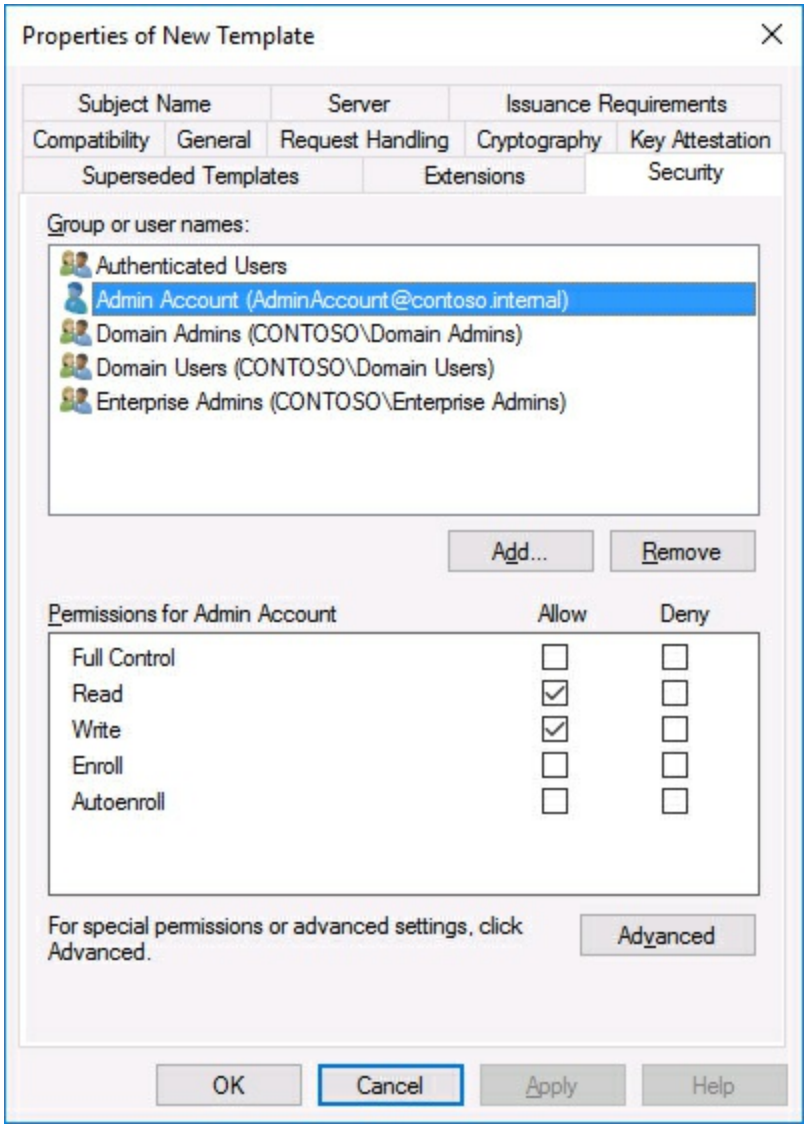


Figure 12-27 Security tab

You use the Issuance Requirements tab, shown in Figure 12-28, to specify whether requests against the template require certificate manager approval. You can also configure a certificate to require approval from more than one manager, which is useful for certificates issued from sensitive templates, such as the Key

Recovery Agent template. In this example, the CA Certificate Manager Approval option is enabled. The This Number Of Authorized Signatures option is set to *1*. The Policy Type Required In Signature option is set to Application Policy. The Application Policy is set to Any Purpose.

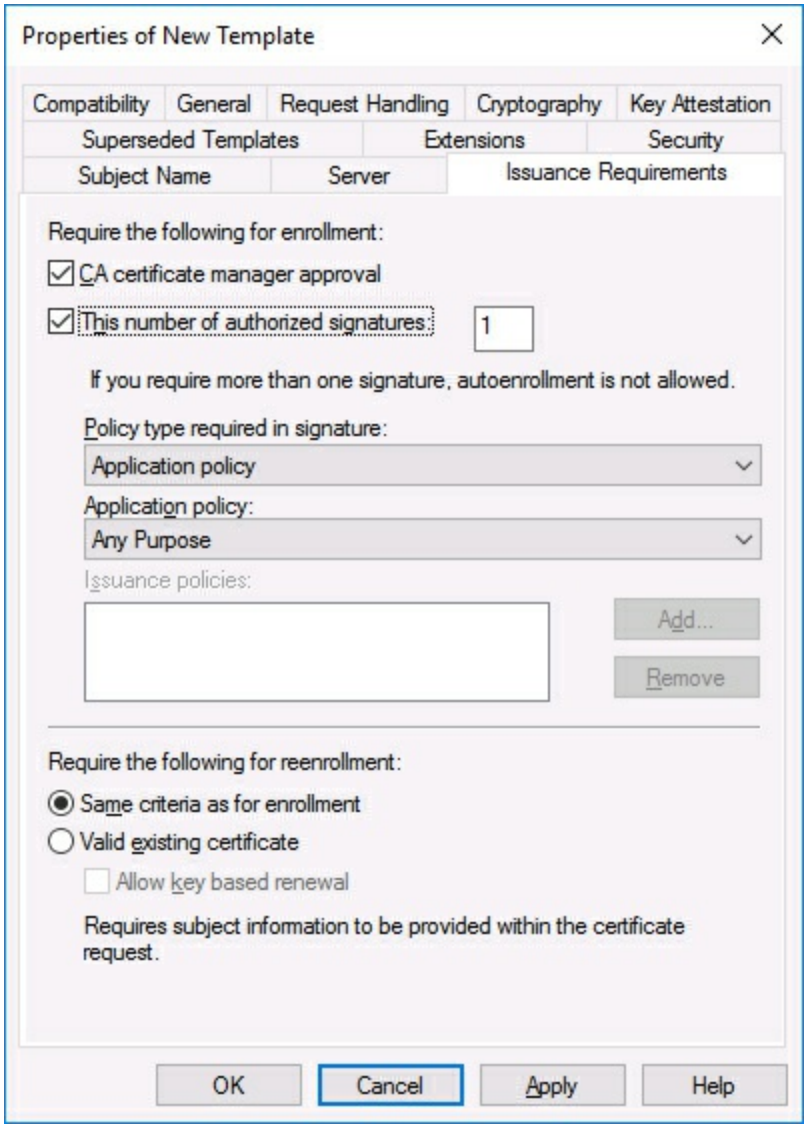


Figure 12-28 Issuance Requirements tab

Adding and editing templates

Windows Server 2019 includes 34 built-in templates. While you can edit some of the default templates should you want to make modifications, the best practice is to create a duplicate of the default template and edit the duplicate. In the event that something goes wrong, you've still got the original. When your new certificate template is ready, you can add the new template to a CA and delete the old one from the CA. To create a duplicate template, right-click the template that you want to duplicate and click Duplicate Template. Templates are stored within Active Directory and automatically replicate to be available to any enterprise CA.

CERTIFICATE AUTOENROLLMENT AND RENEWAL

An advantage of enterprise CAs is that they allow for automatic certificate enrollment and automatic certificate renewal. This means that certificates can be issued automatically to users, services, or devices without requiring them to perform any explicit action. Autoenrollment substantially reduces the amount of administrative effort required to deploy and update certificates.

To configure autoenrollment, perform the following configuration steps for the certificate template

associated with the certificate that you want to autoenroll:

- Ensure that the certificate template for the certificate you want to allow autoenrollment for is configured for a Windows Server 2008 or higher level of compatibility. You configure this on the Compatibility tab of the certificate template's properties.
- Configure the autoenroll permission for the certificate on the Security tab of the certificate template's Properties dialog box, as shown in Figure 12-29. You should configure this permission for the security group that you want to automatically enroll. In this figure, the Authenticated Users group is assigned the Read and Autoenroll (Allow) permissions.

Technet24

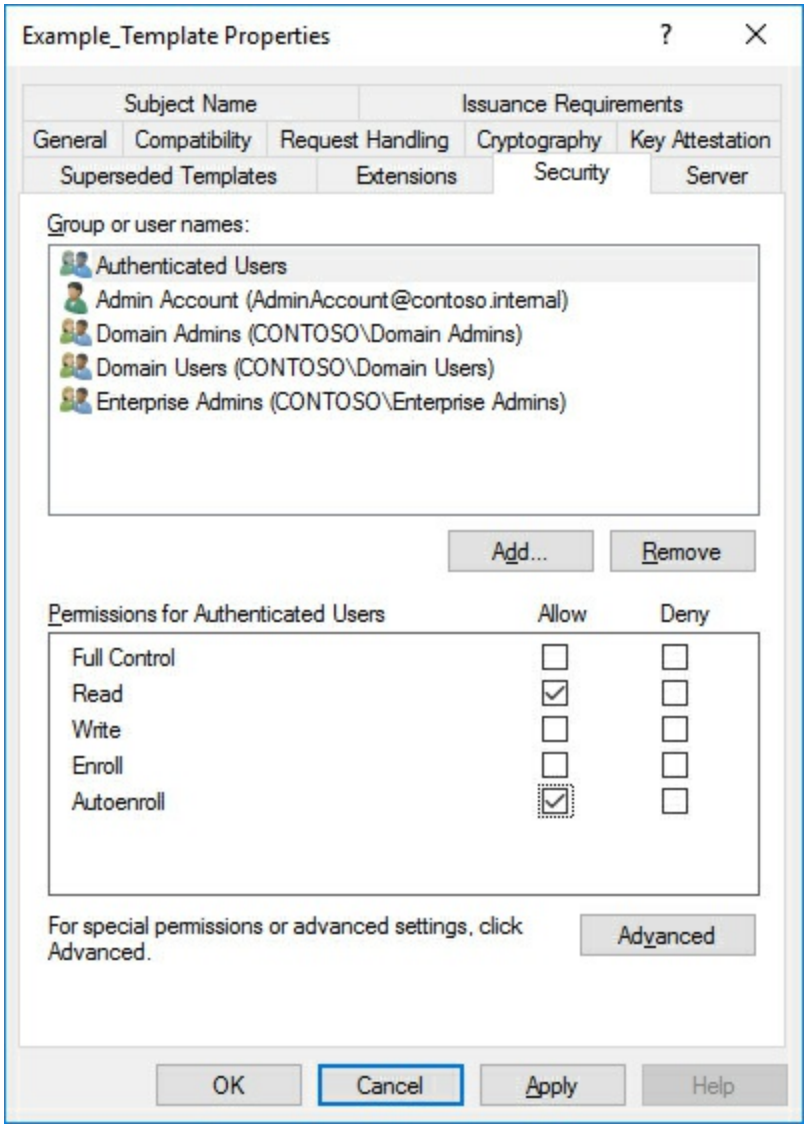


Figure 12-29 Security tab

The other half of configuring autoenrollment is configuring the Certificate Services Client policy in the Default Domain GPO. Depending on whether you want to configure autoenrollment for users or computers, this policy is found in one of the following nodes:

- *Policies\User Configuration\Windows Settings\Security Settings\Public Key Policies*
- *Policies\Computer Configuration\Windows Settings\Security Settings\Public Key Policies*

You can configure the following settings for this policy, as shown in Figure 12-30. In the figure, The Configuration Model, the Renew Expired Certificates, and the Update Certificates That Use Certificate Templates options are set to Enabled.

- **Configuration Model** You can set this to Enabled, Disabled, or Not Configured. Setting this to Enabled allows certificate autoenrollment.
- **Renew Expired Certificates, Update Pending Certificates, And Remove Revoked Certificates** When you configure this setting, certificates that are configured for autoenrollment are automatically renewed before they expire.
- **Update Certificates That Use Certificate Templates** This policy triggers autoenrollment in certificates that use new templates configured to supersede certificates issued from an existing template. Supersede settings are configured within the new certificate's template. The new certificate also needs to be configured with an appropriate autoenrollment permission.
- **Expiration Notification** When enabled, a user is prompted to trigger the reenrollment process. Notifications are triggered based on the remaining percentage of the certificate's lifetime.

Technet24

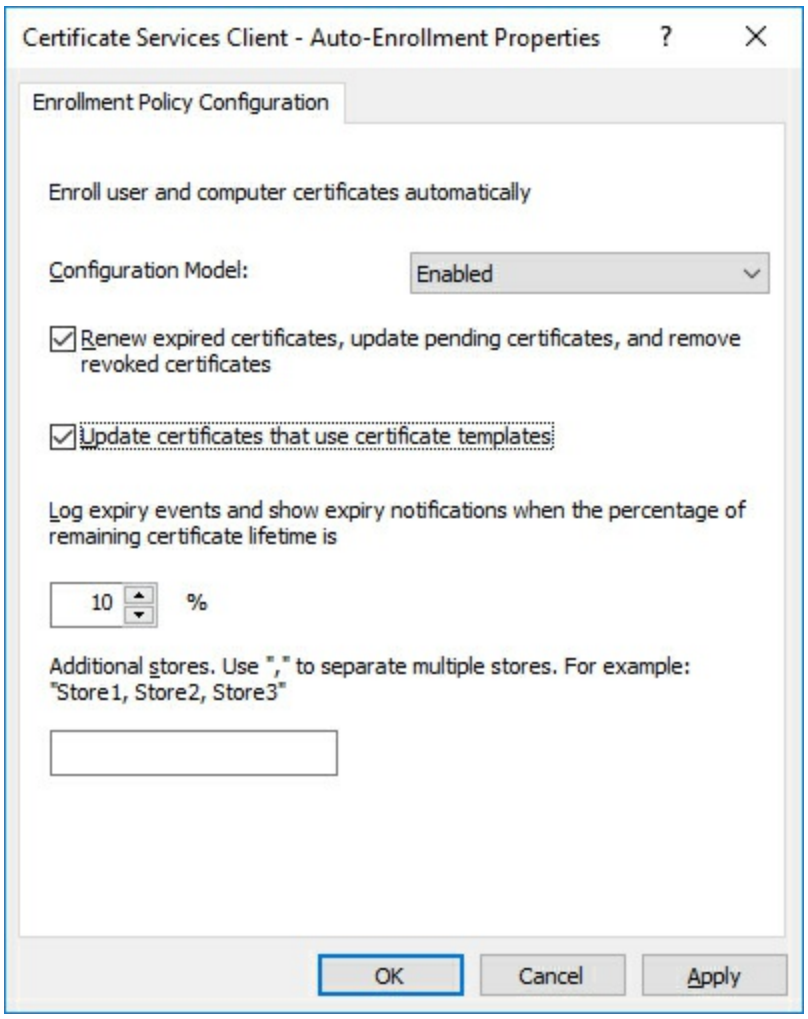


Figure 12-30 Enrollment policy configuration

CA MANAGEMENT

You assign roles on a CA by assigning one of the following permissions, as shown in Figure 12-31. In the figure, The Enterprise Admins group has been assigned the Issue And Manage Certificates and Manage CA permissions.

- **Issue And Manage Certificates.** You assign this permission

when you want to allow someone to issue and revoke certificates.

- **Manage CA.** You assign this permission when you want to allow someone to edit the properties of the CA. This includes the ability to assign CA permissions. You can't directly block someone to whom you have assigned the Manage CA permission from granting themselves the Issue And Manage Certificates permission. Only by configuring auditing and alerts are you able to restrict this action. By default, members of the Enterprise Admins, Domain Admins, and Local Administrators groups have the Manage CA permission on a computer that is a member of a domain.
- **Request Certificates.** You use this permission to control who can request certificates from the CA. Should you have a CA that you've configured to only issue certificates from especially sensitive templates, such as Key Recovery, you might choose to alter this setting.
- **Read.** You use this permission to allow someone to view CA information.

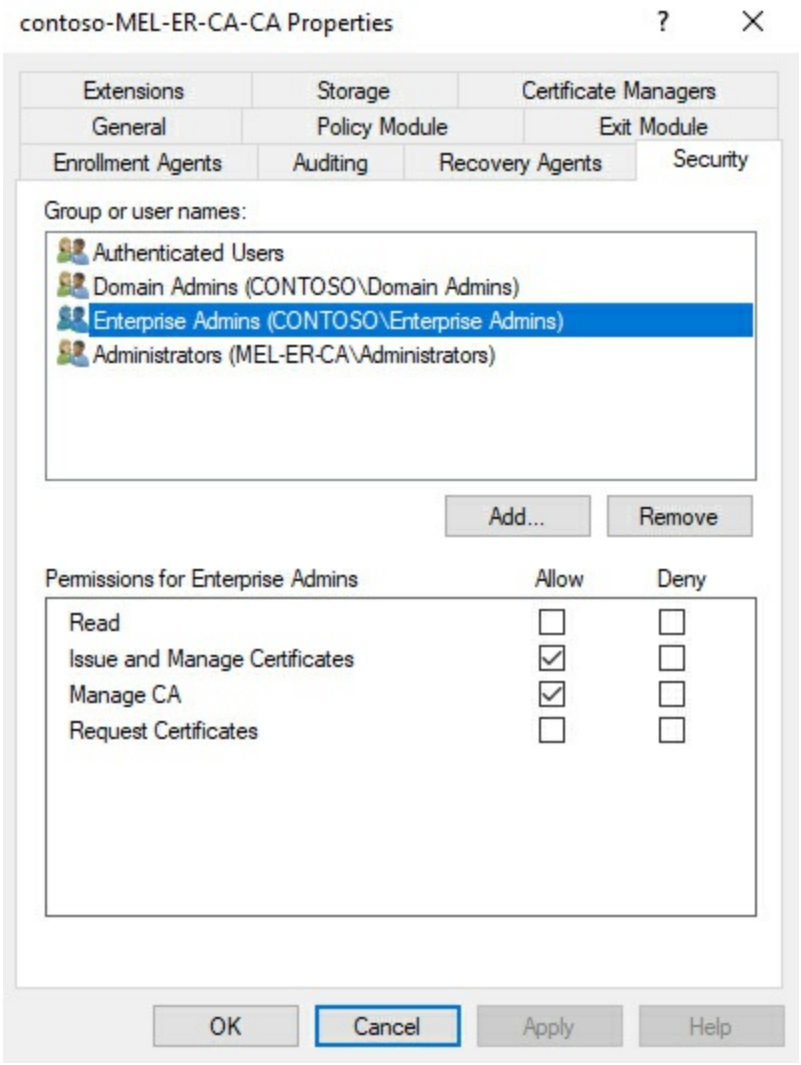
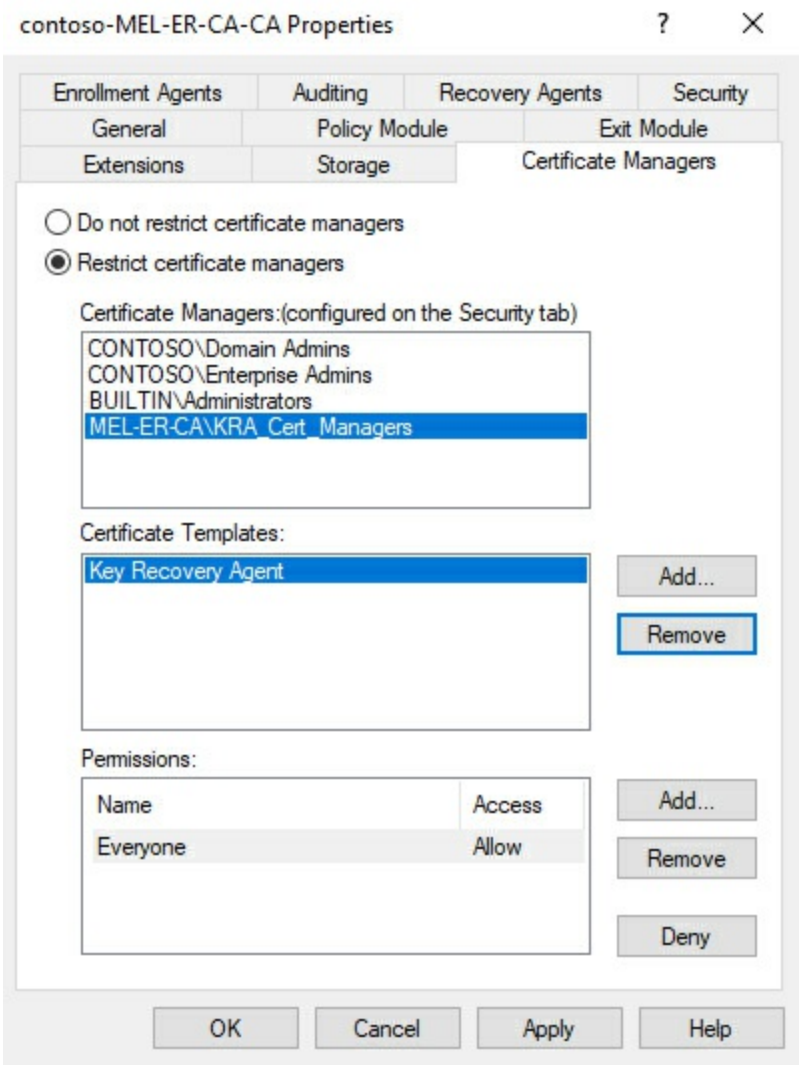


Figure 12-31 Security permissions

You can go further when configuring the ability to issue and manage certificates by configuring this ability based on certificate templates. To configure a security group so it can issue certificates based on a specific template, you first need to assign the group the Issue And Manage Certificates permission on the Security tab of the CA

properties dialog box. Then on the Certificate Managers tab, you need to choose the Restrict Certificate Managers option. After you've done that, you edit the list of certificate templates for which each group can manage certificates. Figure 12-32 shows the local *KRA_Cert_Managers* group configured to manage certificates based off the Key Recovery Agent certificate template.



Technet24

Figure 12-32 Certificate Managers

Handling certificate requests

By default, all certificates on a standalone CA require approval. An enterprise CA uses the properties configured in the CA template. To approve a certificate, navigate to the Pending Requests node of the Certification Authority console, select the Pending Requests node, select the certificate that you want to approve or deny, click All Tasks from the Action menu, and then click either Issue or Deny.

You can alter whether an enterprise CA uses the template to determine whether approval is required—meaning all certificates require approval—by editing the properties of the CA. To do so, click the Properties button on the Policy Module tab. This opens the Properties dialog box, as shown in Figure 12-33, where you can configure this setting. In the figure, Follow The Settings In The Certificate Template, If Applicable. Otherwise Automatically Issue The Certificate option is enabled. The other option, Set The Certificate Request Status To Pending, is not selected.

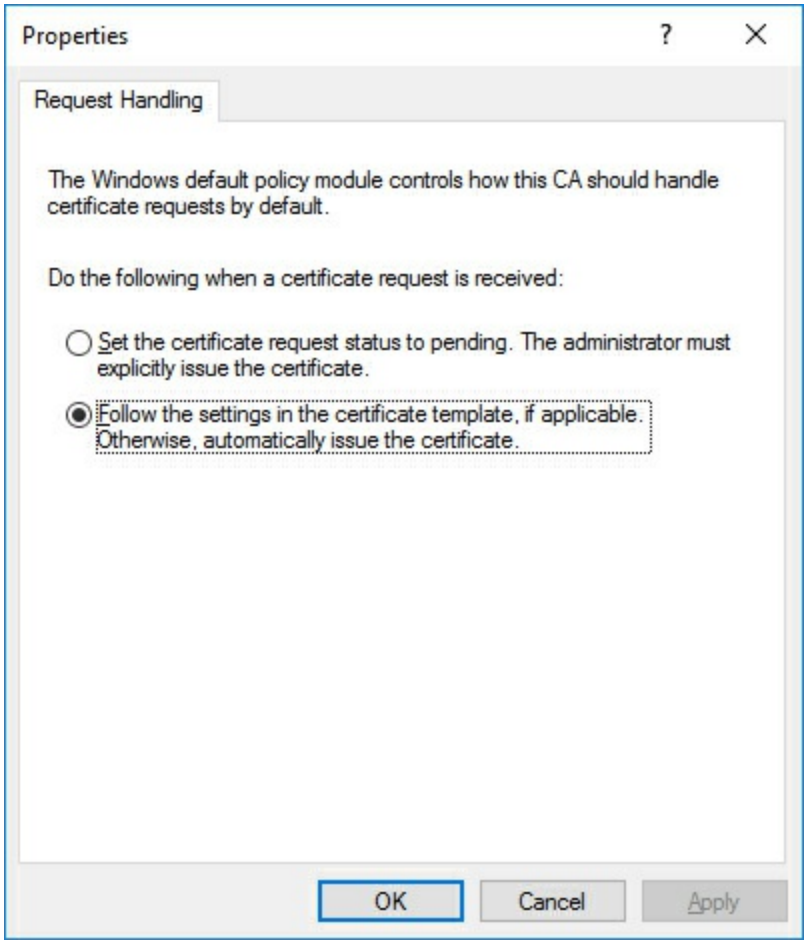


Figure 12-33 Request Handling

CA backup and recovery

You can back up and recover a CA, including the private key, CA certificate, CA database, and CA log from the Certification Authority console. To back up a CA, perform the following steps:

1. In the Certification Authority console, select the CA that you want to back up. On the Action menu, click All Tasks and click Back Up CA.
2. On the Welcome To The Certification Authority Backup page of the Certification Authority Backup Wizard, click Next.

Technet24

3. On the Items To Backup page, choose from the following options, as shown in Figure 12-34. In the figure, The Private Key And CA Certificate and Certificate Database And Certificate Database Log settings are selected. The Back Up To This Location is set to *C:\cabackup*.

    1. **Private Key And CA Certificate** Backs up the CA's private and public keys. Enables you to restore the CA on a different computer in the event that the CA fails.
    2. **Certificate Database And Certificate Database Log** Enables you to recover the public keys of the certificates that the CA has issued. If key archiving is enabled, this option enables you to recover private keys of these certificates.

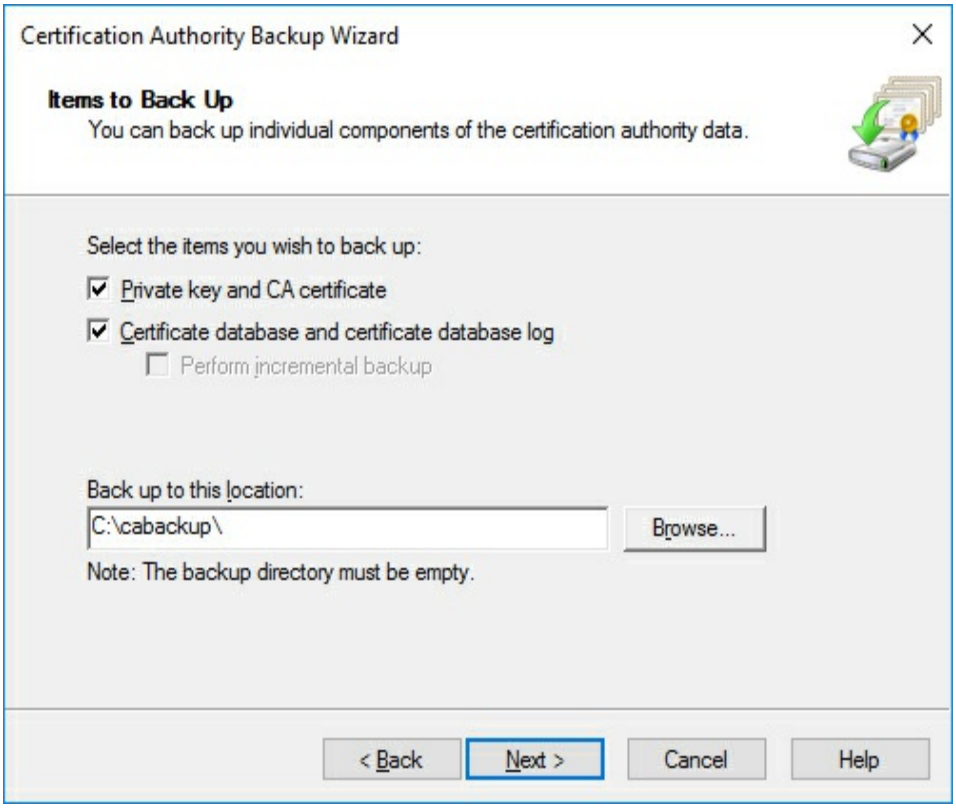


Figure 12-34 CA Backup Wizard

4. Provide a password to encrypt the backup data.

To recover a backup, perform the following steps:

1. In the Certification Authority console, right-click the CA, click All Tasks, and then click Restore CA.
2. When you are warned about stopping Active Directory Certificate services, click OK.
3. On the Items To Restore page, shown in Figure 12-35, select the folder that holds the CA backup data and click Next. In this example, The Private Key And CA Certificate and Certificate Database And Certificate Database Log settings are selected. The Restore From This Location setting is set to *C:\cabackup*.

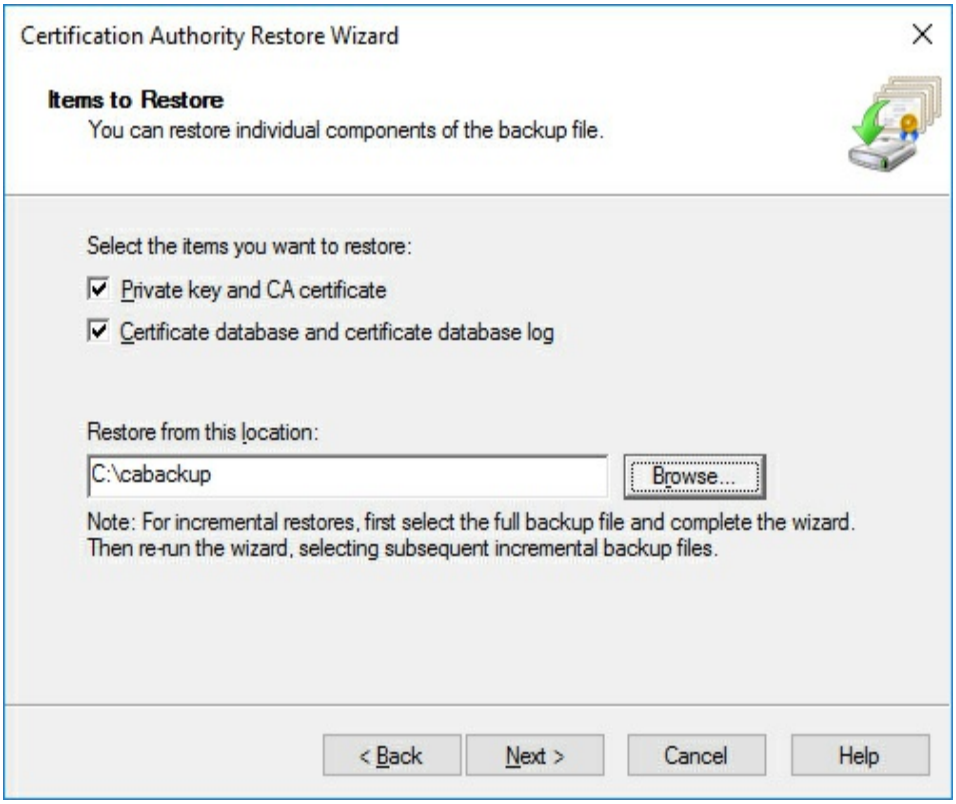


Figure 12-35 CA Restore

4. Provide the password used to protect the backup data and then complete the wizard. When the restoration is complete, you are prompted to restart Active Directory Certificate Services.

One thing that CA backup does not back up is the list of certificate templates that a specific enterprise CA has been configured to issue. If you have configured a CA to issue additional templates beyond the default templates, you need to configure a recovered CA to use the same templates. One method to keep track of which certificate templates a CA is configured to use is to use the *Get-CATemplate* PowerShell cmdlet. You can pipe the output of this cmdlet to a text file, which you can use as a record of which templates a specific CA has been configured to issue.

Key archiving and recovery

A key recovery agent allows an entity holding a properly configured key recovery agent certificate to recover a private key that has been archived in the AD CS database. To perform key recovery on a private certificate, the template the certificate is generated from needs to be configured so that private keys are archived to the AD CS database. This is done on the Request Handling tab of the template's Properties, as shown in Figure 12-36. The Include Symmetric Algorithms Allowed By The Subject, Archive Subject's Encryption Private Key, and Allow Private Key To Be Exported options are enabled.

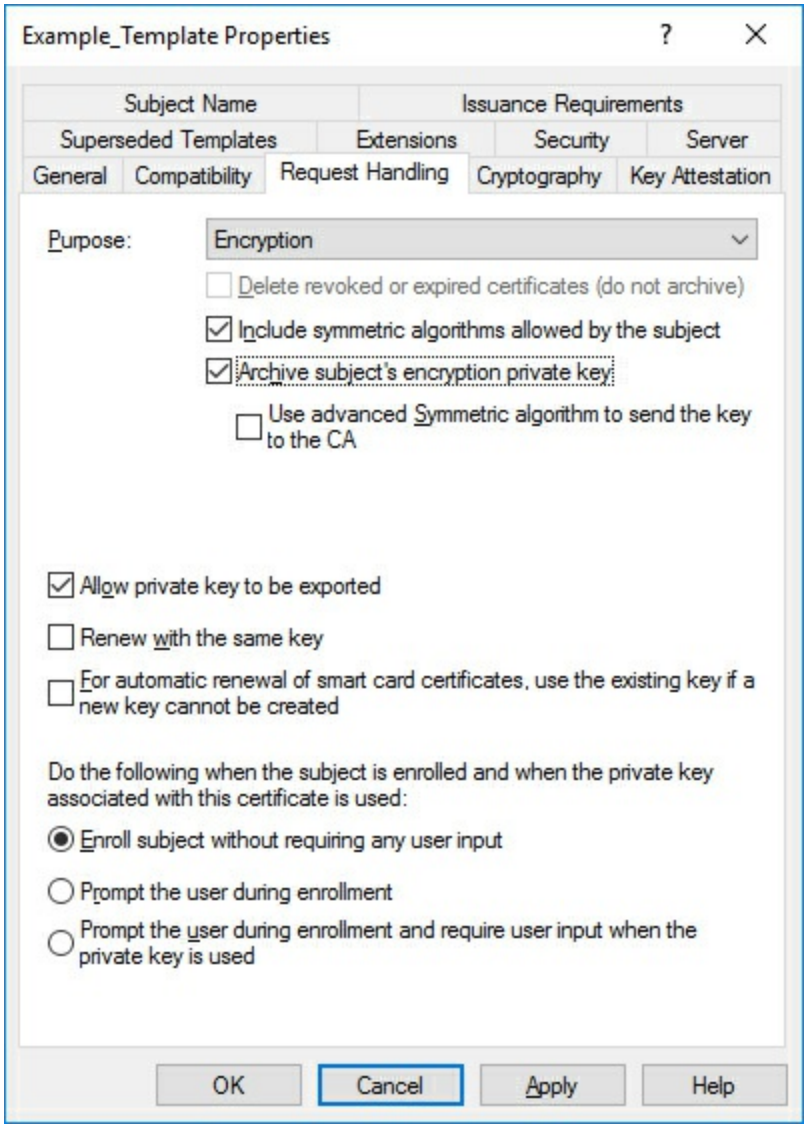


Figure 12-36 Archive private key

The CA also needs to be configured with one or more Recovery Agents. To configure a recovery agent, perform the following steps:

1. In production environments, create an account to be used by the recovery agent. While you can use the built-in administrator account, from a security perspective you're much better off having a separate

account for this role. In an ideal environment, the individual using this account is separate from the individuals who manage the CA. Requiring separate individuals ensures that private keys are only recovered in authorized circumstances. You want to avoid allowing just anyone to recover private keys without oversight, especially in scenarios that involve encryption.

2. Configure a CA to issue Key Recovery certificates. To do this, open the Certification Authority console, right-click Certificate Templates, click New, and then click Certificate Template to Issue. In the Enable Certificate Templates dialog box, select Key Recovery Agent as shown in Figure 12-37 and click OK.

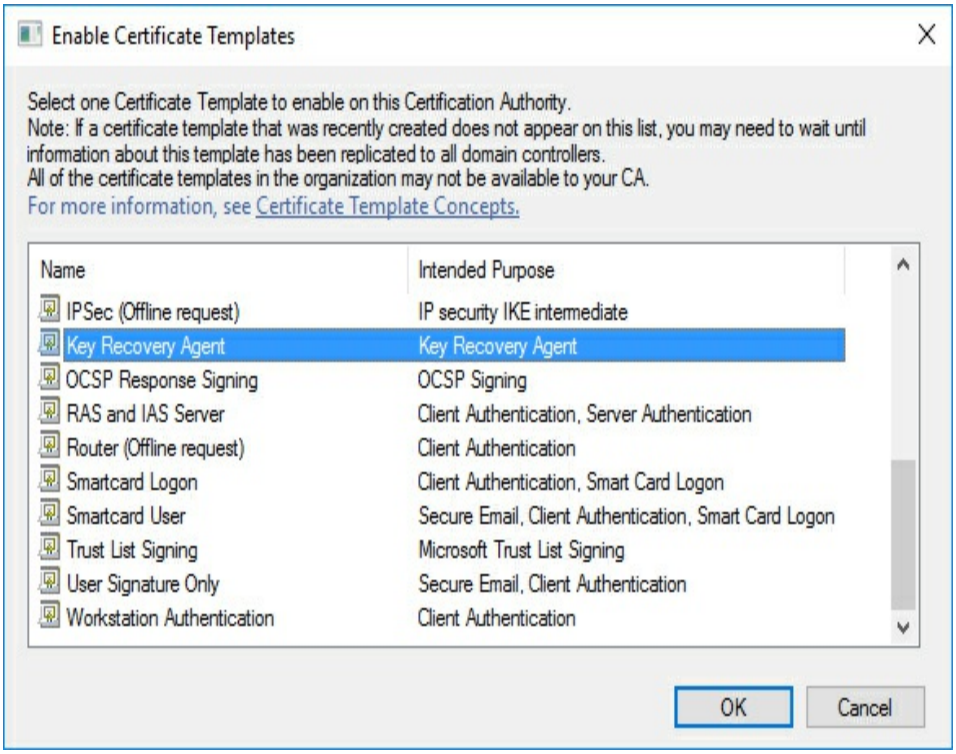


Figure 12-37 Certificate templates

3. The account you use for key recovery needs to then request a certificate based on this template. They can do this using the Certificates snap-in of the Microsoft Management console, as shown in Figure 12-38.

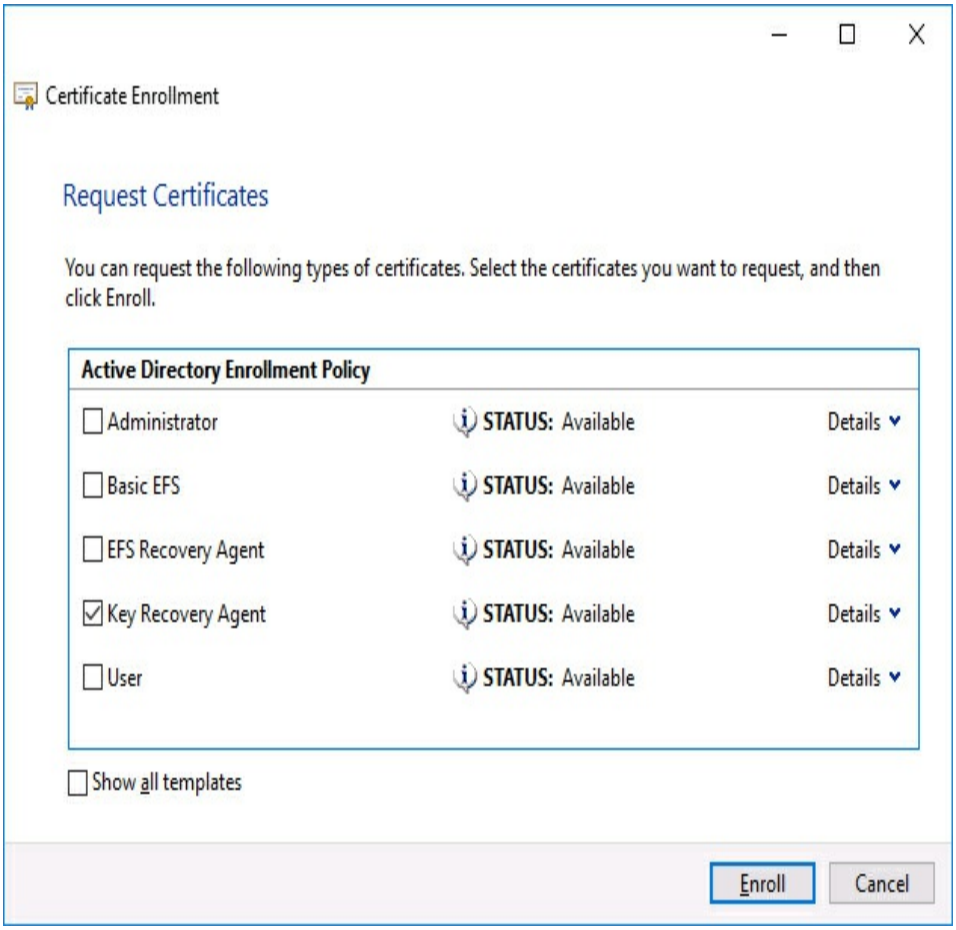


Figure 12-38 Certificate request

4. In the Certification Authority console, you need to approve the certificate request because the default template is configured to require approval, not to have the certificate be issued automatically.
5. After the certificate is issued, export it from the Certificate Enrollment Requests node of the Certificates snap-in, including the private key as shown in Figure 12-39. After you've exported the private key to a *.pfx* file, you should then export the public key to a *.cer* file. Import this public certificate to the CA to enable key recovery.

Technet24

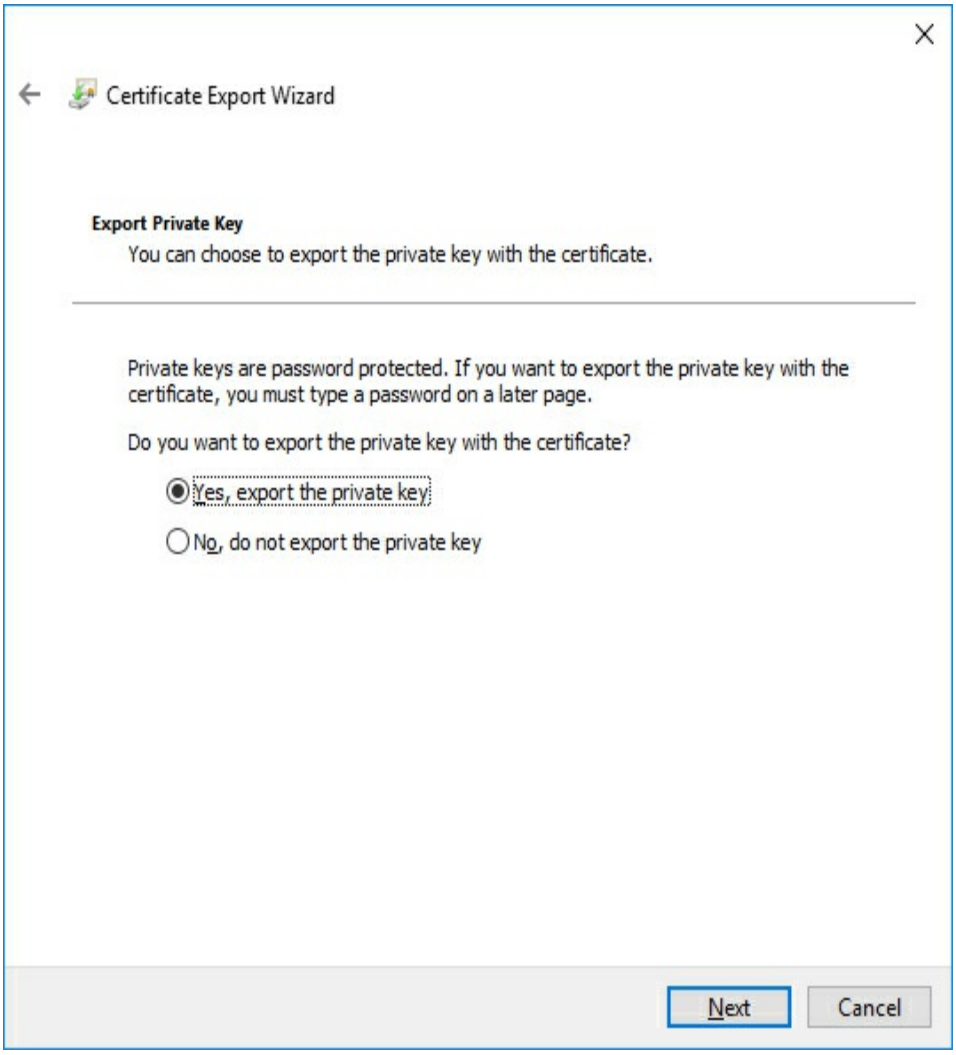


Figure 12-39 Export private key

6. To configure the CA to archive private keys, on the Recovery Agents tab, select Archive the key, click Add, and on the Key Recovery Agent Selection dialog box, shown in Figure 12-40, click the certificate, and then click OK. Note in Figure 12-40 that the name of the user account is *KeyRecovery*. If your key recovery account has a different name, it is listed in this dialog box.

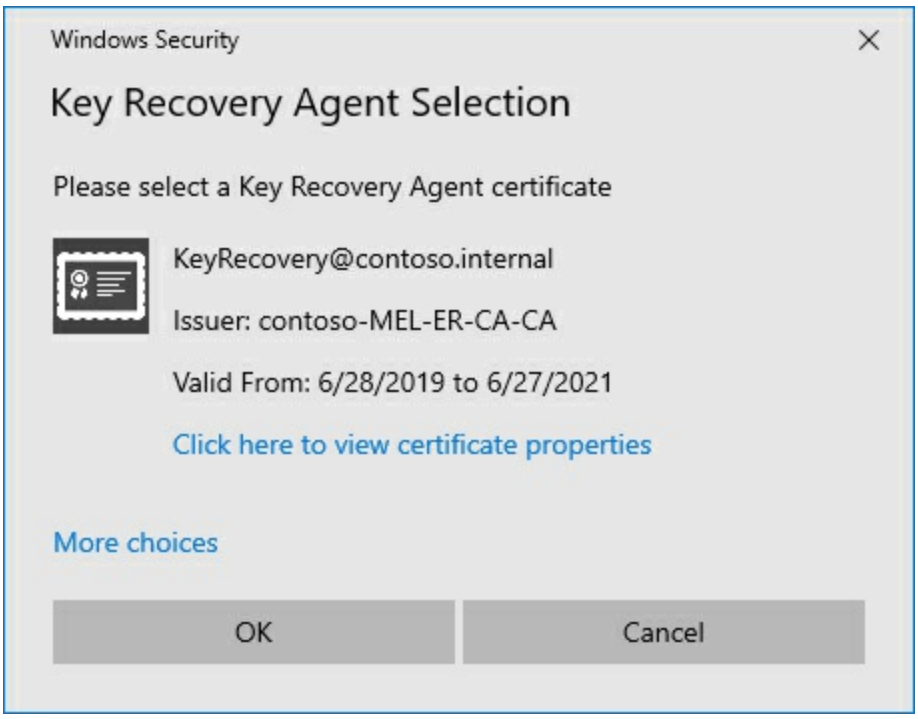


Figure 12-40 KRA selection

7. You need to restart the CA before the changes take place. When you've restarted the CA, the key recovery agent certificate is listed as Valid on the Recovery Agents tab, as shown in Figure 12-41.

Technet24

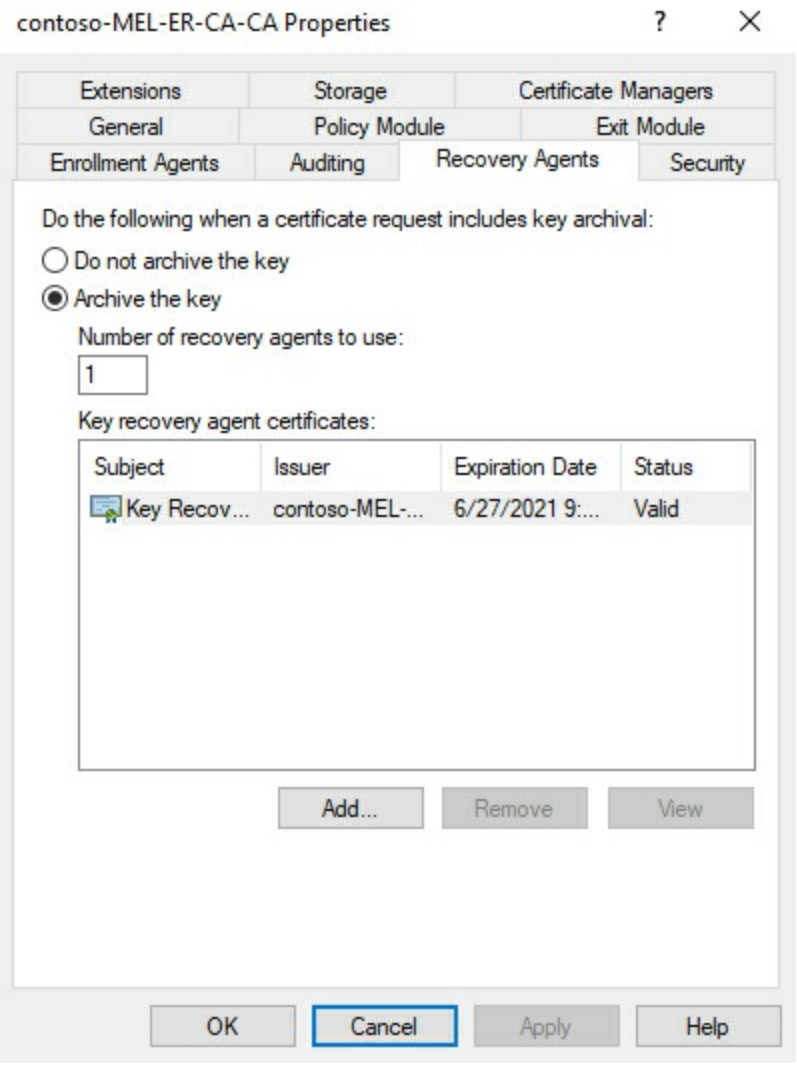


Figure 12-41 Recovery Agents

To recover an archived certificate using a KRA key, perform the following steps:

1. The first step is that a user who has CA Manager permissions needs to extract a blob that contains the private key. To do this, you need the serial number of the certificate that you want to recover. You do this by using the *certutil* command-line utility with the *-getkey* option. For example,

   Click here to view code image

```
Certutil -getkey <certificate_serial_number>
outputblob
```

2. The user who has the key recovery agent certificate in the certificate store then needs to recover the private key. This task is also performed using the *certutil* command-line utility.

   Click here to view code image

```
Certutil -recoverkey outputblob
recoveredkey.pfx
```

CAPolicy.inf

The *CAPolicy.inf* file, located in the *%systemroot%* (usually *C:\Windows*) folder, allows you to customize CA settings. This file is not mandatory and is not present by default, but it does allow you to provide a collection of settings to automatically configure the CA. You would use this primarily when providing certificates used outside your organization. If you have an internal CA that is only being used by organizational clients, and the certificates aren't consumed by anyone outside the organization, you likely don't have to worry about using a CAPolicy.inf file.

The file is a text file. For example, the following *CAPolicy.inf* file could be used for a root CA:

Click here to view code image

```
[Version]
Signature= "$Windows NT$"
[Certsrv_Server]
```

```
RenewalKeyLength=4096
RenewalValidityPeriod=Years
RenewalValidityPeriodUnits=10
LoadDefaultTemplates=0
```

An enterprise CA could use the following *CAPolicy.inf* file:

Click here to view code image

```
[Version]
Signature= "$Windows NT$"
[PolicyStatementExtension]
Policies = LegalPolicy, LimitedUsePolicy
[LegalPolicy]
OID = 1.1.1.1.1.1.1.1.1
URL = "http://www.contoso.com/Policy/LegalPolicy.asp"
URL = "ftp://ftp.contoso.com/Policy/LegalPolicy.txt"
[LimitedUsePolicy]
OID = 2.2.2.2.2.2.2.2.2
URL = "http://www.contoso.com/Policy/LimitedUsePolicy.asp"
URL = "ftp://ftp.contoso.com/Policy/LimitedUsePolicy.txt"
LoadDefaultTemplates=0
```

The OIDs listed above are just examples. In a real *CAPolicy.inf* file, you would use OIDs that you obtained from the ISO Name Registration Authority.

More Info

CAPolicy.inf

You can find out more about the syntax of CAPolicy.inf files at http://blogs.technet.com/b/askds/archive/2009/10/15/windows-server-2008-r2-capolicy-inf-syntax.aspx.

Managing Certificate Services using PowerShell

There are a number of PowerShell cmdlets that you can

use to install and configure Certificate Services. These cmdlets are included in the *ADCSAdministration* and *ADCSDeployment* modules and are listed in Table 12-2 and Table 12-3.

Table 12-2 ADCSAdministration PowerShell Module cmdlets

Cmdlet	Function
Add-CAAuthorityInformationAccess	Add a Uniform Resource Identifier (URI) for the Authority Information Access (AIA) or Online Responder OCSP address for a CA.
Add-CACrlDistributionPoint	Add a CRL distribution point to all subsequently issued certificates.
Add-CATemplate	Add a certificate template stored in Active Directory to the CA so that the CA can issue certificates from that template.
Backup-CARoleService	Backs up both the CA database and the CA private key.
Confirm-CAAttestationIdentityKeyInfo	Checks whether a CA trusts the attestation identity key for specific secure hardware.

Technet24

Confirm-CAEndorsementKeyInfo	Checks whether a CA trusts specific secure hardware for key attestation.
Get-CAAuthorityInformationAccess	View the AIA and OCSP URI information configured for the CA.
Get-CACrlDistributionPoint	View the CRL distribution point information configured for the CA.
Get-CATemplate	View the list of templates that the CA is configured to issue. This is not the entire list of templates stored within AD DS.
Remove-CAAuthorityInformationAccess	Remove a Uniform Resource Identifier (URI) for the Authority Information Access (AIA) or Online Responder OCSP address for a CA.
Remove-CACrlDistributionPoint	Remove a CRL distribution point. This applies only to all subsequently issued certificates. Existing certificates still retain distribution point information that was valid at the time the certificate was issued.
Remove-CATempl	Removes a certificate template from a CA.

ate	
Restore-CARoleService	Restore a CA database and private key from an existing backup.

Table 12-3 ADCSDeployment PowerShell Module cmdlets

Cmdlet	Function
Install-AdcsCertificationAuthority	Configures a CA. This cmdlet does not install the binaries, which need to be added separately using the *Install-WindowsFeature* cmdlet.
Install-AdcsEnrollmentPolicyWebService	Configures the Enrollment Policy Web Service on a computer where the binaries are already present.
Install-AdcsEnrollmentWebService	Configures the Enrollment Web Service on a computer where the binaries are already present. Can also be used to add additional instances of the service on the computer.
Install-AdcsNetworkDeviceEnrollmentService	Configure the Network Device Enrollment Service on a computer where the binaries are already present.
Install-AdcsOnlineResponder	Configures the Online Responder role service on a computer where the binaries are already present.

Technet24

Install-AdcsWebEnrollment	Configures the Web Enrollment role service on a computer where the binaries are already present.
Uninstall-AdcsCertificationAuthority	Removes the CA role service and deletes associated configuration information.
Uninstall-AdcsEnrollmentPolicyWebService	Removes the Certificate Enrollment Policy Web Service.
Uninstall-AdcsEnrollmentWebService	Removes the Certificate Enrollment Web Service, or specific instances of the service if more than one is present.
Uninstall-AdcsNetworkDeviceEnrollmentService	Removes the Network Device Enrollment service from a computer it has been deployed on.
Uninstall-AdcsOnlineResponder	Removes the Online Responder service from a computer it has been deployed on.
Uninstall-AdcsWebEnrollment	Removes the Web Enrollment role service from a computer on which it has been deployed.

Managing Certificate Services using

Certutil.exe and Certreq.exe

Not all Certificate Services functionality is available through PowerShell in Windows Server 2019. There are still several tasks that you need to use the *certutil.exe* and *certreq.exe* command-line tools to perform.

Certutil.exe

Certutil.exe allows you to manipulate local certificates as well as to configure and manage Certificate Services. Some *certutil.exe* functionality is present in available PowerShell cmdlets. Over time, it is likely that all functionality that is present in *certutil.exe* will be available through PowerShell.

Important *certutil.exe* options are listed in Table 12-4.

Table 12-4 Certutil.exe options

Option	Function
-deny	Use this to deny a pending certificate request.
-resubmit	Resubmit a pending certificate request.
-revoke	Revoke a certificate.
-CRL	Publish a new CRL.
-renewCert	Renew the CA certificate.

Technet24

-backup	Backup the AD CS.
-backupDB	Backup the AD CS database.
-backupKey	Backup the CA certificate and private key.
-restore	Restore AD CS from a backup.
-restoreDB	Restore the AD CS database from a backup.
-restoreKey	Restore the CA certificate and private key from a backup.
-dsPublish	Publish a certificate or CRL to AD.
-ADTemplate	List templates stored in AD.
-TemplateCAs	List CAs configured to issue a template.

More Info

Certutil.exe

You can find out more about the certutil.exe command at https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/cc732443(v=ws.11).

Certreq.exe

You can use *Certreq.exe* to import a certificate request file and submit it to a CA. If the user running *certreq.exe* has the appropriate permissions on the CA, they can

save the issued certificate. Important *certreq.exe* options are listed in Table 12-5.

Table 12-5 Certreq.exe options

Option	Function
-Submit	Submits a request to a CA.
-Retrieve	Retrieves a response from a CA.
-Enroll	Enrolls for a new certificate or renews a certificate.
-Accept	Accepts and installs a certificate request response.

More Info

Certreq.exe

You can find out more about the certreq.exe command at https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/dn296456(v=ws.11).

Chapter 13

Active Directory Federation Services

AD FS components

Claims, claim rules, and attribute stores

Claims provider

Relying party

Relying party trust

Claims provider trust

Configuring certificate relationship

Attribute stores

Claim rules

Configure Web Application Proxy

Workplace Join

Multifactor authentication

Managing AD FS with PowerShell

Managing Web Application Proxy with PowerShell

Active Directory Federation Services (AD FS) is Microsoft's identity federation solution. Identity federation allows identification, authorization, and authentication to occur across organizational boundaries. When a federated trust is established, users can use their local credentials to access resources hosted by another organization. When configured properly, federation also enables users to access resources hosted in the cloud. Although it is possible to configure full-forest or domain-trust relationships when users in one organization need to use their credentials to access resources in another organization, these trust relationships are often more comprehensive than is necessary. AD FS makes it possible to configure highly restricted access to information and resources between partner organizations while still allowing each partner to authenticate using its own credentials.

AD FS COMPONENTS

AD FS deployments in Windows Server 2019 consist of two components:

- **Federation Server.** The computer that hosts the federation server role manages requests involving identity claims. There must be at least one federation server in an Active Directory forest when deploying AD FS.

- **Web Application Proxy.** When you want to provide AD FS functionality to clients on untrusted networks, such as on the Internet, you deploy the computer that functions as a web application proxy on a perimeter network. This server relays connections to the Federation Server on the internal network, and this role can also be installed on a standalone computer. In Windows Server 2012 and earlier versions of AD FS, this role was known as the Federation Proxy.

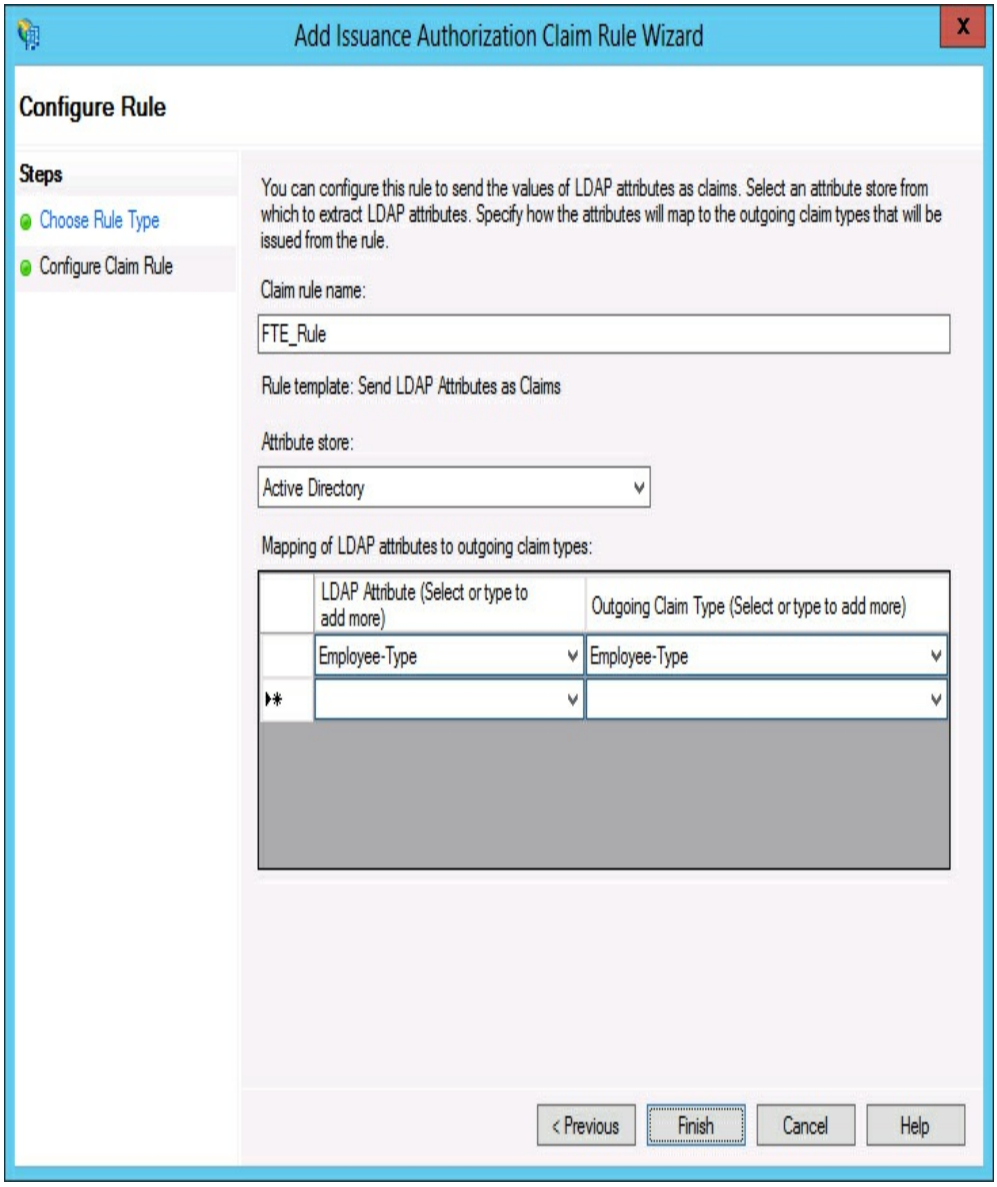
CLAIMS, CLAIM RULES, AND ATTRIBUTE STORES

AD FS provides claims-based authentication. Claims-based authentication works on the basis of a claim about the user, such as "Allow access to this web application if this user is a full-time employee of the partner organization." AD FS uses the following when building tokens that contain claim data:

- **Claim.** Claims are descriptions made about an object based on the object's attributes. For example, a user account name, employee type, or a security group membership can constitute a claim. Claims are also used with Dynamic Access Control, which is covered in Chapter 14, "Dynamic Access Control and Active Directory Rights Management Services."
- **Claim rules.** Claim rules determine how a federation server processes a claim. You can have a simple rule, such as treating a user's email address as a valid claim. Claim rules can also be more complex where an attribute, such as job title from one organization, is translated into a security group membership in the partner organization.
- **Attribute store.** An attribute store holds the values used in claims. AD FS generally uses Active Directory as an attribute store. For

example, a lookup of a user's Managed By attribute returns a value that lists the user account assigned as that user's manager.

Figure 13-1 shows how to create a claim rule where Active Directory functions as the attribute store and the Employee-Type attribute is mapped to an outgoing claim.



Technet24

Figure 13-1 Claim Rule Name

CLAIMS PROVIDER

A claims provider is a federation server that provides users with claims. These claims are stored within digitally encrypted and signed tokens. When a user needs a token, the claims provider server contacts the Active Directory deployment in its native forest to determine whether the user has authenticated. If the user is properly authenticated against the local Active Directory deployment, the claims provider then builds a user claim by using attributes located within Active Directory and within other attribute stores. The attributes that are added to the claim are dependent on the attributes required by the partner.

RELYING PARTY

The relying party server is a member of the Active Directory forest that hosts the resources that the user in the partner organization wants to access. The relying party server accepts and validates the claims contained in the token issued by the claims provider. The relying party server then issues a new token that the resource uses to determine what access to grant the user from the partner organization.

Inside OUT

Non-exclusive roles

It is possible for a single AD FS server to function as both a claims provider server and a relying party server. This enables users in each partner organization to access resources in the other organization through a federation trust.

RELYING PARTY TRUST

You configure the relying party trust on the AD FS server that functions as the claims provider server. Although this might initially seem counterintuitive, it makes sense when you consider it as a statement: “A relying party trust means that a claims provider trusts a specific relying party.” Figure 13-2 shows the Add Relying Party Trust Wizard where the address of the relying party server to be trusted is adl-dc.wingtiptoys.internal.

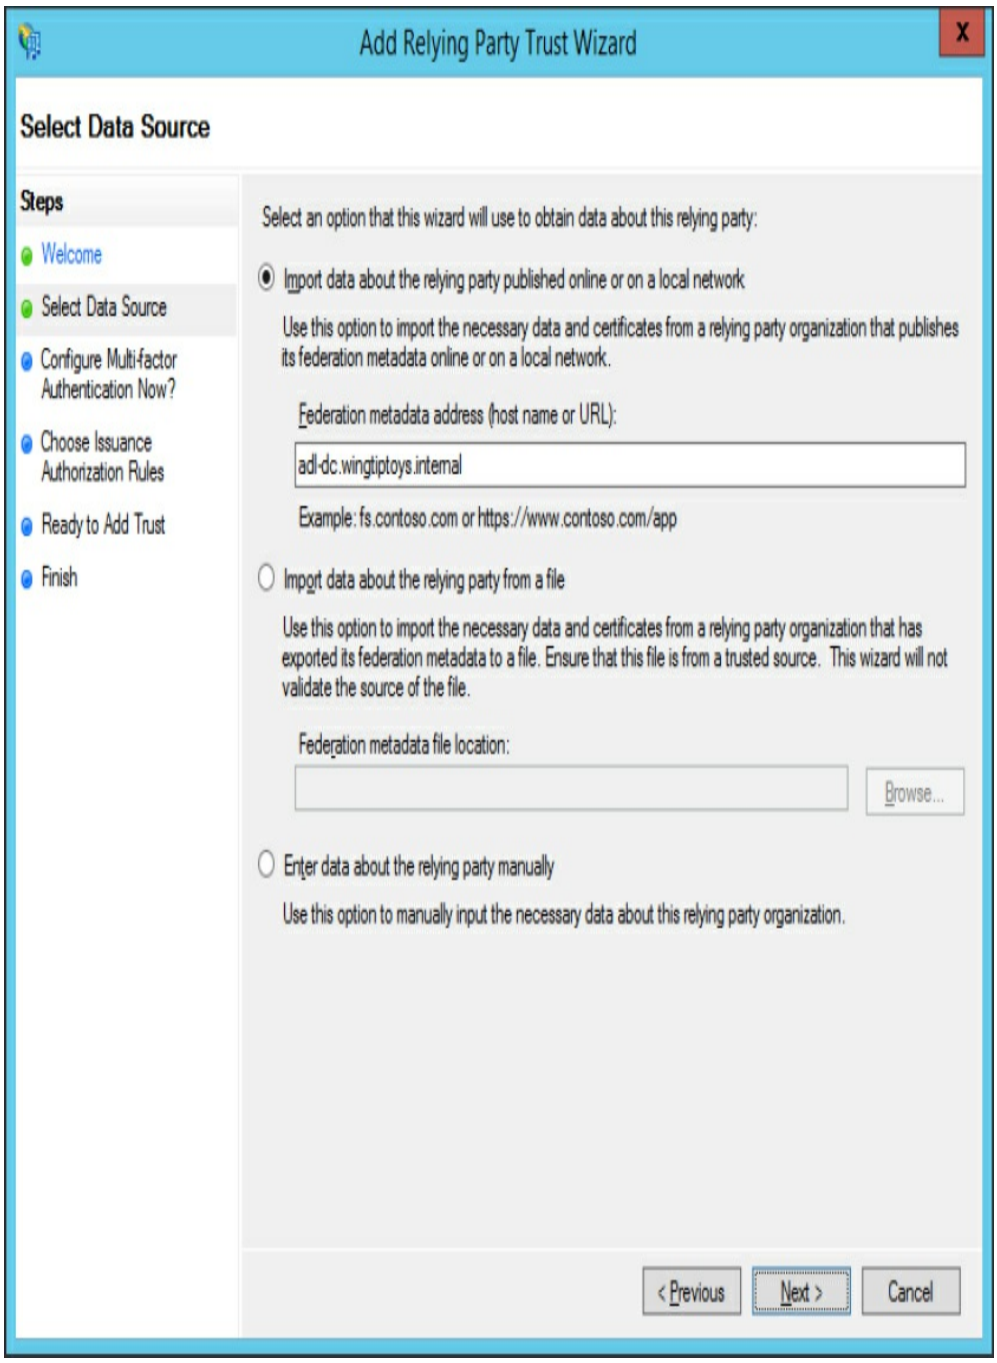


Figure 13-2 Select Data Source for relying party

CLAIMS PROVIDER TRUST

You configure the claims provider trust on the Federation Server that functions as the relying party. This also seems counterintuitive until you consider the claims provider trust as a statement: “A claims provider trust means that a relying party trusts a specific claims provider.” Figure 13-3 shows how to set up a claims provider trust where the address of the claims provider to be trusted is cbr-dc.contoso.com.

Technet24

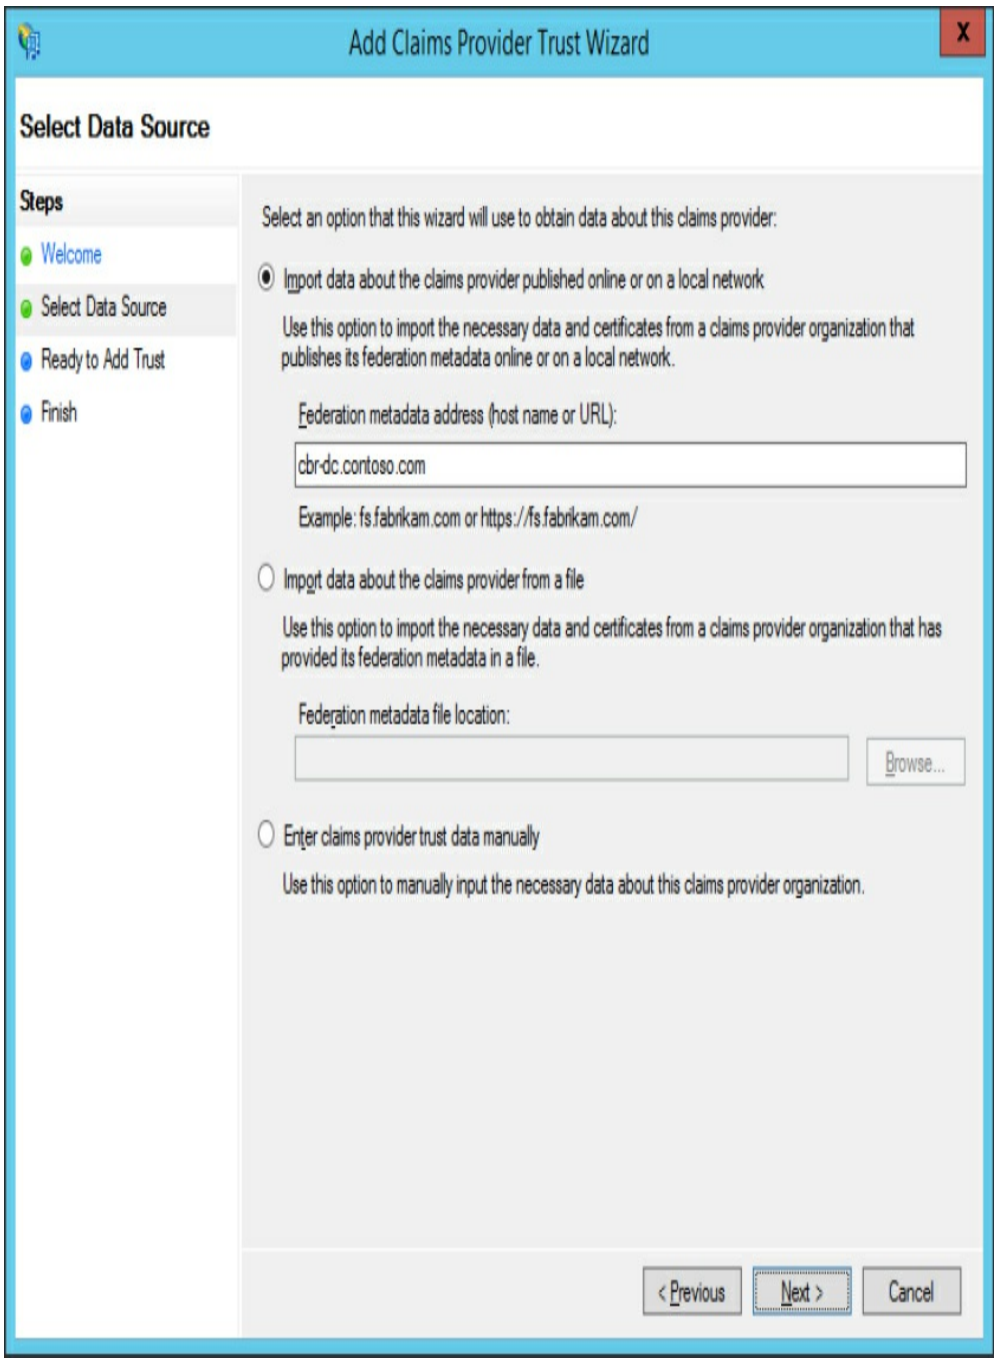


Figure 13-3 Select Data Source for claims provider

CONFIGURING CERTIFICATE

RELATIONSHIP

You need to configure partners so that the certificates issued by the opposite partner are trusted. You can accomplish this by using a trusted third-party CA or by configuring CA trusts between partners. To configure a CA trust, you need to import the CA certificate of the partner organization's CA into the Trusted Root CA store. You can do this either directly on the AD FS server, or you can configure Group Policy so that all computers within the policy's scope trust the partner organization's CA.

You can use a certificate issued from the computer certificate template, which is available through an Active Directory Certificate Services enterprise CA, to secure the federation server endpoint. When using a trusted third-party CA, you can also use a typical Transport Layer Security (TLS) certificate.

AD FS uses the following certificates:

- **Token-signing certificates.** The federation server uses the token-signing certificate to sign all the tokens that it issues. The server that functions as the claims provider uses the token-signing certificate to verify its identity. The relying party uses the token-signing certificate when verifying that the token was issued by a trusted federation partner.
- **Token-decrypting certificates.** The claims provider uses the public key from this certificate to encrypt the user token. When the relying party server receives the token, it uses the private key to decrypt the user token.

ATTRIBUTE STORES

The attribute store holds information about users. The AD FS server uses the information contained in the attribute store to build claims after the user is authenticated. In the majority of AD FS implementations, Active Directory functions as the attribute store. You can configure additional attribute stores based on the following technologies:

- Active Directory Lightweight Directory Services (AD LDS) on computers running Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019
- All supported editions of Microsoft SQL Server 2012, SQL Server 2014, SQL Server 2016, SQL Server 2017, and SQL Server 2019

You add attribute stores to AD FS in the Add An Attribute Store dialog box, as shown in Figure 13-4.

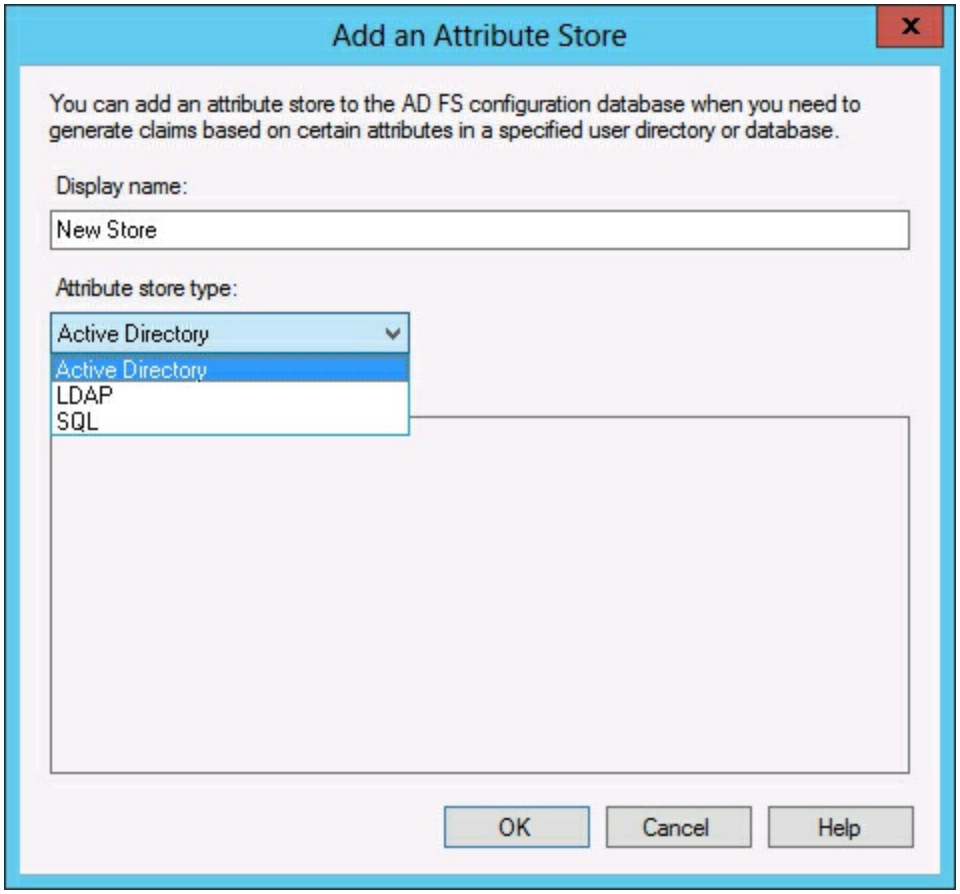


Figure 13-4 Add An Attribute Store

CLAIM RULES

Claim rules determine how AD FS servers send and consume claims. AD FS supports two different types of claim rules: relying party trust claim rules and claims provider trust claim rules.

Relying party trust claim rules

Relying party trust claim rules determine how the claims about a user are forwarded to the relying party.

You can configure these rules by editing the relying party trust's claim rules on the AD FS server that functions as the claims provider. There are three types of relying party trust claim rules:

- **Issuance transform rules.** Determine how claims are sent to the relying party.
- **Issuance authorization rules.** Determine which users have access to the relying party. An issuance authorization claim rule is shown in Figure 13-5.
- **Delegation authorization rules.** Determine if users can act on behalf of other users when accessing the relying party.

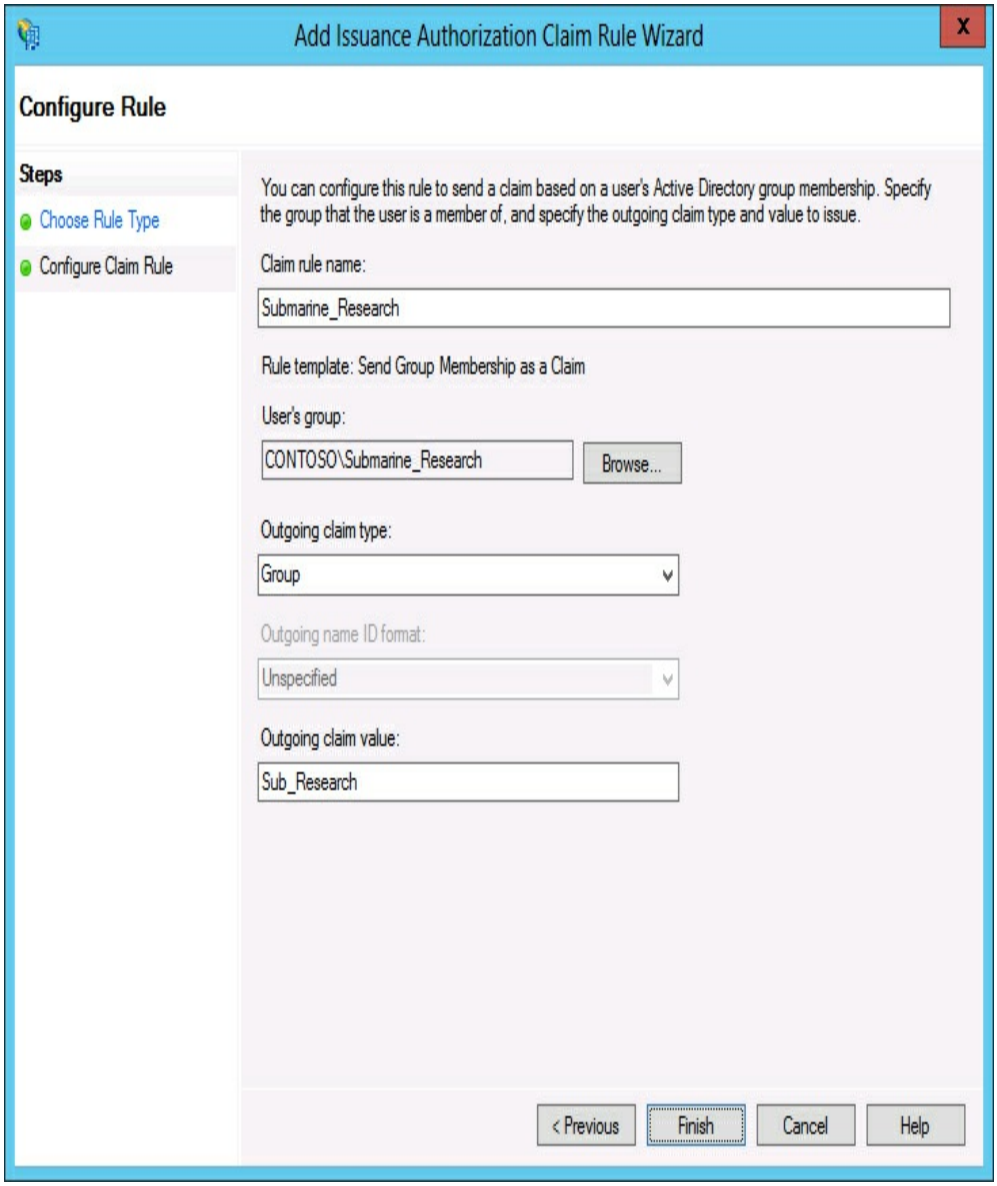


Figure 13-5 Choose a rule type

Claims provider trust claim rules

Claims provider trust rules determine how the relying party filters incoming claims. You can configure these rules by editing the claims provider trust's claim rules.

All claims provider trust claim rules are acceptance transform rules. Figure 13-6 shows an incoming claim type related to group membership.

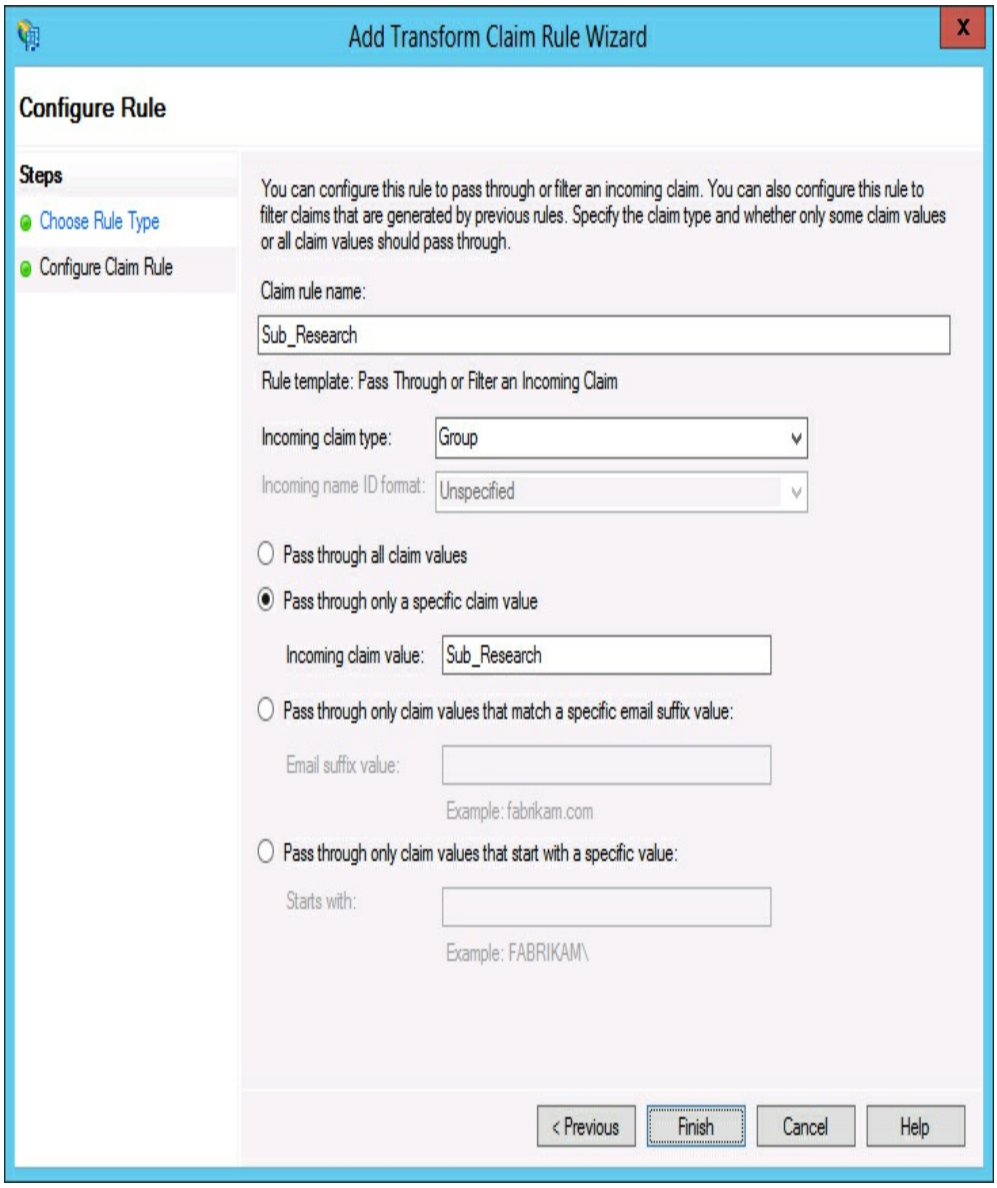


Figure 13-6 Claim type related to group membership

CONFIGURE WEB APPLICATION

PROXY

Web application proxy servers can be deployed on perimeter networks as a way to increase security for AD FS deployments. Clients communicate with the AD FS proxy server, which then communicates with federation servers on the internal network. For example, on the claims provider side of the federated trust, the web application proxy forwards authentication data from the client to the AD FS server. The AD FS server confirms the authentication and issues the token. The token is then relayed through the proxy to the relying party web application proxy server. The relying party web application proxy server relays the token to the relying party AD FS server, which issues a new token that is sent back through the proxy to the original client. At no point does a web application proxy server actually create claims or generate tokens. All communication between a web application proxy server and a federation server occurs using the HTTPS protocol.

In Windows Server 2012 and earlier versions of AD FS, the web application proxy server role was performed by an AD FS proxy server. A web application proxy server in Windows Server 2016 and Windows Server 2019 differs from the AD FS proxy role service available in Windows Server 2012 because it has the following features:

- It is deployed as a Remote Access role service.

- It provides secure remote access to web-based applications hosted on the internal network.
- It functions as a reverse proxy for web-based applications.
- It can be used to publish access to Work Folders.

When deploying a web application proxy server, you need to ensure that the following certificates are present in the server's certificate store:

- A certificate that includes the federation service name. If the web application proxy server must support Workplace Join, the certificate must also support the following subject alternative names:
  - <federation servicename>.<domain name>
  - Enterpriseregistration.<domainname>
- A wildcard certificate, a subject alternative name certificate, or individual certificates to cover each web application that is accessible through the web application proxy.
- A copy of the certificate that external servers use if you are supporting client certificate preauthentication.

The server hosting the web application proxy role must also trust the certificate authority that issues these certificates. After deploying the web application proxy role service, run the Web Application Proxy Configuration Wizard. Running the wizard involves specifying the federation service name and the credentials of an account that has AD FS administrative privileges, as shown in Figure 13-7.

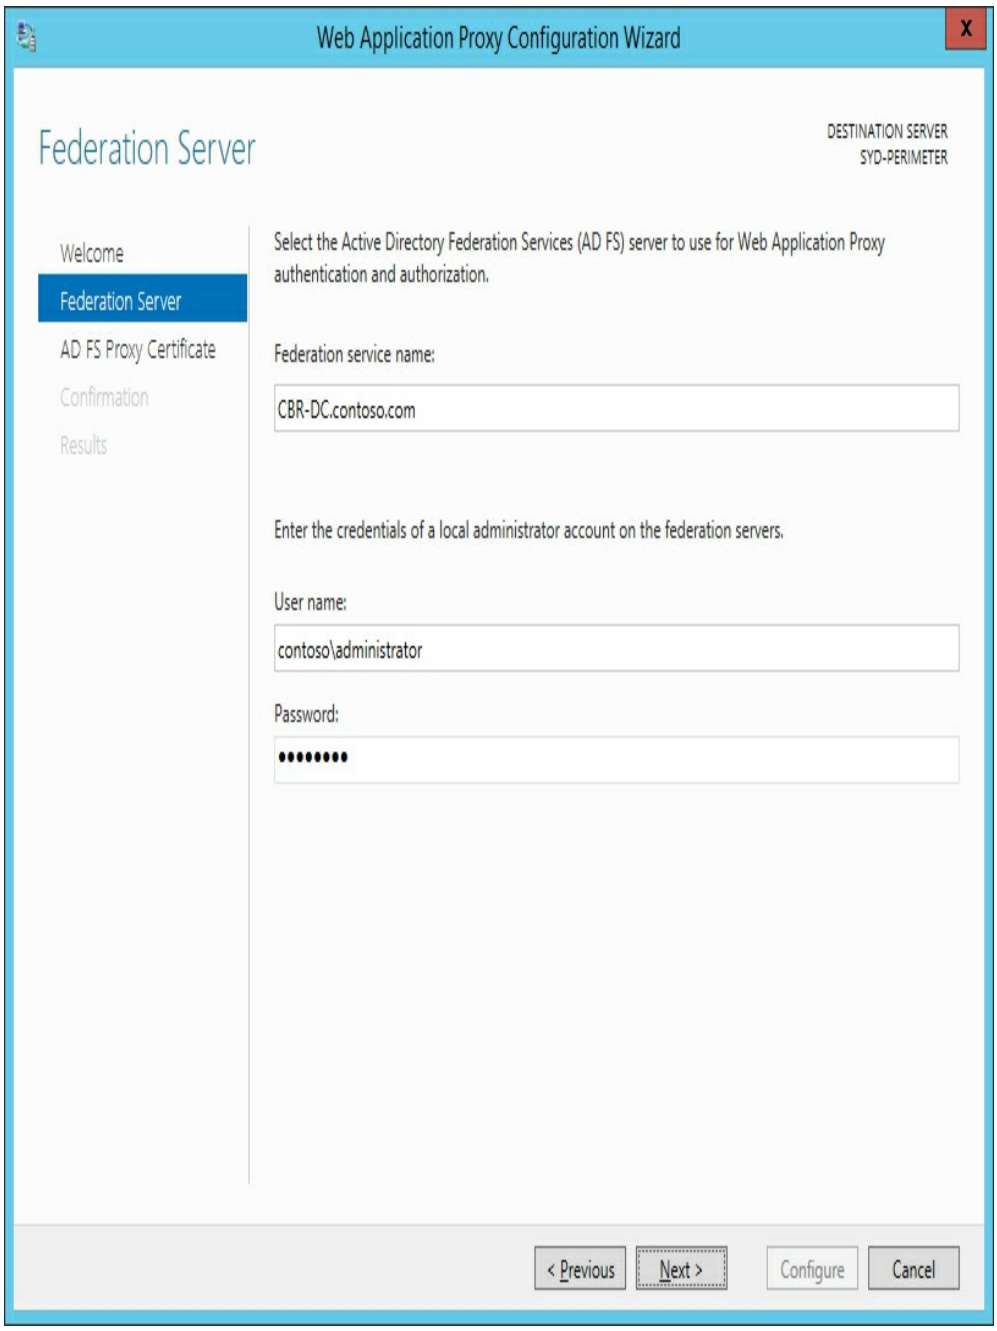


Figure 13-7 Web application proxy configuration

You also need to specify the certificate that the AD FS proxy component of the web application proxy uses, as

shown in Figure 13-8. You must have access to the certificate's private key as well as the public key when configuring the web application proxy's AD FS proxy component.

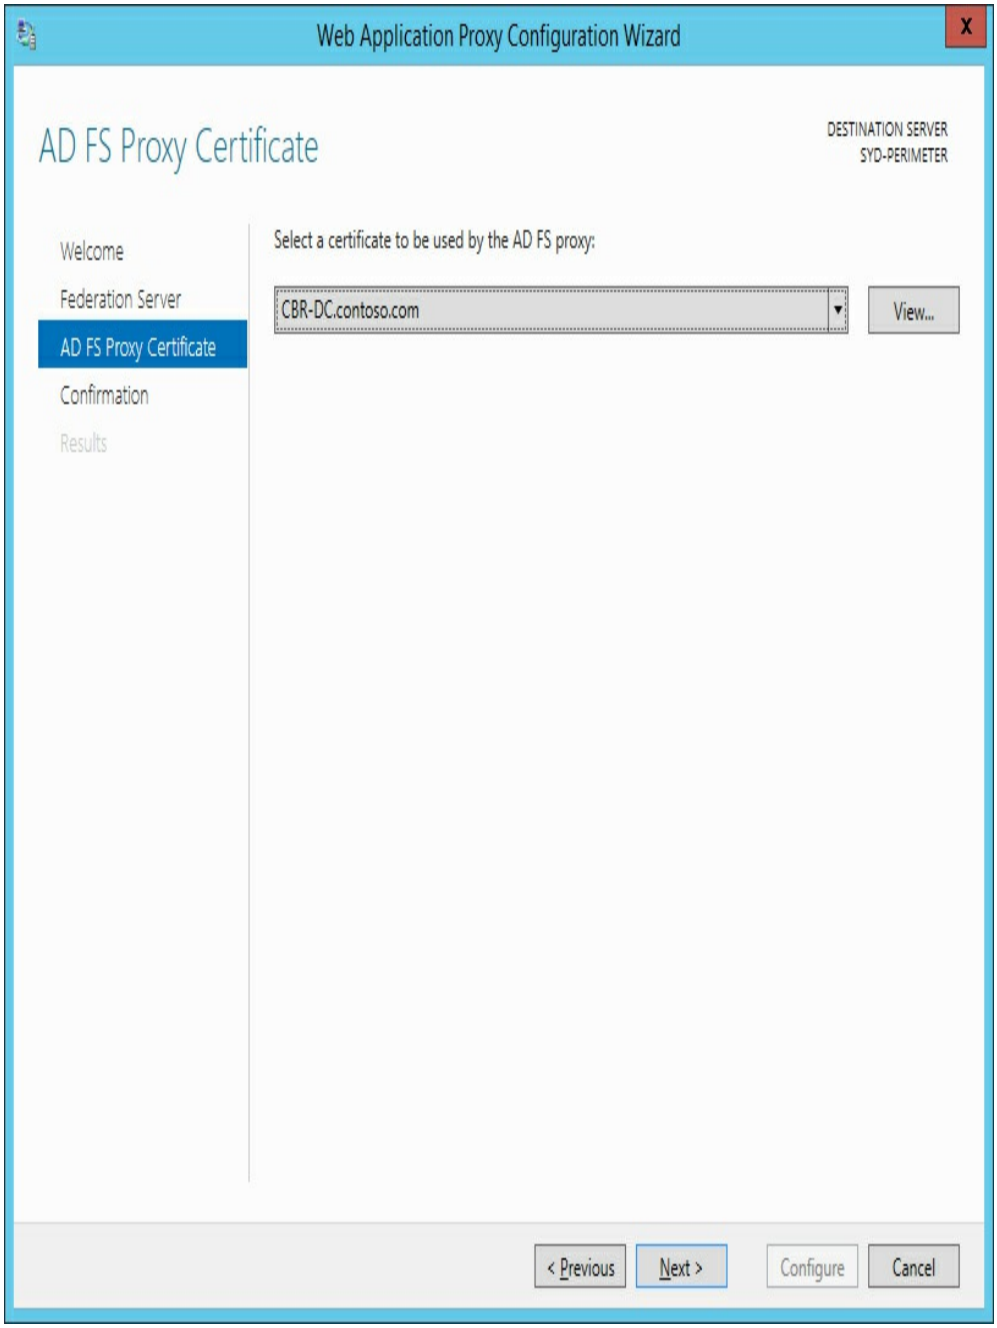


Figure 13-8 AD FS Proxy Certificate

WORKPLACE JOIN

Workplace Join is a feature available in Windows Server 2019, Windows Server 2016, and Windows Server 2012 R2 that you can use to allow non-domain joined devices and computers to access domain resources and claims-enabled applications in a secure manner. When a non-domain joined device or computer performs a Workplace Join, a special object representing the device is created in AD DS. The object stores attributes that describe the device that you can use when configuring access rules.

Workplace Join requires the following components:

- Windows 8.1, Windows 10, Android, or iOS 7 or later devices
- Claims-aware application
- AD FS deployment on internal network
- Web application proxy deployed on the perimeter network
- Device registration service

Workplace Join also supports single sign-on (SSO). When a user authenticates to access one claims-aware application, that authentication carries over to other claims-aware applications hosted by the organization as well. When supporting Workplace Join, consider acquiring certificates from trusted third-party CAs because the majority of the devices performing

Technet24

Workplace Join are managed by their owners rather than centrally by the IT department, which means that getting those devices to trust certificates from an internal CA requires extra administrative effort.

Enable Workplace Join by running the following Windows PowerShell cmdlets:

Click here to view code image

```
Initialize-ADFSDeviceRegistration
Enable-ADFSDeviceRegistration
```

After you've enabled Workplace join, you can enable device authentication by editing the Global Authentication Policy through the AD FS console, as shown in Figure 13-9.

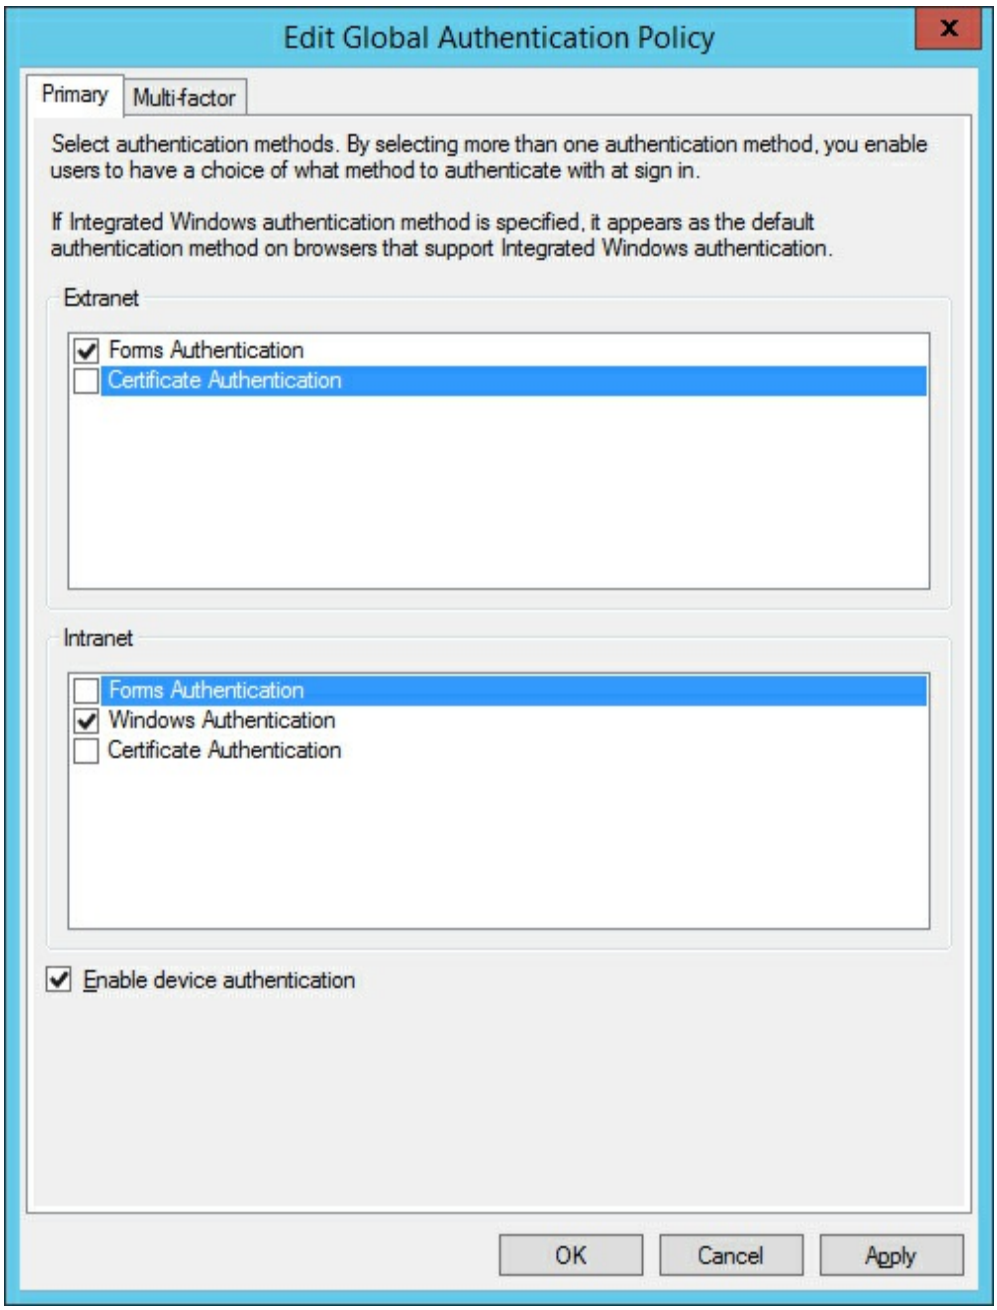


Figure 13-9 Enabling device authentication

Users authenticate by using their UPNs. When configuring Workplace Join, ensure that you configure

your organization's external DNS zone with a record that maps enterpriseregistration.upndomainname.com (where upndomainname.com is the UPN suffix) to the Web Application Proxy server's IP address or the AD FS server that you've configured to support Workplace Join.

MULTIFACTOR AUTHENTICATION

AD FS supports multifactor authentication. When you implement multifactor authentication, more than one form of authentication is required. For example, a username, password, and a code from an authenticator application running on a mobile device might be required. You can either configure multifactor authentication in AD FS globally, as shown in Figure 13-10, or on a per-relying party trust basis.

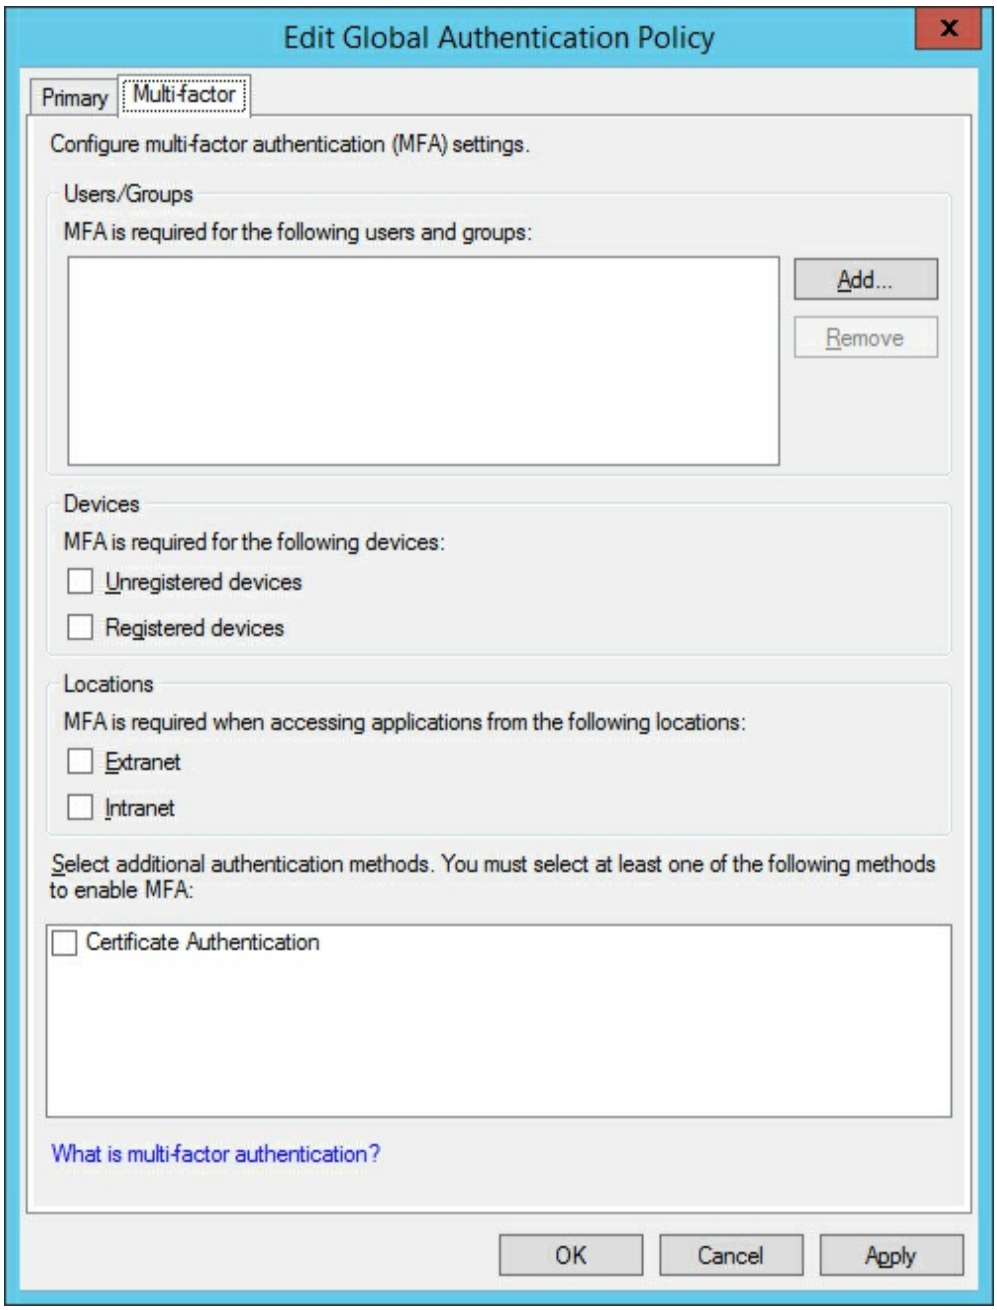


Figure 13-10 Multifactor authentication

You can configure multifactor authentication by using a third-party vendor's product, or you can integrate

Microsoft's Azure Multi-Factor authentication service. Azure Multi-Factor authentication allows the following authentication methods:

- **Telephone call.** An automated telephone call is made to the user's registered telephone number. The user enters the data provided in the call when authenticating.
- **Text message.** A text message is sent to the user's registered mobile telephone. The user enters the data provided in the text message when authenticating.
- **Mobile app.** The user installs a mobile authentication app on his or her mobile device that generates a periodically changing code. The user enters this code when authenticating.

MANAGING AD FS WITH POWERSHELL

You can use the Windows PowerShell cmdlets from the AD FS module, listed in Table 13-1, to manage AD FS on Windows Server 2019.

Table 13-1 AD FS PowerShell cmdlets.

Noun	Verbs	Function
ADDeviceRegistration	Initialize	Initializes Device Registration service config in the forest
AdfsAccessControlPolicy	Get, Remove, New, Set	Manage AD FS access control policy
AdfsAdditional	Get, Set	Manage rules that trigger

Authentication Rule		additional authentication providers to be invoked
AdfsAlternateTlsClientBinding	Set	Configure TLS certificate client binding
AdfsApplicationGroup	Get, Disable, Set, Remove, New, Enable	Manage AD FS application groups
AdfsApplicationPermission	Grant, Set, Get, Remove	Manage AD FS application permissions
AdfsAttributeStore	Set, Add, Get, Remove	Manage AD FS attribute store
AdfsAuthenticationProvider	Register, Get, Unregister	Configure AD FS authentication providers
AdfsAuthenticationProviderConfigurationData	Export, Import	Import and export AD FS authentication provider configuration data
AdfsAuthenticationProviderWebContent	Get, Remove, Set	Configure authentication provider for web content
AdfsAzureMfaConfigured	Get	Get the Azure Multi-Factor authentication configuration
AdfsAzureMfaTenant	Set	Configure an existing Azure Multi-Factor authentication tenant

Technet24

AdfsAzureMfaTenantCertificate	New	Configure a new AD FS Azure Multi-Factor authentication tenant certificate
AdfsCertificate	Get, Update, Remove, Set, Add	Configure an AD FS certificate
AdfsCertificateAuthority	Disable, Get, Set	Configure an AD FS certificate authority
AdfsCertSharingContainer	Set	Configure an AD FS certificate sharing container
AdfsClaimDescription	Remove, Add, Get, Set	Manage claim descriptions
AdfsClaimRuleSet	New	Create a new AD FS claim rule set
AdfsClaimsProviderTrust	Set, Get, Update, Disable, Enable, Remove, Add	Configure an AD FS claims provider trust
AdfsClaimsProviderTrustsGroup	Remove, Get, Add	Configure an AD FS claims provider trusts group
AdfsClient	Remove, Set, Disable, Add,	Configure an AD FS client

	Get, Enable	
AdfsContactPerson	New	Configure a new AD FS contact person
AdfsDeploymentSQLScript	Export	Export an SQL script for AD FS deployment
AdfsDeviceRegistration	Enable, Set, Disable, Get	Configure AD FS device registration
AdfsDeviceRegistrationUpnSuffix	Get, Set, Remove, Add	Configure an AD FS device registration UPN suffix
AdfsEndpoint	Disable, Set, Get, Enable	Manage an AD FS endpoint
AdfsFarm	Install	Install an AD FS farm
AdfsFarmBehaviorLevel	Restore	Restore a farm behavior level to a previous level
AdfsFarmBehaviorLevelRaise	Test, Invoke	Test or perform a farm behavior level raise
AdfsFarmBehaviorLevelRestore	Test	Check whether prerequisites have been met to restore the existing AD FS farm behavior level
AdfsFarmInformation	Set, Get	View and configure information about the AD FS farm

Technet24

AdfsFarmInstallation	Test	Performs checks to test whether you can install a new federation server farm
AdfsFarmJoin	Test	Performs checks required before you run the Add-AdfsFarmNode cmdlet to add a computer to a farm
AdfsFarmNode	Remove, Add	Manage AD FS farm nodes
AdfsGlobalAuthenticationPolicy	Set, Get	Configure AD FS global authentication policy
AdfsGlobalWebContent	Get, Set, Remove	Configure AD FS global web content
AdfsLdapAttributeToClaimMapping	New	Configure AD FS attribute claim mapping
AdfsLdapServerConnection	New	Create a connection to an LDAP server for AD FS
AdfsLocalClaimsProviderTrust	Get, Enable, Disable, Add, Set, Remove	Configure AD FS Local Claims Provider Trust
AdfsNativeClientApplication	Remove, Get, Add, Set	Configure AD FS client application
AdfsNonClaim	Set, Enable,	Configure non-claims aware

sAwareRelyingPartyTrust	Disable, Add, Remove, Get	relying party trust
AdfsOrganization	New	Create a new AD FS organization
AdfsProperties	Set, Get	Configure AD FS properties
AdfsProxyTrust	Revoke	Remove the proxy trust
AdfsRegistrationHosts	Get, Set	Configure AD FS registration hosts
AdfsRelyingPartyTrust	Set, Add, Enable, Update, Disable, Get, Remove	Configure relying party trusts
AdfsRelyingPartyTrustsGroup	Remove, Add, Get	Configure relying party trust group
AdfsRelyingPartyWebContent	Set, Remove, Get	Manage web content for the relying party
AdfsRelyingPartyWebTheme	Get, Set, Remove	Configure the web theme for the relying party
AdfsSamlEndpoint	New	Create a new AD FS SAML endpoint

AdfsScopeDescription	Add, Remove, Set, Get	Manage AD FS scope description
AdfsServerApplication	Remove, Set, Get, Add	Manage AD FS server applications
AdfsSslCertificate	Set, Get	Configure AD FS TLS certificates
AdfsSyncProperties	Set, Get	Manage AD FS sync properties
AdfsTrustedFederationPartner	Get, Remove, Add, Set	Manage AD FS trusted federation partners
AdfsWebApiApplication	Get, Remove, Set, Add	Manage AD FS web API applications
AdfsWebApplicationProxyRelyingPartyTrust	Add, Set, Disable, Get, Remove, Enable	Manage relying party trust for Web Application Proxy
AdfsWebConfig	Get, Set	Configure AD FS web customization configuration settings
AdfsWebContent	Import, Export	Import and export properties of all web content objects

AdfsWebTheme	New, Get, Export, Set, Remove	Manage AD FS web themes

MANAGING WEB APPLICATION PROXY WITH POWERSHELL

You can use the Windows PowerShell cmdlets from the WebApplicationProxy module, listed in Table 13-2, to manage AD FS on Windows Server 2019.

Table 13-2 Web Application Proxy PowerShell cmdlets.

Noun	Verbs	Function
WebApplicationProxy	Install	Install the Web Application Proxy role
WebApplicationProxyApplication	Remove, Set, Add, Get	Configure Web Application Proxy applications
WebApplicationProxyAvailableADFSRelyingParty	Get	View Web Application Proxy available in AD FS relying party configuration information
WebApplicationProxyConfiguration	Get, Set	View and configure Web Application Proxy configuration information
WebApplicationP	Updat	Configure AD FS SSL endpoint listeners

Technet24

roxyDeviceRegistration	e	that correspond to UPN suffixes
WebApplicationProxyHealth	Get	View health status of the Web Application Proxy
WebApplicationProxySslCertificate	Get, Set	Configure the AD FS TLS certificate for the federation server proxy component of the Web Application Proxy

Chapter 14

Dynamic Access Control and Active Directory Rights Management Services

Dynamic Access Control

Configuring Group Policy to support DAC

Configuring User and Device Claims

Configuring Resource Properties

Central access rules

Central access policies

Staging

Access Denied Assistance

Installing AD RMS

AD RMS certificates and licenses

Technet24

- AD RMS Templates
- AD RMS Administrators and Super Users
- Trusted User and Publishing Domains
- Exclusion policies
- Managing AD RMS with Windows PowerShell

Dynamic Access Control (DAC) provides a way to dynamically assign access permissions to content based on the content's properties and information about the user and device that are attempting to access the content. Active Directory Rights Management Services (AD RMS) provides a way of securing content through encryption and through rules that are applied to the operating system and applications on what actions the user can perform with that content. Both technologies, if correctly implemented, can minimize the chance of information that should stay within organizational boundaries from finding its way outside of those boundaries.

DYNAMIC ACCESS CONTROL

If you've worked with NTFS permissions for a while, you are most likely aware that they are rarely properly implemented. At first glance, the system seems logical. You create a collection of groups that represent ways of

describing a user or computer's place in the organization. You then use those groups to apply permissions to restrict access to files and folders. This is great in theory, but it requires that security groups are kept up to date and that the permissions themselves are accurately configured. In many organizations, the process of ensuring group membership is erratic and piecemeal, with users only being added to groups after they lodge service desk tickets.

Dynamic Access Control (DAC) allows you to configure security using the properties of a user account or a computer account and the properties of the file. For example, you can configure DAC so that only people who have Don Funk as a manager are able to open files that contain the word Cake. You can do this by setting up classification rules and claims so that every file is checked to see if it contains the word Cake, and if it does, a custom attribute is configured to reflect that status. Another rule can be configured to set permissions on the file so that anyone whose Active Directory user account has the Managed By attribute set to Don Funk has access to open the file. Access is still mediated by NTFS permissions, but those NTFS permissions are configured based on the properties of the file and the accessing user, not using the traditional method of right-clicking the file or parent folder and setting permissions manually.

DAC has the following requirements:

- Windows Server 2019, Windows Server 2016, Windows Server 2012 R2, or Windows Server 2012 with the File Server Resource Manager (FSRM) role service installed on the file servers that host files protected through DAC. DAC is not supported on Windows Server 2008 R2 and earlier file servers.
- Windows 10, Windows 8.1, or Windows 8 client computers to support device claims. Clients using Windows 7 to access files that have security applied through DAC are still able to access files, but device claims are ignored.
- A domain controller running Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, or Windows Server 2019 in the domain. It's also possible to apply an update to the Active Directory Schema if no Windows Server 2012 domain controllers are present, and all domain controllers are running Windows Server 2008 R2.

CONFIGURING GROUP POLICY TO SUPPORT DAC

Configuring Active Directory to support DAC requires that you configure a Group Policy that applies to all domain controllers in the domain. You can do this in a policy that applies to the domain or just to domain controllers. To enable support for DAC, enable the KDC Support For Claims, Compound Authentication And Kerberos Armoring policy. You can find this policy in the Computer Configuration\Policies\Administrative Templates\System\KDC node.

CONFIGURING USER AND DEVICE CLAIMS

Before you can configure access rules, you need to configure claims. Claims are bits of information about users and computers and are usually derived from Active Directory attributes. You can edit Active Directory attributes by using Active Directory Administrative Center as well as by using other tools. DAC supports the following types of claims:

- **User Claim.** This is information about the user and can be based on a user account's attributes, such as EmployeeType. You can also use claims related to group membership.
- **Device Claim.** This is information about the computer that the user is accessing the file from. For example, you could edit the computer account's Location attribute and set it to Secure. This allows you to configure DAC so a user can access a file from a computer that has the Location attribute set to Secure, but cannot access the same file from another computer that does not have the Location attribute set to Secure.

You create claims in Active Directory Administrative Center by navigating to the Claim Types section under the Dynamic Access Control section. When creating a claim type, select an existing Active Directory attribute as the basis for the claim. In Figure 14-1, you can see that the Department attribute forms the basis of a claim type. When you configure a claim, select whether the claim relates to Users or Computers. You can also specify suggested values to associate with the claim. In

Technet24

the case of a claim related to the Department attribute, this might be a list of departments within your organization.

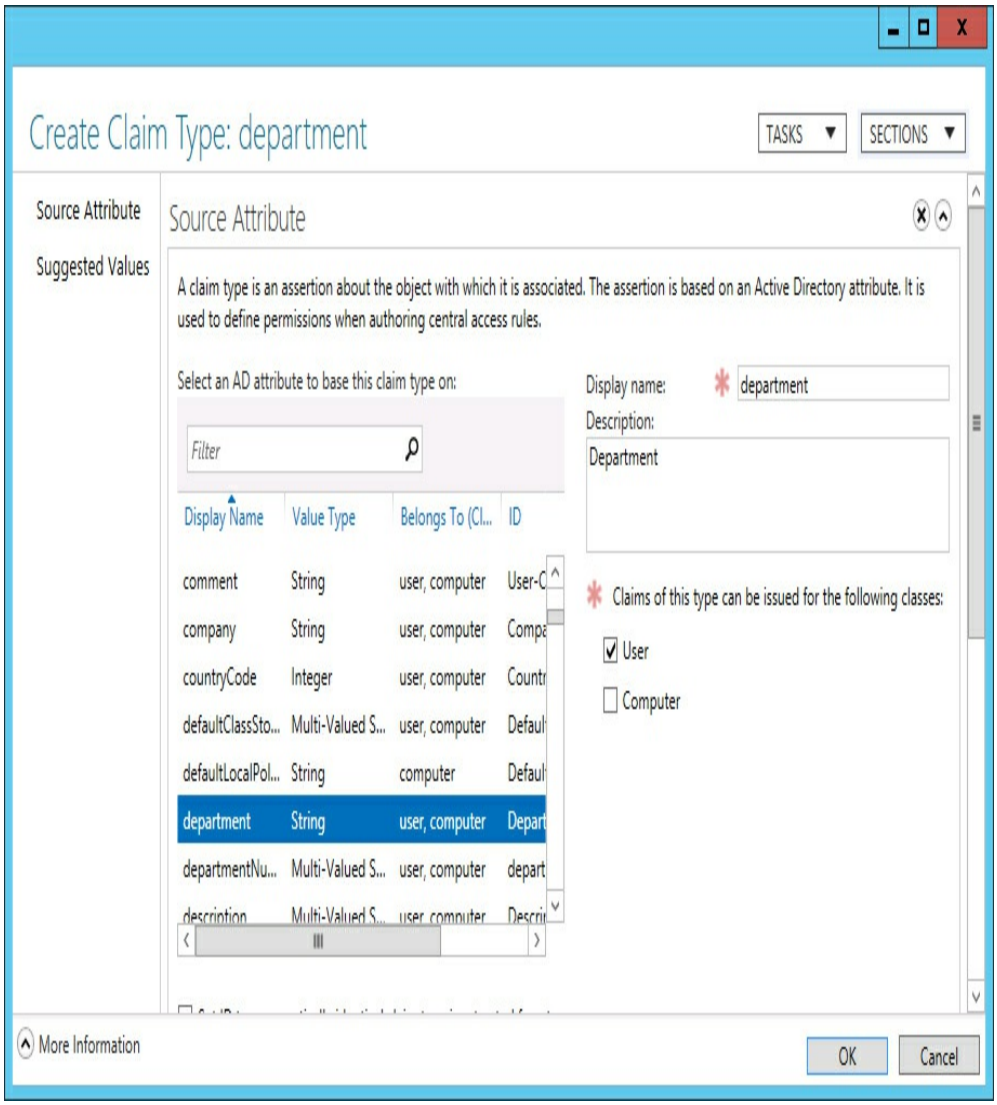


Figure 14-1 Create Claim Type: Department

CONFIGURING RESOURCE PROPERTIES

Resource Properties determine the resource attributes that you can use when configuring central access rules. Windows Server 2019 ships with a collection of default resource properties, but you can also add extra resource properties based on available attributes. Figure 14-2 shows a custom resource property named Project and two value options: Hovercraft or Submarine. Resource Properties can be assigned to files either manually or automatically by configuring File Server Resource Manager File Classification Rules.

Technet24

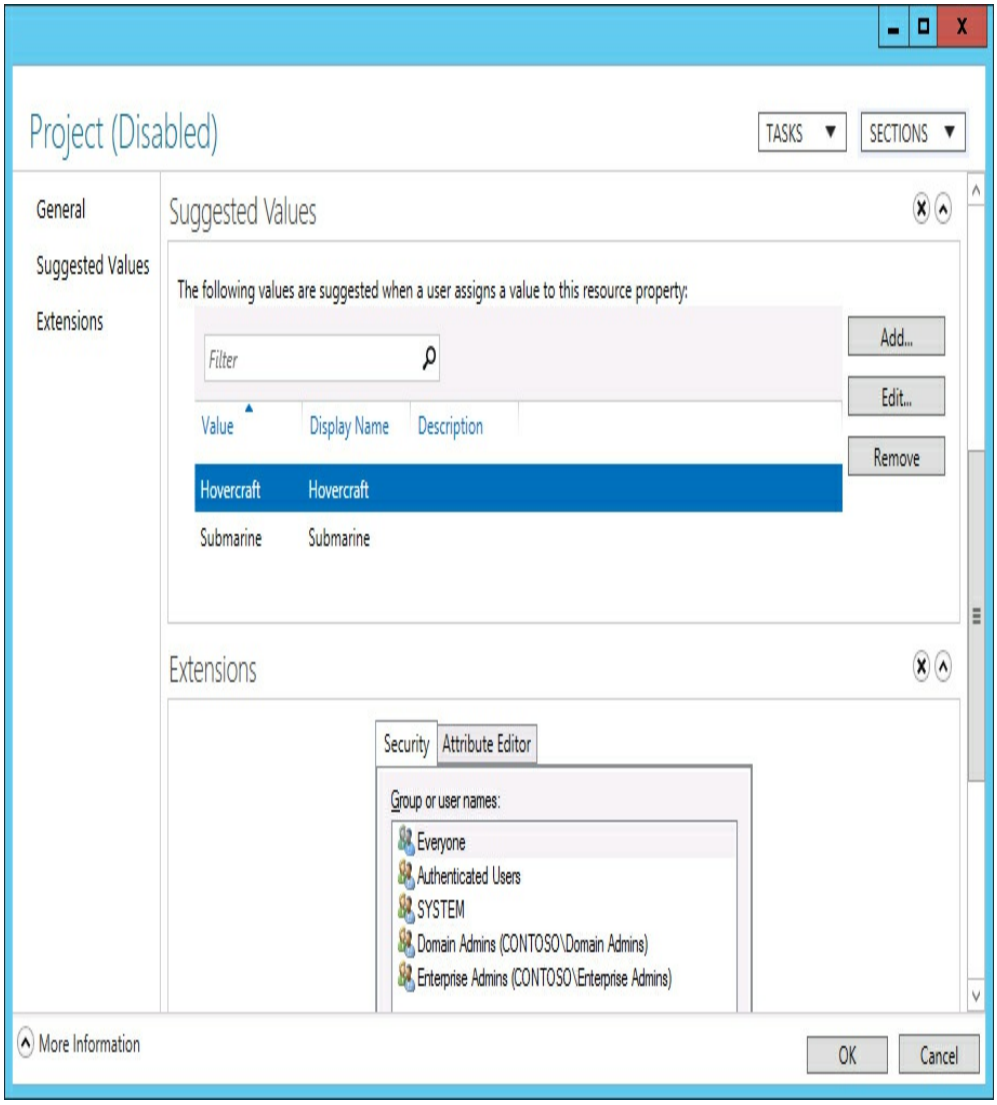


Figure 14-2 Hovercraft project

The Global Resource Property List, shown in Figure 14-3, is a list of all resource properties that you can use when configuring Central Access Rules. You can add and remove these properties as necessary. When you publish a property list through Active Directory, you can then assign these properties to files and folders either

manually or automatically using File Server Resource Manager.

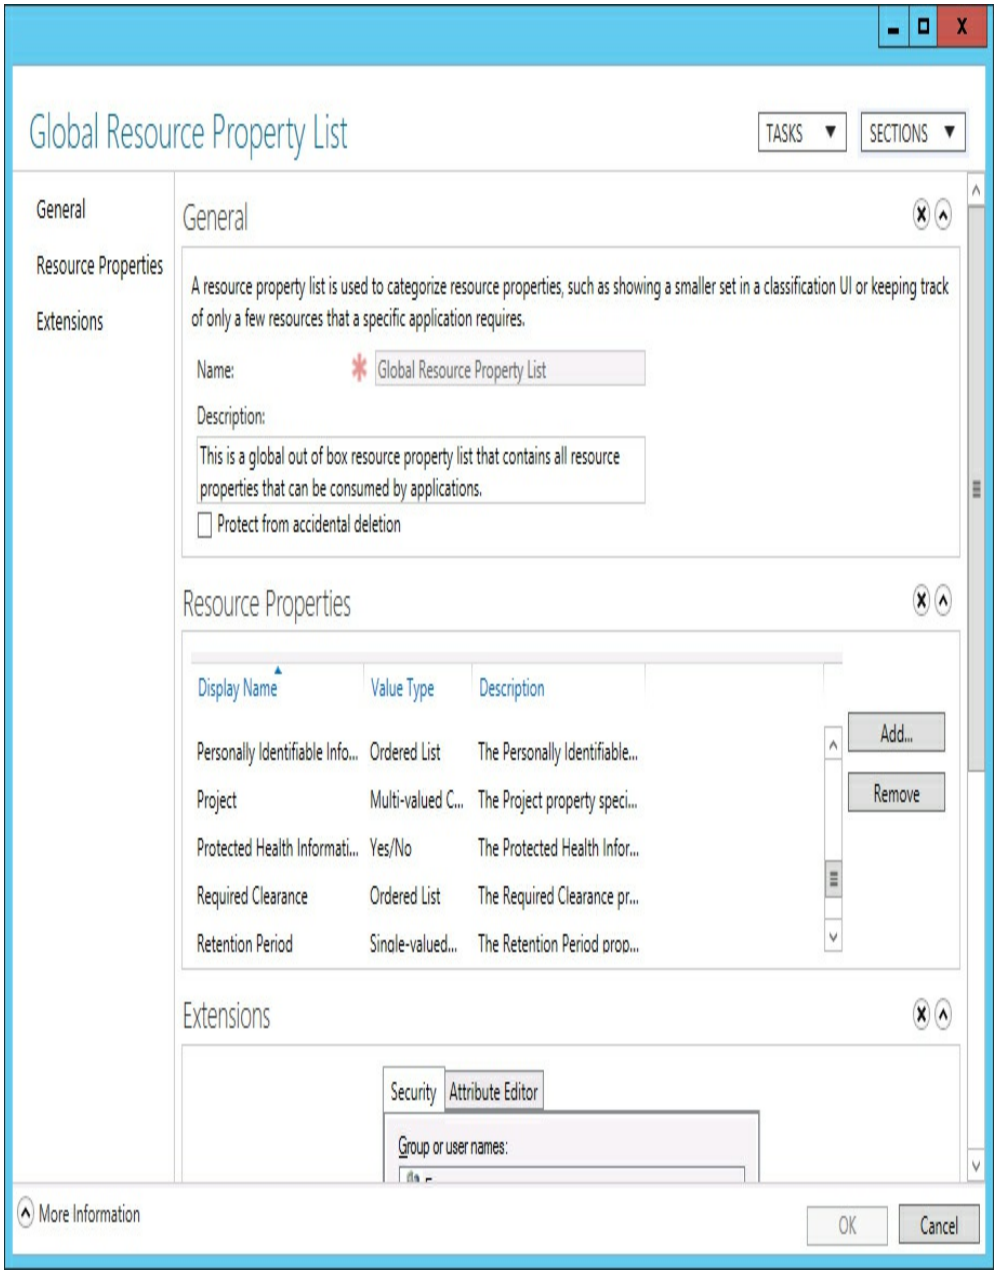


Figure 14-3 Global Resource Property List

You use File Server Resource Manager to apply

properties to files through File Classification Rules. You can configure a classification rule to look for a particular string of text in a file and then to assign a particular property to that file based on the string. The example in Figure 14-4 shows that the Submarine value is assigned to the Project property for files that meet the classification requirements, the classification requirements in this case being any file that contains the text string submarine.

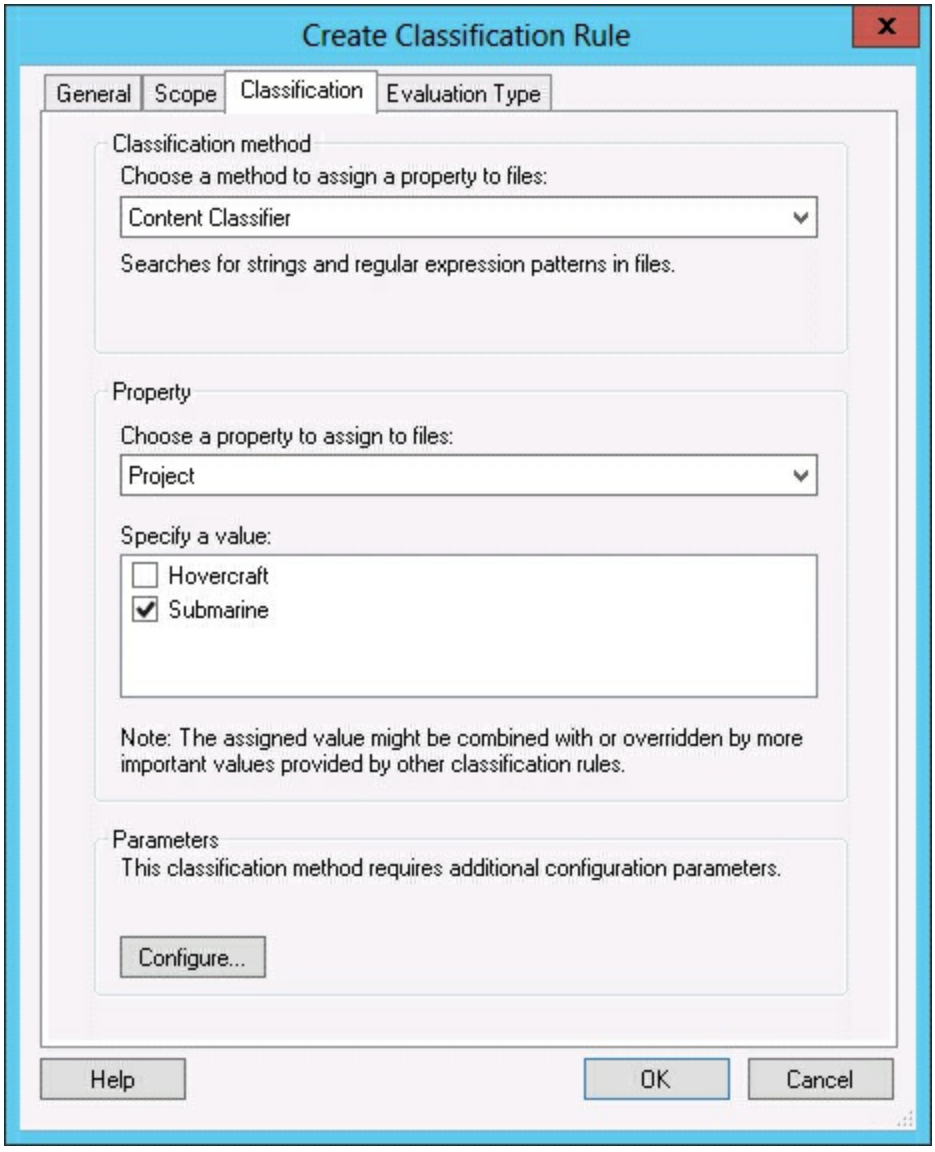


Figure 14-4 Create Classification Rule

CENTRAL ACCESS RULES

Central access rules include a set of permissions and the conditions under which those permissions are applied. For example, in Figure 14-5, the rule applying the permissions spelled out in the permissions entry is

applied to the file or folder if the file or folder has the Project resource property set to Hovercraft. You can have multiple conditions in a central access rule. You can, for example, require that the Project resource property be set to Hovercraft and the Confidentiality resource property be set to High for the permissions configured in the permissions entry to be applied.

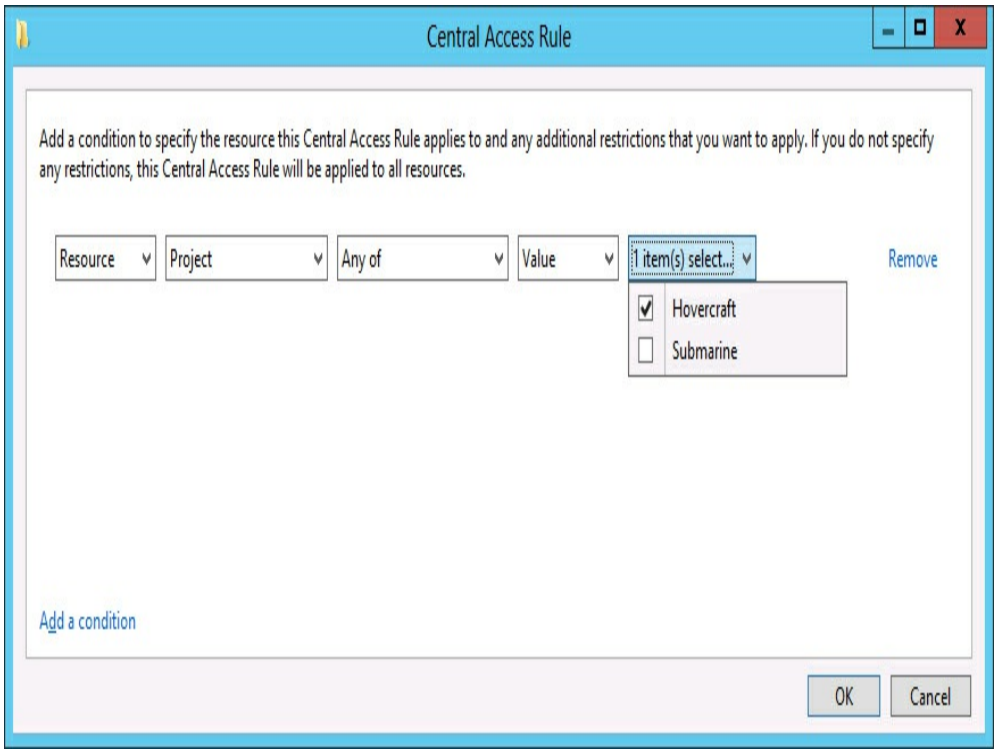


Figure 14-5 Central Access Rule

After you configure the conditions that trigger the Central Access Rule, you can then specify the set of permissions that is applied. Unlike standard NTFS permissions, permission entries allow you to apply permissions that are conditional upon user and device

claims. For example, the permissions entry shown in Figure 14-6 is conditional upon the user attempting access to not only be a member of the Hovercraft_Project security group, but also have the EmployeeType attribute in their user account set to the value FTE. If the user's EmployeeType attribute isn't set to FTE, the permissions assigned through the user's Hovercraft_Project group membership are not granted. You can configure multiple conditions based on user and device claims when configuring a permissions entry. For example, in the case of sensitive documents you might also require that the computer account have an attribute set to indicate that it is a secure computer.

Technet24

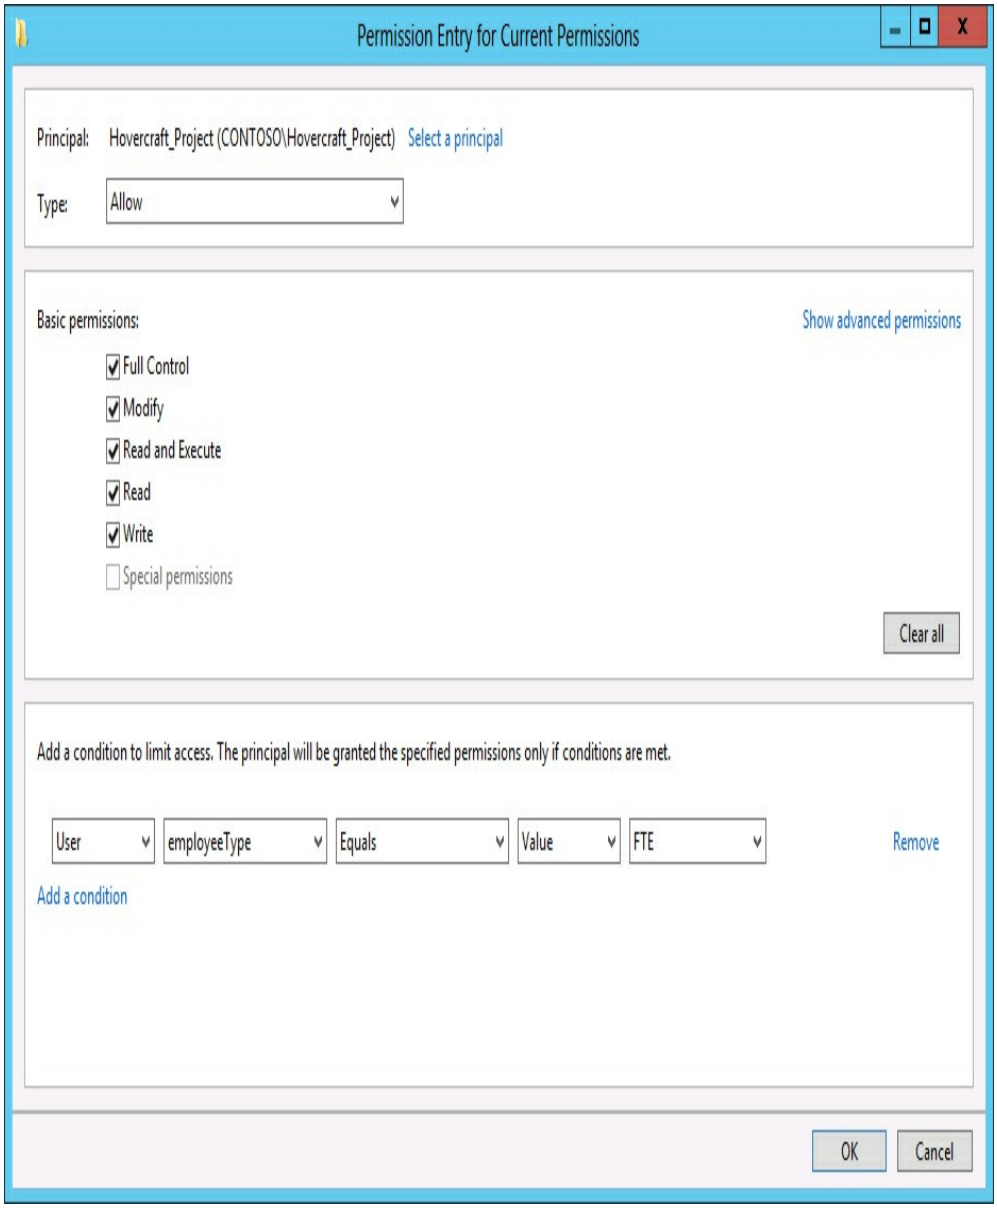


Figure 14-6 Permission Entry for Current Permissions

CENTRAL ACCESS POLICIES

A central access policy is a collection of central access rules. For example, the Contoso Policy Central Access Policy, shown in Figure 14-7, publishes two central

access rules: Research_Projects and Secret_Projects. Only file servers running Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, or Windows Server 2019 that are within the scope of the central access policy apply the rules that are contained within the policy.

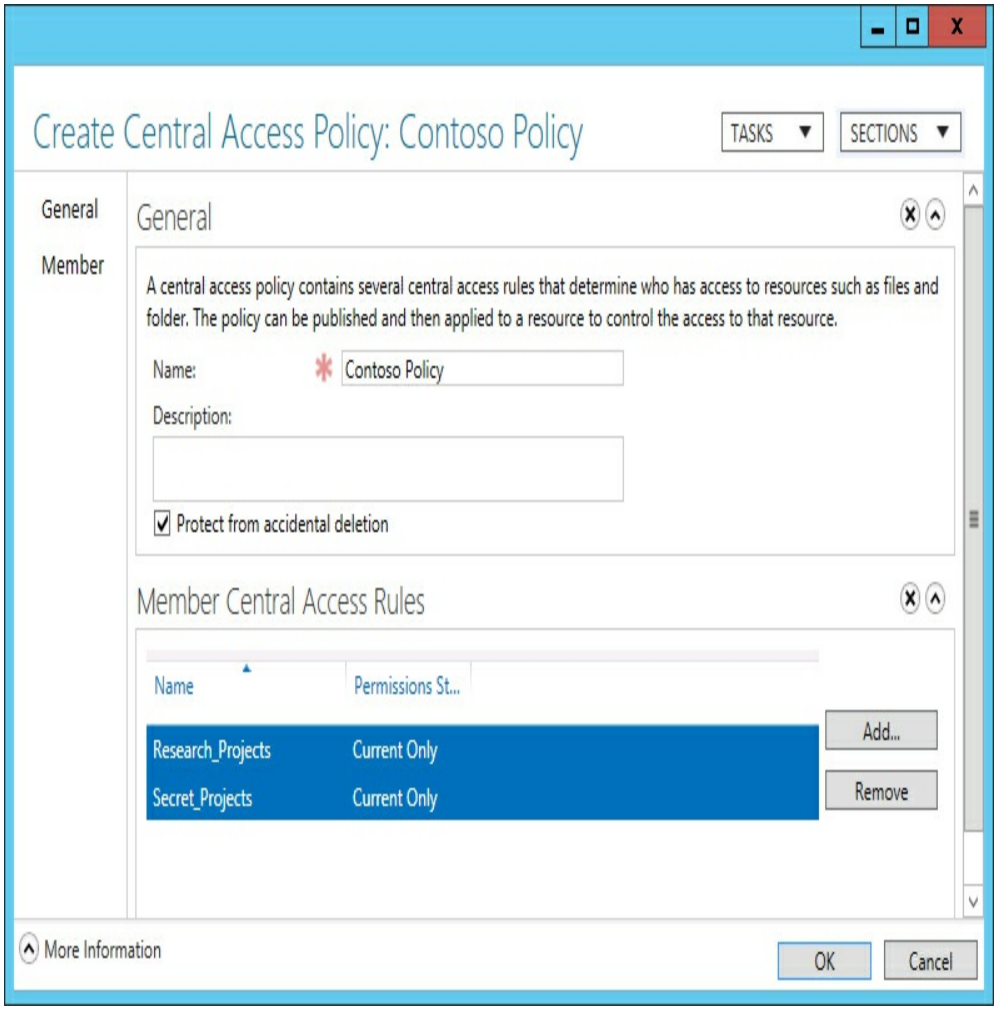


Figure 14-7 Create Central Access Policy: Contoso Policy

You distribute central access policies through Group Policy. You do this by configuring the Manage Central Access Policies policy, which is located in the Computer Configuration\Policies\Windows Settings\Security Settings\File System node. This policy is shown in Figure 14-8.

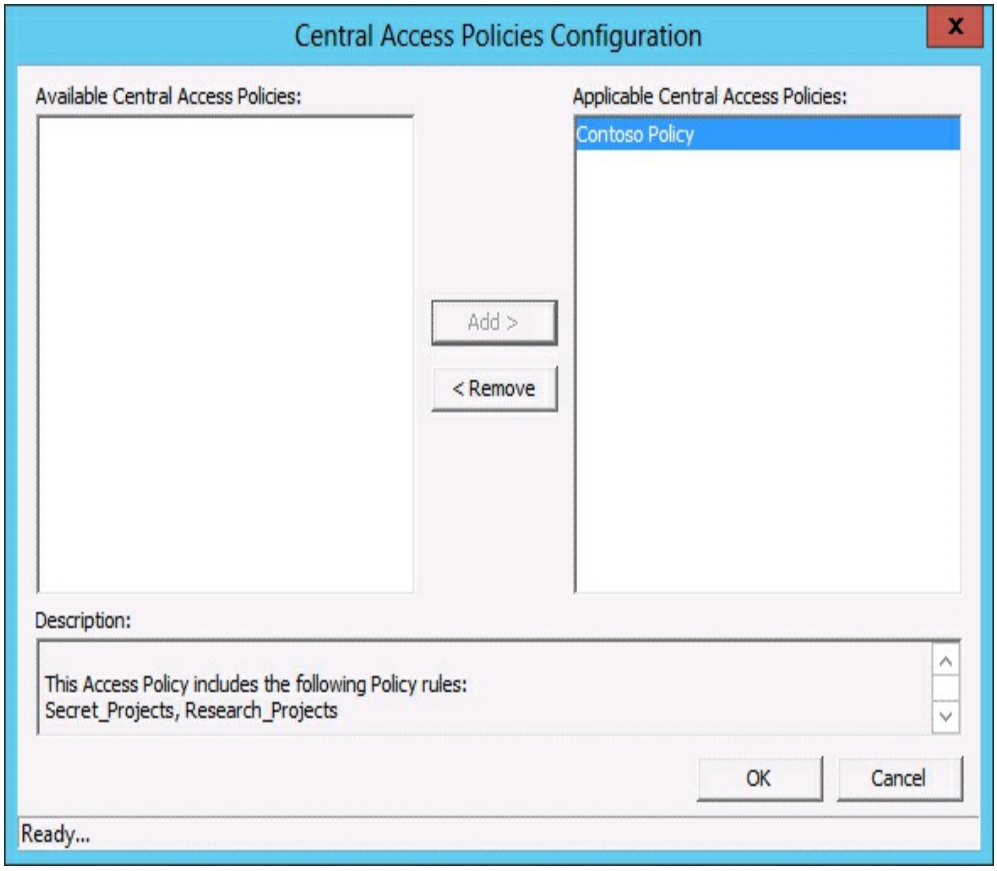


Figure 14-8 Central Access Policies Configuration

Inside OUT

Lots of work to get it running

You must do quite a bit of planning and configuration to get DAC up and running. After you do get it up and running, you'll find that the amount of time you spend mucking about with permissions is substantially reduced. You'll also find that file

system permissions are applied consistently and correctly, which is something that is very hard to achieve when they are applied manually. On the other hand, Microsoft is now spending far more effort on Azure Information Protection, which mostly does the same thing, so if you're starting from scratch in this area, you should probably read up on that topic in Chapter 17, "Azure IaaS and hybrid services."

STAGING

Staging allows you to configure a set of proposed rather than applied permissions. You use auditing to determine the results of these staged permissions before implementing them. You must enable auditing through Group Policy to determine the results of staged permissions. The policy you need to enable is the Audit Central Access Policy Staging Properties policy, which is located in the Computer Configuration\Policies\Windows Settings\Security Settings\Advanced Audit Policy Configuration\Audit Policies\Object Access node of a Group Policy Object (GPO).

You configure staged permissions in the Proposed area of a Central Access Rule by selecting the Enable Permission Staging Configuration option, as shown in Figure 14-9. You can verify the functionality of the proposed permissions by checking the Security event log. To do this, sign in to the file server that hosts the files and search for events with event ID 4818.

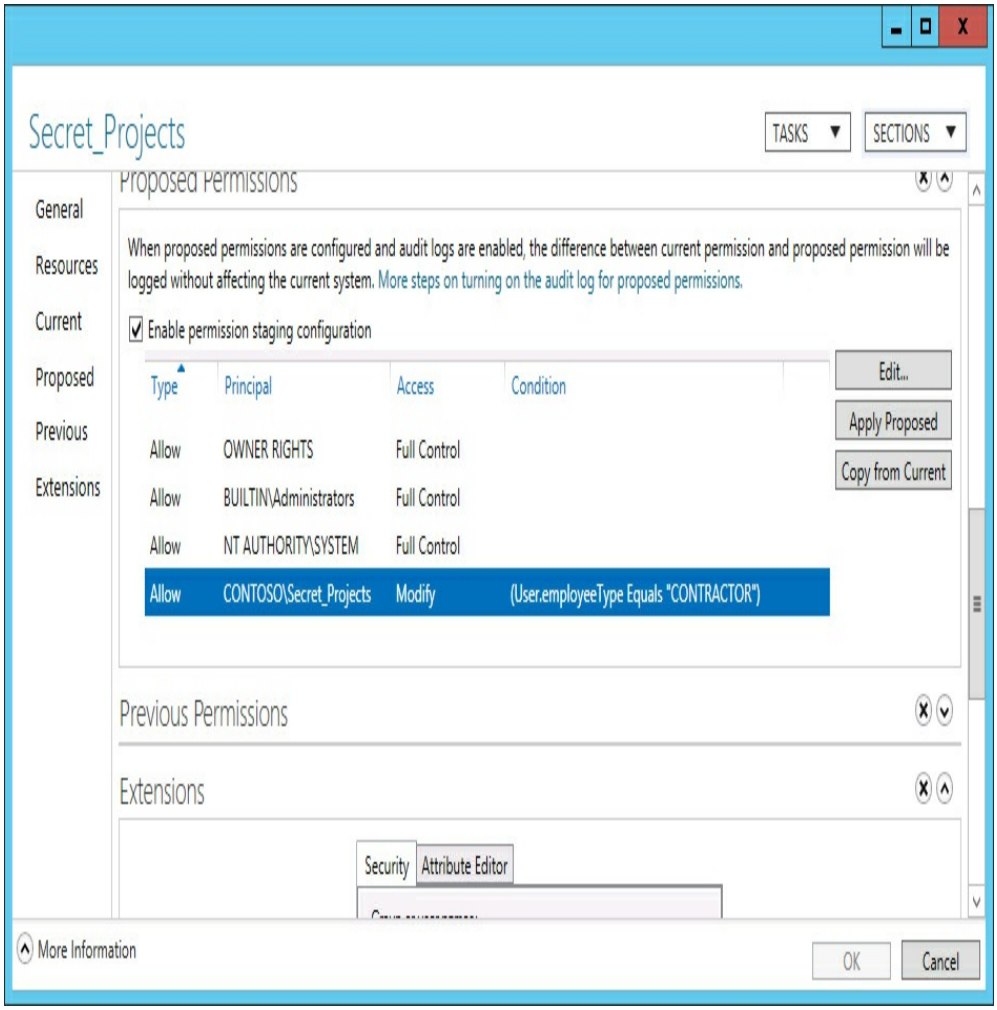


Figure 14-9 Proposed permissions.

ACCESS DENIED ASSISTANCE

Access Denied Assistance is a feature available to Windows 10, Windows 8.1, and Windows 8 clients that provides an informational dialog box explaining to users that they are unable to access a file because they do not have appropriate permissions. It's also possible to configure Access Denied Assistance so that an email can

be forwarded to the support desk if the help desk's assistance is needed to untangle permissions and allow the user access to a file.

You configure the Access Denied Assistance by configuring the Customize Message For Access Denied Errors policy, as shown in Figure 14-10. This policy is located in the Computer Configuration\Policies\Administrative Templates\System\Access-Denied Assistance node of a GPO.

Technet24

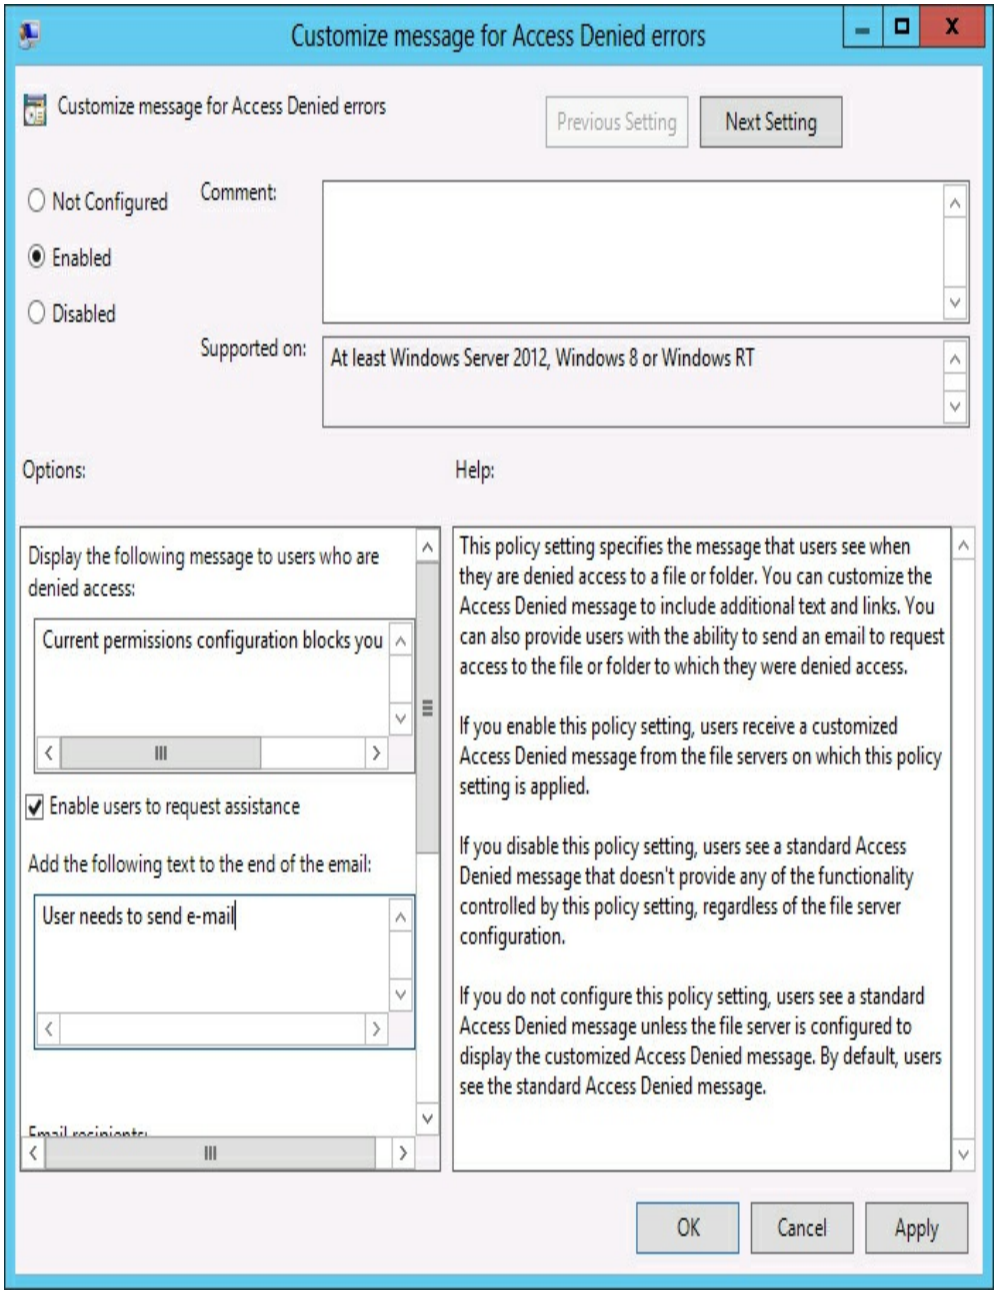


Figure 14-10 Customized Message For Access Denied Errors

INSTALLING AD RMS

AD RMS uses the term “cluster” to describe an AD RMS deployment, even though it has nothing to do with failover clustering or network load balancing. When you deploy AD RMS, you first deploy a root cluster. An AD RMS root cluster is responsible for managing all the AD RMS licensing and certificate traffic for the forest in which it is installed. You should only have one AD RMS root cluster per forest. Organizations that have multiple forests should deploy multiple AD RMS root clusters. After you have deployed a root cluster, you can configure additional licensing-only clusters. Licensing-only clusters distribute licenses that clients use to consume and publish content.

Installing AD RMS involves performing the following steps:

- **Specifying the database that AD RMS uses to store configuration information.** You can use a Microsoft SQL Server instance to perform this role or deploy the Windows Internal Database. You should only use Microsoft SQL Server in large AD RMS deployments. AD RMS can use any supported version of SQL Server 2008.
- **Specifying a service account.** This account needs to be a domain account, and best practice is to use a specially configured group managed service account for this role. Using a group managed service account ensures that the account password is managed by Active Directory and does not need to be manually updated on a periodic basis by an administrator.
- **Choose a cryptographic mode.** Mode 2 uses RSA 2048-bit keys and SHA-256 hashes. Mode 1 uses RSA 1024-bit keys and SHA-1

hashes. Mode 2 is more secure than mode 1 and is therefore the recommended choice.

- **Specify Cluster Key Storage.** This determines where the cluster key is stored. The default is to have the key stored within AD RMS, although you can also use a cryptographic service provider (CSP) if one is available. When you use a CSP, you must perform manual key distribution when adding additional AD RMS servers.
- **Specify a cluster key password.** This password is used to encrypt the cluster key. You need to provide this password when joining additional AD RMS servers to the cluster and when recovering an AD RMS cluster from backup.
- **Input the cluster address.** This is a website address in FQDN format that is usually hosted on the AD RMS server. Best practice is to configure an SSL certificate with the FQDN of the AD RMS server. Although it is possible to specify a non-SSL protected address, doing so removes the ability to use AD RMS with identity federation. The cluster address and port cannot be altered after you deploy AD RMS.
- **Specify a Licensor certificate name.** This is the name used with the licensor certificate and should represent the certificate's functionality.
- **Determine whether to register the service connection point (SCP) in Active Directory.** The SCP allows domain members to locate the AD RMS cluster automatically. A user account must be a member of the Enterprise Admins security group to register an SCP. SCP registration can occur after AD RMS is deployed if the account used to deploy AD RMS is not a member of this security group.

AD RMS CERTIFICATES AND LICENSES

AD RMS uses four specific types of certificates. These certificates have the following functions:

- **Server licensor certificate (SLC).** This certificate is created when you install the AD RMS role on the first server in the AD RMS cluster. This certificate is valid for 7150 days and is used to issue the following certificates and licenses:
    - SLCs to additional servers that join the cluster
    - Rights account certificates
    - Client licensor certificates
    - Publishing licenses
    - Use licenses
    - Rights policy templates
- **AD RMS machine certificate.** This certificate identifies a trusted device. The machine certificate public key encrypts Rights Account Certificate private keys, and the machine certificate private key decrypts Rights Account Certificates.
- **Rights account certificate (RAC).** This certificate identifies a user. AD RMS can only issue RACs to AD DS users whose user accounts are configured with an email address. By default, a RAC has a validity of 365 days. Temporary RACs are issued when a user accesses content from a device that is not a member of a trusted forest and has a validity of 15 minutes.
- **Client licensor certificate.** This certificate allows AD RMS–protected content to be published to computers that cannot connect directly to the AD RMS cluster. These certificates are tied to a user RAC. Other computer users are unable to publish AD RMS–protected documents until a new connection to the AD RMS cluster is established from the computer.

In addition to the four certificate types, AD RMS uses two license types:

- **Publishing license.** A publishing license determines the rights that apply to AD RMS content. This license contains the content

key, URL, and digital signature of the AD RMS server.

- **End-user license.** The end-user license allows a user to access AD RMS–protected content. An end-user license is issued per user, per document. End-user licenses are cached by default, although it's possible to disable caching, so that an end-user license must be obtained each time the user attempts to access protected content.

AD RMS TEMPLATES

Rights Policy Templates allow you to apply rights policies to documents. When an author creates a document or sends an email message, the author can apply a template to that document. It's also possible to use File Server Resource Manager to automatically apply templates to documents based on the properties of those documents, such as the document having a particular resource property or containing a specific text string.

Rights Policy Templates are created by using the AD RMS Management Console. When creating a template, you can enable the following rights on a per user group or per user basis, with any right not granted unavailable to the user:

- **Full Control.** The user has full control over the AD RMS–protected content.
- **View.** Gives the user the ability to view the AD RMS–protected content.
- **Edit.** Allows the user to modify the AD RMS–protected content.
- **Save.** Allows the user to save the AD RMS–protected content.

- **Export.** Allows the user to use the Save As function with the AD RMS–protected content.
- **Print.** Allows the user to print the AD RMS–protected content.
- **Forward.** Used with Microsoft Exchange, allows the user to forward a protected message.
- **Reply.** Used with Microsoft Exchange, allows the user to reply to a protected message.
- **Reply All.** Used with Microsoft Exchange, allows the recipient of a protected message to use the Reply All function.
- **Extract.** Allows a user to copy data from the AD RMS–protected content.
- **Allow Macros.** Allows the user to use macros with the AD RMS–protected content.
- **View Rights.** Allows the user to view rights assigned to the AD RMS–protected content.
- **Edit Rights.** Allows the user to modify rights assigned to the AD RMS–protected content.

Figure 14-11 shows the rights assigned to the submarine_project@contoso.com group. You can assign different rights to multiple groups. If a user is a member of more than one group, the rights are cumulative.

Technet24

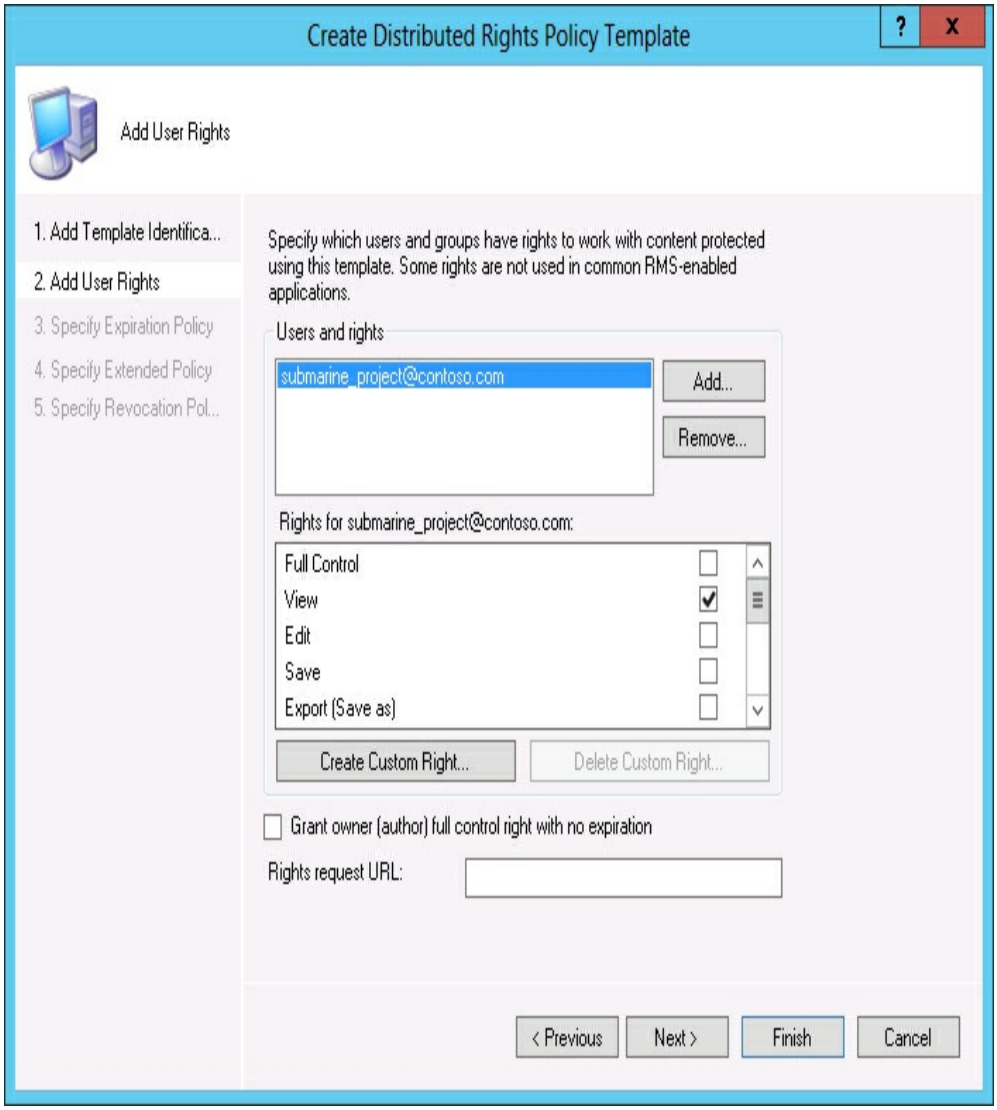


Figure 14-11 Create Distributed Rights Policy Template

When configuring an AD RMS Template, you can configure content expiration settings. Content expiration settings allow you to have content expire either on a certain date or after a certain number of days. The example in Figure 14-12 shows content

expiration configured to expire 14 days after content publication. An additional setting allows you to configure use license expiration, which allows you to configure how often a user must connect to the AD RMS cluster to obtain a new license to access the content.

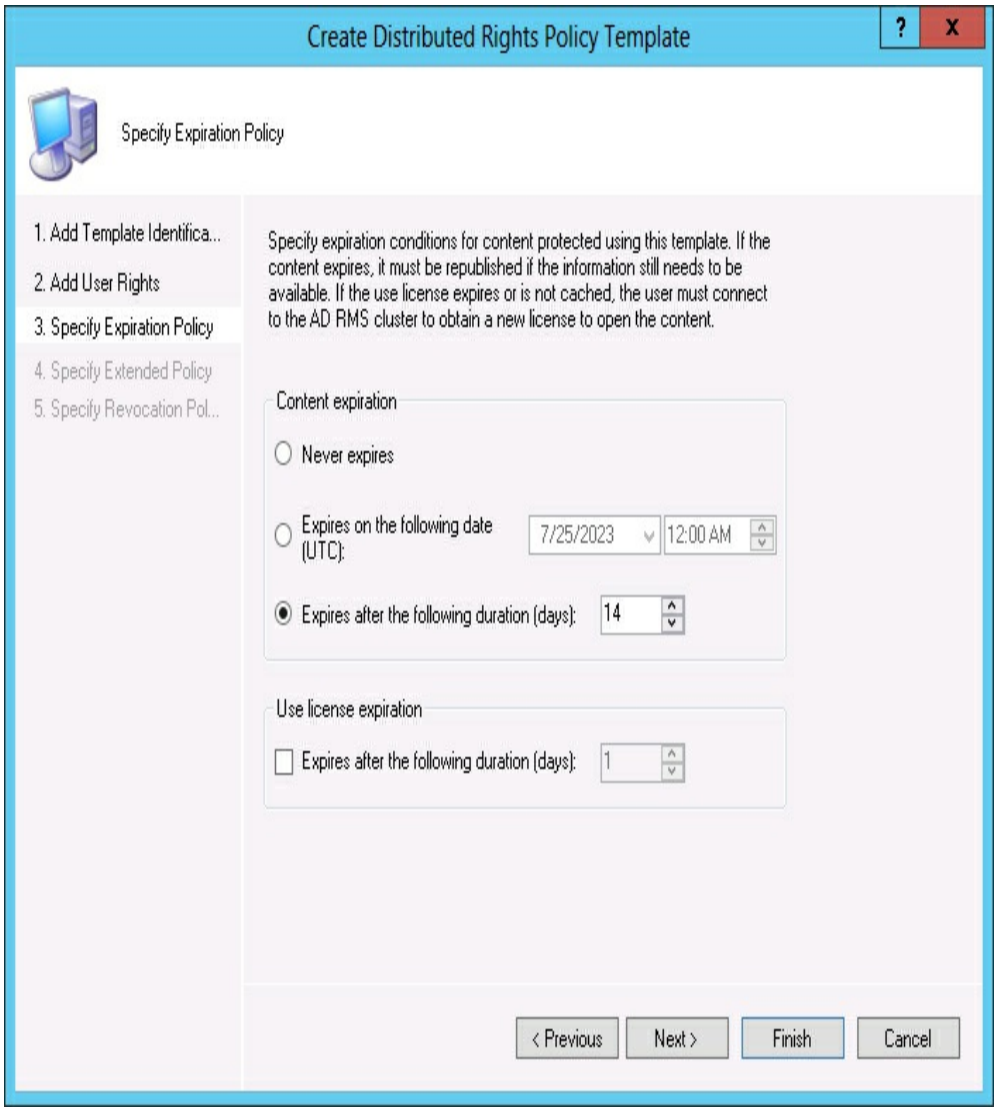


Figure 14-12 Create Distributed Rights Policy Template

The Extended Policy settings, as shown in Figure 14-13, allow you to configure whether AD RMS content can be viewed using a browser add-on and whether a new license must be obtained each time content is consumed.

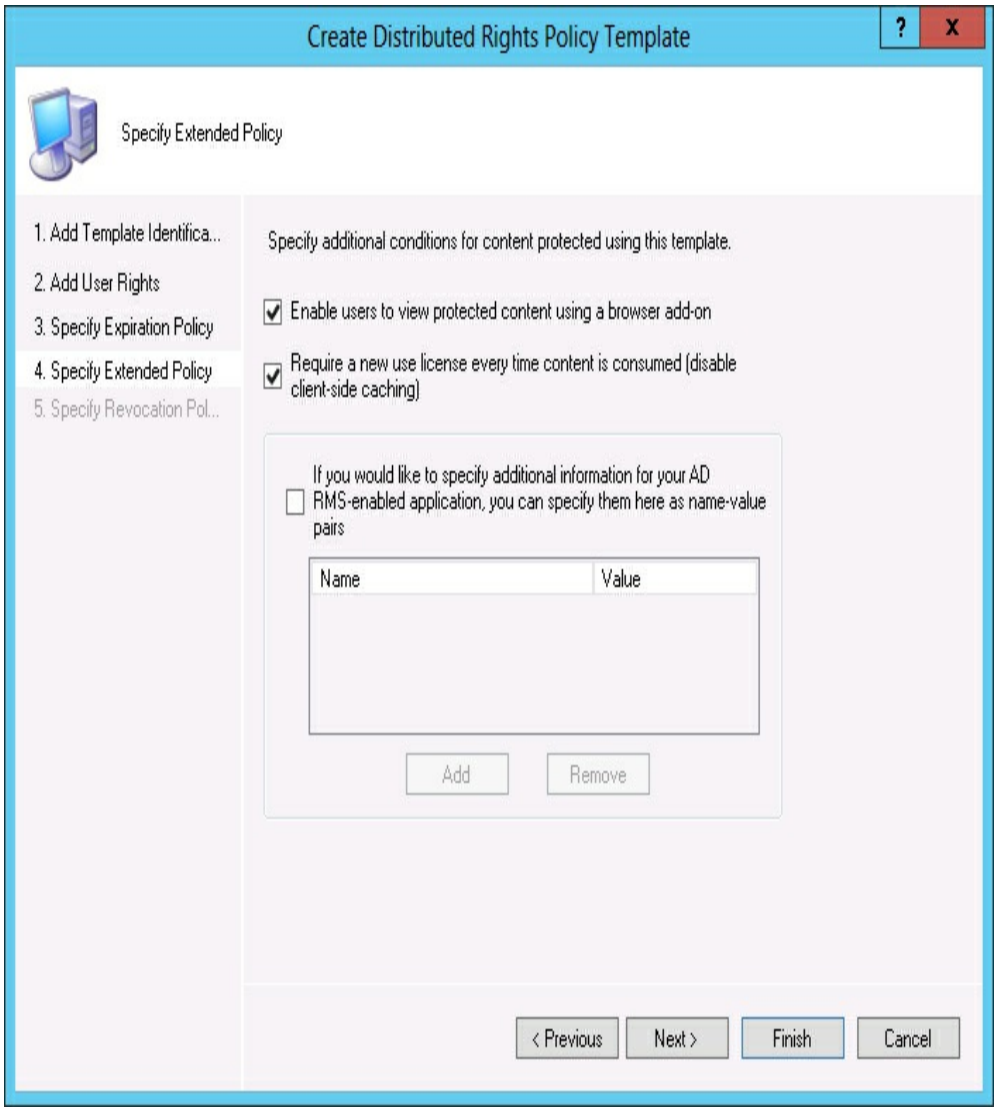


Figure 14-13 Configure browser and content settings.

AD RMS ADMINISTRATORS AND SUPER USERS

There are three separate local groups on an AD RMS server that you can add users to when you want to assign them privileges within AD RMS. These groups are as follows:

- **AD RMS Enterprise Administrators.** Members of this group can perform any task within AD RMS, including enabling the AD RMS Super Users group.
- **AD RMS Template Administrators.** Users who are members of this group can configure and manage AD RMS templates.
- **AD RMS Auditors.** Users who are members of this group are not able to modify AD RMS server settings and templates but can view the properties of the server and templates.

The AD RMS Super Users group is a special group that you can configure and enable on the AD RMS server. Members of the AD RMS Super Users group have full owner rights over all use licenses issued by the AD RMS cluster. Members of the Super Users group can:

- Recover expired content
- Recover content when a template is deleted
- Recover content without requiring author credentials

Because members of this group have access to all content, you should have strict policies about managing and auditing this group's membership. A Super Users group is not configured by default, and the group that you configure as the Super Users group must have an

associated email address, as shown in Figure 14-14.

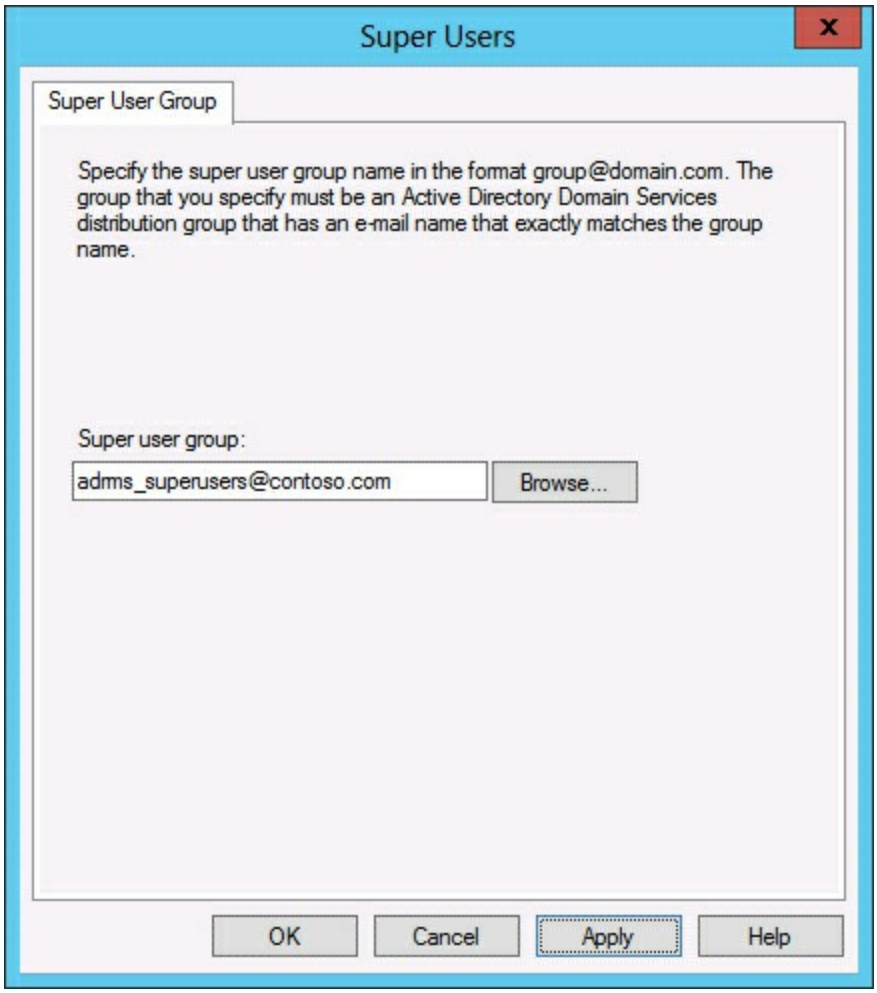


Figure 14-14 Super User Group

TRUSTED USER AND PUBLISHING DOMAINS

Trusted User Domains (TUDs) allow you to configure an AD RMS cluster to manage requests for CLCs for users who have been issued RACs from a different AD RMS cluster. For example, if an organization has two separate

Active Directory forests, and each forest has its own AD RMS deployment, you'd configure Trusted User Domains so that clients from one forest can issue CLCs to clients with RACs issued from the other forest. TUDs can be one-way or bidirectional. When configuring TUDs, you must export the TUD from the partner before importing the TUD locally.

Trusted Publishing Domains (TPDs) allow the AD RMS cluster in one forest to issue end-user licenses to content published with licenses issued from an AD RMS cluster in another forest. You must export the TPD file and have it imported by the partner AD RMS cluster before the AD RMS cluster in the partner forest can issue end-user licenses to local AD RMS clients.

EXCLUSION POLICIES

Exclusion policies allow you to deny specific entities the ability to interact with AD RMS. You can configure exclusions on the basis of application, user, and lockbox version. Exclusion works in the following ways:

- **User Exclusion.** You can exclude a user based on email address or based on the public key assigned to the user's RAC. Use email-based exclusions for users in the forest and public key-based exclusions for external users.
- **Lockbox Exclusion.** Lockbox exclusion allows you to exclude specific client operating systems. Each version of the Windows operating system has a specific lockbox identity. If you want to block clients running Windows 7 from interacting with AD RMS,

configure an exclusion where the minimum lockbox version to the version available is Windows 8.

- **Application Exclusion.** This allows you to exclude specific applications from interacting with AD RMS. You must specify the application file name, the minimum version, and the maximum version.

When you configure an exclusion, the exclusion only applies to new certificate or licensing requests. Licenses and certificates that were issued during the exclusion period still exclude the application, user, or lockbox version. If you remove an exclusion, the removal only applies to new licenses or certificates.

Apply AD RMS templates automatically

You can use File Server Resource Manager to automatically apply AD RMS templates to files. You do this by performing the following general steps:

1. Create a new file management task with an appropriate name related to the template.
2. On the Scope tab, set the scope of the task to the folders that host the files to which you want to apply the AD RMS template.
3. On the Condition tab, specify the condition that allows the rule to recognize the files that you want to apply the AD RMS template to. For example, you can create a rule that is triggered if the file contains the text "SECRET."
4. On the Action tab of the Create File Management Task dialog box, shown in Figure 14-15, specify RMS Encryption as the action type and the AD RMS template that you want to apply to the file.

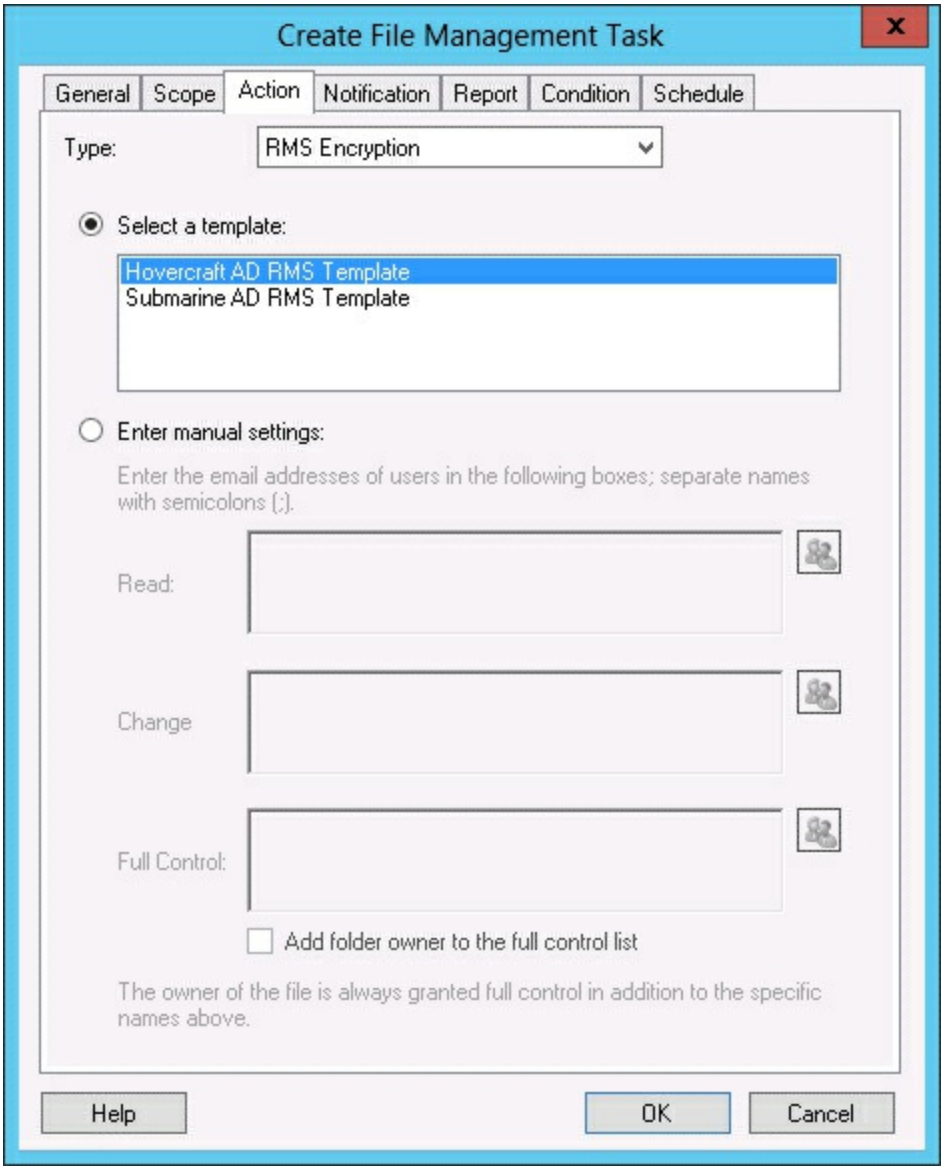


Figure 14-15 Create File Management Task

5. On the Schedule tab, configure how often the file management task should run and whether it should classify only new files or periodically attempt to reclassify existing files.

MANAGING AD RMS WITH WINDOWS POWERSHELL

AD RMS PowerShell cmdlets are available in the ADRMS and ADRMSADMIN Windows PowerShell module. The functionality of these cmdlets is described in Table 14-1.

Table 14-1 AD RMS PowerShell cmdlets

Noun	Verbs	Function
ADRMS	Install, Uninstall, Update	Install, remove, and update AD RMS after the binaries have been installed on the target server
RmsCertChain	Get	Generates a report detailing information on the certificate chain of a specific user request on the AD RMS server
RmsCertInfo	Get	Generates a report about a specific certificate used in a user request for the AD RMS cluster
RmsChildCert	Get	Returns all child certificates from a parent certificate used in a user request to an AD RMS cluster
RmsCluster	Update	Updates RMS cluster information
RmsCryptoMode2	Initialize	Configure an RMS cluster to transition to cryptographic mode 2
RmsEncryptedIL	Get	Determine use-license information from an issuance license associated with a user request

RmsMfgEnrollment	Install, Update, Uninstall	Configure Microsoft Federation Gateway service enrollment
RmsMfgSupport	Install, Uninstall	Install or remove Microsoft Federation Gateway configuration
RmsReportDefinitionLanguage	Export	Export AD RMS report definition files
RmsRequestInfo	Get	Generate a report on a specific user request
RmsSvcAccount	Set, Get	Configure RMS service account
RmsSystemHealthReport	Get	View RMS system health
RmsTPD	Export, Import	Import and export trusted publishing domain information
RmsTUD	Import, Export	Import and export trusted user domain
RmsUserRequestReport	Get	View information about an RMS user request

Technet24

Dynamic Access Control cmdlets

You can use the cmdlets listed in Table 8-6 to manage Dynamic Access Control.

Table 14-2 Dynamic Access Control cmdlets

Noun	Verb	Function
ADCentralAccessPolicy	Get, New, Remove, Set	Manage central access policies
ADCentralAccessPolicyMember	Add. Remove	Add and remove central access rules to a central access policy in AD DS
ADCentralAccessRule	Get, New, Remove, Set	Manages central access rules
ADClaimType	Get, New, Remove, Set	Manage claim types
ADResourceProperty	Get, New, Remove, Set	Manage resource properties
ADResourcePropertyList	Get, New, Remove, Set	Manage resource property lists
ADResourcePro	Add,	Manage resource properties in a

pertyListMembe r	Remove	resource property list
ADResourcePro pertyValueType	Get	View resource property value type

Chapter 15

Routing and Remote Access

Remote Desktop Gateway

Virtual private networks

LAN routing

Network Address Translation (NAT)

DirectAccess

Managing Remote Access using PowerShell

Today, we expect to be able to access work resources whether we're in the office, in an airline lounge, in a coffee shop, or at home on the couch. Remote access technologies allow authorized users to make a connection from their laptops and tablet computers to resources on their organization's internal network from any location on the public Internet.

In this chapter, you learn about Windows Server 2019 remote access technologies and how you can leverage those technologies to enable remote access to your organization's internal network. You learn how each technology is appropriate for a specific use and why you might want to use several technologies to support all the clients in your organization.

REMOTE DESKTOP GATEWAY

Remote Desktop (RD) Gateway allows users to make a connection from a host on the Internet to a host on the local area network using the Remote Desktop Connection (RDC) Client software. Remote Desktop Client software is included in all currently supported versions of Windows and is also available for macOS X computers, as well as iOS and Android devices. Although RD Gateway is typically used to connect to Remote Desktop Session Host servers or VDI virtual machines, you can also configure it to connect to desktop workstations. This allows a user to use an RD Gateway server to connect to her work PC from her laptop computer without having to connect through a VPN.

RD Gateway servers, like any other remote access servers, need to be placed on perimeter networks and need to have at least one public IP address. Connections to RD Gateway servers occur by using the Remote Desktop Protocol (RDP) over HTTPS protocol. When

configuring the external firewall, network administrators need to allow access on port 443 from hosts on the Internet to the RD Gateway server. Connections from the RD Gateway server to RDP hosts on the internal network occur on port 3389.

RD Gateway servers don't need to be members of an Active Directory domain, but it is substantially more difficult to configure the RD Resource Authorization Policy (RAP) to limit connections to specific hosts without such membership. You can accomplish something similar through firewall rules by restricting access on the basis of network address. The drawback to domain membership, however, is the risk involved in placing a domain-joined computer on a perimeter network and configuring the internal firewall to support the configuration.

RD Gateway connection and resource policies

RD Gateway connection and resource policies allow you to specify under what conditions users can connect to an RD Gateway server. RD Connection Authorization Policies (CAPs) can be stored on the RD Gateway server or in a Central RD CAP Store. Some of the benefits of using a Central CAP Store are that the same CAPs can apply to multiple RD Gateway servers without having to reproduce them on each server and that they can also be

centrally updated as needed. If your organization has only a single RD Gateway server however, there is no need to use a central RD CAP Store.

On an RD Gateway server that is not a member of a domain, you might have to create a separate local user group with local user accounts to allow users to authenticate against the RD Gateway server prior to authenticating against the session target.

To create an RD Gateway CAP, perform the following general steps:

1. In the RD Gateway Manager console, right-click Connection Authorization Policies, click Create New Policy, and then click Custom.
2. On the New RD CAP dialog box, enter a policy name on the General tab.
3. On the Requirements tab, select Password or, if your organization supports it, Smart Card. Next, click Add Group to specify the User Group to which you want to grant access to RD servers through the RD Gateway server. You should create a separate security group for users you want to grant access to, which makes granting and revoking access just a matter of removing the user's account from this group.
4. On the Device Redirection tab, choose whether you want to enable device redirection for all client devices or disable it for specific types, such as volumes, clipboards, and plug-and-play devices. You also have the ability to restrict connections to only clients that enforce RD Gateway device redirection policies. You might do this if you want to block access to clients running non-Microsoft operating systems.
5. On the Timeouts tab, you can enable an idle timeout, which disconnects users who aren't interacting with their session, and a session timeout, which limits all sessions including active sessions. Click OK to finish creating the policy.

After you have configured an RD CAP, you need to configure an RD Resource Authorization Policy (RAP). An incoming connection must meet the conditions specified in one RD CAP and one RD RAP before it can be successfully established. To create an RD RAP, perform the following steps:

1. In the RD Gateway Manager console, right-click Resource Authorization Policies, click Create New Policy, and then click Custom.
2. On the New RD RAP page, enter the policy name.
3. On the User Groups page, click Add to add the user groups that the policy applies to.
4. On the Network Resource page, choose between the following options:
    1. Select an Active Directory Domain Services group
    2. Select an existing RD Gateway managed group
    3. Allow users to connect to any network resource
5. On the Allowed Ports page, choose to restrict connections from the RD Gateway server to port 3389, which is the default, to another specific port, or through any port. Unless the RD servers use alternate ports, the default is appropriate.

Configuring server settings

RD Gateway servers need a TLS certificate to encrypt and authenticate the RDP over an HTTPS connection. As is always the case with TLS certificates, obtaining a certificate from a trusted, third-party certification authority (CA) may cost money and time (at the time of writing, the most popular CA provides TLS certificates free of charge), but this investment means that any

client can use the RD Gateway without problems. Using an internal certificate is cheaper because it is not necessary to purchase a certificate, but it does require you to configure RD Gateway clients to trust the issuing authority. Given the organizational cost to support the distribution of an internal or self-signed certificate, it is usually just easier to obtain a certificate from a trusted, third-party CA.

Configuring clients to use RD Gateway

You can configure clients to use an RD Gateway server either by manually configuring the server's address by editing the advanced properties of Remote Desktop Connection, as shown in Figure 15-1, or by configuring the server's address through Group Policy.

Technet24

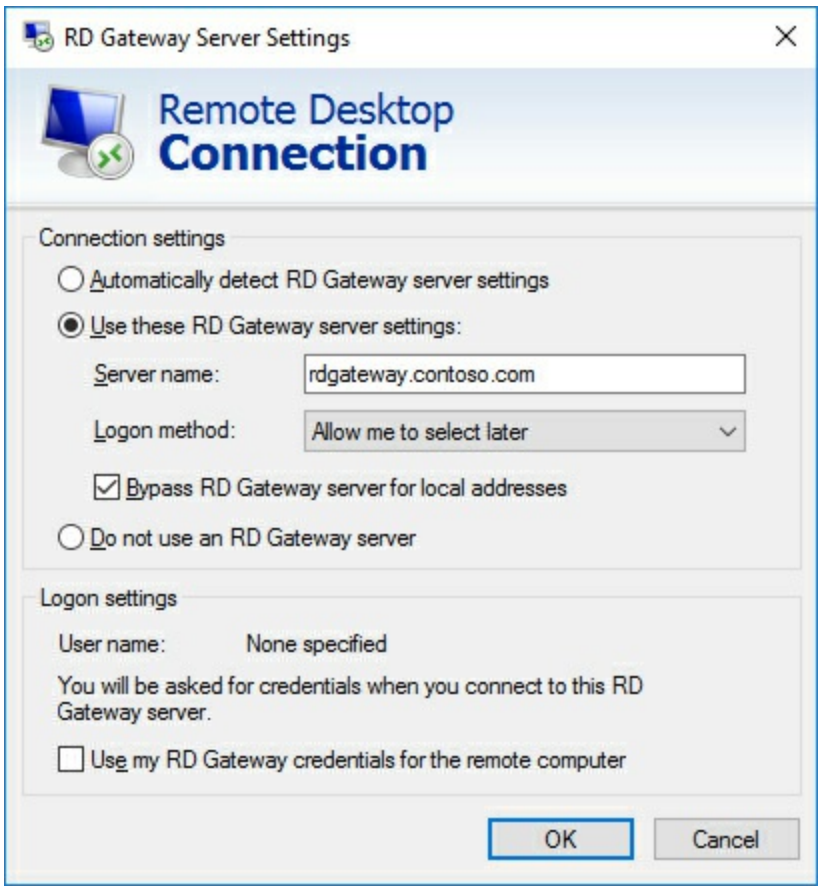


Figure 15-1 RD Gateway Settings

There are three RD Gateway-related Group Policy items that can be configured to simplify the process of ensuring that clients connect to the correct server. These policy items are located in the User Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\RD Gateway node of a standard GPO. These policies can be used to configure the following:

- **Set RD Gateway Authentication Method.** This policy allows you to force a user to use a particular authentication method, but it

also allows you to set a default authentication method that the user can override. The options are:

- Ask For Credentials; Use NTLM Protocol
- Ask For Credentials; Use Basic Protocol
- Use Locally Logged-In Credentials
- Use Smart Card
- Negotiate Protocol

- **Enable Connection Through RD Gateway.** When this policy is enabled, the client attempts to connect through the configured RD Gateway server if it is unable to establish a connection to the specified target of the RD Connection. Set this policy and the RD Gateway Address Policy to ensure that domain clients connect through RD Gateway when they are not connected to the internal network.
- **Set RD Gateway Server Address.** This policy allows you to configure the RD Gateway server address, either on the basis of IP address or fully qualified domain name (FQDN). This policy should be set in conjunction with the Enable Connection Through RD Gateway policy to ensure that remote domain clients transparently connect through the RD Gateway server without having to configure anything themselves.

VIRTUAL PRIVATE NETWORKS

Windows Server 2019 supports four separate virtual private network (VPN) technologies: PPTP, L2TP/IPsec, SSTP, and IKEv2. PPTP is the simplest of these technologies to set up because it does not require integration with certificate services. The drawback of PPTP is that it is an older protocol, has known

vulnerabilities, and is not as secure as the other VPN protocols available on Windows Server 2019. Some organizations use SSTP or IKEv2, but most use L2TP/IPsec.

Inside OUT

VPN Certificates

If you are using VPN protocols that rely on certificates generated by an internal enterprise CA, ensure that a Certificate Revocation List (CRL) Distribution Point (CDP) or Online Responder is accessible to clients on the Internet. Remember that clients need to be able to check the revocation status of certificates, which is difficult if the CDP is accessible only to clients on the internal network. You can deploy a CRL DP on the perimeter network or host it in the cloud, which is more common today.

IKEv2 Always On VPN protocol

The IPsec Tunnel Mode with Internet Key Exchange version 2 (IKEv2) VPN protocol was introduced with the release of the Windows Server 2008 R2 and Windows 7 operating systems. From Windows Server 2012 R2 onward, this protocol also went by the name Always On VPN. The main benefit of Always On VPN is its capability to allow automatic reconnection without requiring reauthentication when the connection is disrupted. Automatic reconnection can occur even if the client's connection address changes and reconnection is possible for up to eight hours. For example, imagine that a user is connected to the internal network using an Always On connection through a laptop computer's cellular modem and then boards an airplane and

connects to the airplane's Wi-Fi. Under an Always On scenario, the connection is automatically reestablished without requiring the user to manually reauthenticate, even though the connection method, from cellular network to airplane Wi-Fi, has changed. Connections only have to be reestablished when a user puts the computer into hibernation mode.

Always On/IKEv2 has the following features:

- Supports IPv6
- Enables VPN reconnect
- Supports EAP and Computer certificates for client authentication
- Does not support PAP or CHAP
- Only supports MS-CHAPv2 with EAP
- Supports data origin authentication, data integrity, replay protection, and data confidentiality
- Uses UDP port 500

To configure Always On VPN server running Windows Server 2019, do the following:

1. Deploy an enterprise CA.
2. Create a new certificate template based on the IPsec template. You need to modify the application policies of this template to ensure that it supports the IP security IKE intermediate policy and Server Authentication certificate.
3. Install a certificate generated from this template on the VPN server.
4. Ensure that at least one Certificate Revocation List Distribution Point is accessible to clients on the Internet. You can place this CRL DP on the perimeter network or host it in the cloud.

5. Ensure that all clients trust the enterprise CA that issued the VPN server's certificate.

Windows clients running the Windows 8 operating system or later attempt to establish an Always On VPN session before they try SSTP or L2TP/IPsec. If those options are unsuccessful, the client falls back to PPTP.

More Info

Always On VPN

You can find out more about Always On VPN at https://docs.microsoft.com/en-us/windows-server/remote/remote-access/vpn/always-on-vpn/always-on-vpn-technology-overview.

SSTP VPN protocol

Some locations block the ports that some VPN protocols use. The advantage of the SSTP VPN protocol is that it allows computers that might not normally be able to establish a remote access connection to do so through TCP port 443, which is usually used for TLS connections to websites. While it doesn't have the performance of other protocols, SSTP almost always works, even if it has to traverse NAT, firewalls, and proxy servers. SSTP doesn't work, however, if there is a web proxy in place that requires manual authentication. For example, if you are in a hotel and you have to sign on to the Wi-Fi network through a browser, you need to perform this action before you attempt to establish an SSTP connection.

SSTP has the following requirements:

- You need to install a TLS certificate on the VPN server.
- All clients need to trust the CA that issued the TLS certificate.
- You need to configure the TLS certificate with a name that matches the FQDN that maps to an IP address associated with the external network interface of the VPN server.
- The Certificate Revocation List Distribution Point of the CA that issued the TLS certificate needs to be accessible to clients on the Internet.

L2TP/IPsec protocols

Layer Two Tunneling Protocol with IPsec (L2TP/IPsec) is a VPN protocol that has been used with Windows Server operating systems and clients since the release of Windows 2000. It works with almost every Windows and third-party operating system.

When configuring L2TP/IPsec you need to ensure that you have configured a CA that is able to automatically issue certificates to connecting clients. As is the case with other VPN protocols, the CA needs to be trusted by the client, and the CDP must be accessible to clients making connections from the Internet. You would primarily deploy L2TP/IPsec as a way of providing secure access to clients that are not able to leverage the simpler to implement Secure Socket Tunneling Protocol (SSTP) VPN protocol. Although L2TP/IPsec usually requires you to deploy digital certificates, it is possible,

with special configuration, to get L2TP/IPsec to work with pre-shared keys.

PPTP VPN protocol

PPTP is the oldest VPN protocol supported by Windows Server 2019 and is also the least secure. It is most often used when organizations need to support clients running unsupported operating systems, such as Windows XP, or in situations where it isn't possible to deploy the certificate infrastructure required to implement L2TP/IPsec. PPTP connections provide data confidentiality but do not provide data integrity or data origin protection. This means that captured data can't be read, but you also can't be sure that the transmitted data was the same data sent by the client.

> **Inside OUT**
>
> **PPTP**
>
> PPTP requires no additional configuration beyond setting up the VPN server and ensuring that port 1723 is open for clients from the Internet. Just remember that PPTP has been around for a long time and isn't a secure VPN protocol.

VPN authentication

Although Windows Server 2019 supports many protocols that have been in use for some time, these protocols are often less secure than more recently developed protocols. Windows Server 2019 supports the following protocols, listed from most secure to least

secure:

- **Extensible Authentication Protocol-Transport Level Security (EAP-TLS).** Use this protocol with Smart Cards or digital certificates. You can only use this protocol if you are using RADIUS authentication or if the remote access server performing authentication is domain-joined.
- **Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2).** Provides mutual authentication, which means that not only is the user authenticated, but the service that the user is connecting to is also authenticated. Allows for the encryption of the authentication process and the session.
- **Challenge Handshake Authentication Protocol (CHAP).** Authentication data is encrypted through MD5 hashing. The data is not encrypted.
- **Password Authentication Protocol (PAP).** This protocol does not encrypt authentication data, meaning that if the authentication is captured, then attackers will have access to credentials.

Deploying a VPN server

To configure a Windows Server 2019 server that has one network card connected to the perimeter network, which is accessible to clients on the Internet, and a second network card that is connected to the internal network to support VPNs, perform the following general steps:

1. Ensure that the prospective VPN server is connected to the Active Directory domain and log in with an account that is a Domain Admins group member.
2. Open an elevated PowerShell prompt and enter the following commands:

Click here to view code image

```
Install-WindowsFeature -IncludeAllSubfeature
RemoteAccess
```

3. Open the Routing And Remote Access console, right-click the server name, and click Configure And Enable Routing and Remote Access. This opens the Routing And Remote Access Server Setup Wizard. Click Next.
4. On the Configuration page, select Remote Access (Dial-Up Or VPN) and then click Next.
5. On the Remote Access page, select VPN and then click Next.
6. On the VPN Connection page, select the network connection that is connected to the Internet and then click Next.
7. On the IP Address Assignment page, choose between assigning IP addresses automatically from an existing scope or from a specified range of addresses.
8. On the Manage Multiple Remote Access Servers page, select No, Use Routing And Remote Access To Authenticate Connection Requests, click Next, and then click Finish.

Disable VPN protocols

By default, Windows Server 2019 server supports SSTP, PPTP, L2TP, and IKEv2 remote access connections. If you want to limit incoming connections to a specific protocol or disable an existing protocol, you need to edit the port properties related to that protocol. If you don't disable a specific protocol, clients might use that protocol to connect to the VPN server. For instance, if you've gone to the effort of configuring IKEv2, you don't want to find out a few weeks later that all your clients

are still falling back to using PPTP! To configure which VPN protocols are available, perform the following general steps:

1. Open the Routing and Remote Access console, right-click the Ports node under the VPN server, and then click Properties.
2. On the Ports Properties dialog box, select the VPN protocol for which you want to limit connections, and then click Configure.
3. On the Configure Device dialog box, set the Maximum Ports to 1 and ensure that the Remote Access Connections (Inbound Only) item is not selected. Click OK.

Granting access to a VPN server

If you have only a small number of users, you can grant access to those users on an individual basis by editing their user account properties in Active Directory Users And Computers. To grant access to an individual user, choose Allow access under Network Access Permission on the Dial-In tab of the user's Account Properties.

The default setting for network access for access to a Windows Server 2019 VPN server is for it to be controlled through a Network Policy Server (NPS) network policy. To be granted access, a user must meet the conditions of a connection request policy and a network policy. A connection request policy includes type of access, time of access, and which authentication protocols are available. A network policy can be as simple as specifying that a particular user group has access.

When you install a VPN server, a default connection request policy called Microsoft Routing And Remote Access Service Policy is created. You can view the policy's properties under the Policies\Connection Request Policies node of the Network Policy Server console on the VPN server. To configure this policy, perform the following steps:

1. Open the Network Policy Server console on the VPN server.
2. Expand the Policies\Connection Request Policies node. Right-click Microsoft Routing And Remote Access Service Policy and then click Properties.
3. On the Conditions tab, verify that the day and time restrictions match those that you want to enforce for the remote access policy. You can also use this tab to add conditions to the policy. Clients must meet all specified conditions, so keep in mind that the more conditions you add, the fewer clients the policy applies to. Conditions that you can add include the following:
    1. **User Name.** Specifies a specific HCAP user name.
    2. **Access Client IPv4 Address.** Specifies the RADIUS client IPv4 address.
    3. **Access Client IPv6 Address.** Specifies the RADIUS client IPv6 address.
    4. **Framed Protocol.** Specifies the protocol used for framing incoming packets, such as Point-to-Point Protocol (PPP) or Serial Line Internet Protocol (SLIP).
    5. **Service Type.** Limits clients to a specific service, such as Point-to-Point protocols.
    6. **Tunnel Type.** Restricts clients to a specific tunnel type, such as PPTP, L2TP, and SSTP.
    7. **Day and Time Restrictions.** Specifies when connections can be

active.

8. **Calling Station ID.** Specifies the Network Access Server telephone number.
9. **Client Friendly Name.** Specifies the RADIUS client name.
10. **Client IPv4 Address.** Specifies the RADIUS client IPv4 address.
11. **Client IPv6 Address.** Specifies the RADIUS client IPv6 address.
12. **Client Vendor.** Specifies the RADIUS client vendor.
13. **Called Station ID.** Specifies the NAS identification string or telephone number.
14. **NAS Identifier.** Specifies the NAS identification string.
15. **NAS IPv4 Address.** Specifies the IPv4 address of NAS.
16. **NAS IPv6 Address.** Specifies the IPv6 address of NAS.
17. **NAS Port Type.** Specifies the RADIUS client access type, such as ISDN, VPN, IEEE 802.11 wireless, and so on.

4. Settings are a set of categories that are required. You can select from the following:

    1. **Authentication Methods.** This setting enables you to specify which protocols clients can use to connect to the VPN server. The VPN server attempts to use the most secure protocol first and then, in order of most to least secure, attempts less secure protocols until it reaches the least secure protocol supported by the policy.
    2. **Authentication.** This setting determines whether the VPN server performs authentication or whether authentication credentials are passed to another server.
    3. **Accounting.** This setting enables you to record authentication information to a log file or SQL server database. Recording this information gives you a record of which people have gained access through VPN connections.
    4. **Attribute, RADIUS Attributes Standard & Vendor Specific.**

These settings are necessary only if you are placing the VPN server into an existing RADIUS infrastructure.

Inside OUT

Default conditions and settings

In most situations, the default conditions and settings are likely to be adequate. Most administrators are just worried about ensuring that clients have remote access and are less concerned about whether a client is coming from a particular IPv6 address. Many of the available conditions are really only appropriate for site-to-site links rather than client VPN links. Because most organizations use dedicated hardware devices for site-to-site VPN links, this topic is not covered in this book.

After you ensure that the details of the connection request policies meet your needs, you can configure a network policy. Network policies are useful when you want to give remote access to a large number of users, but you do not want to configure access on a per-account basis. To create a new network policy that allows this access, perform the following general steps:

1. Open the Network Policy Server console. Right-click the Policies\Network Policies node and then click New. The Specify Network Policy Name And Connection Type dialog box opens.
2. Enter a name for the policy and set the Type of Network Access Server drop-down menu to Remote Access Server (VPN-Dial Up) and then click Next.
3. On the Conditions page, click Add and then Add Conditions. The conditions that you can add are as follows:
    1. **Windows Groups.** Limits access to user or computer groups.
    2. **Machine Groups.** Limits access to members of the specified computer group.
    3. **User Groups.** Limits access to members of the specified user

group.

4. **Day and Time Restrictions.** Restricts when the policy can be used for access.
5. **Access Client IPv4 Address.** Specifies the IPv4 network address of the client.
6. **Access Client IPv6 Address.** Specifies the IPv6 network address of the client.
7. **Authentication Type.** Defines acceptable authentication protocols.
8. **Allowed EAP Types.** Defines allowed EAP authentication protocols.
9. **Framed Protocol.** Specifies the protocol used for framing incoming packets, such as Point-to-Point Protocol (PPP) or Serial Line Internet Protocol (SLIP).
10. **Service Type.** Limits clients to a specific service, such as Point-to-Point protocols.
11. **Tunnel Type.** Restricts clients to a specific tunnel type, such as PPTP, L2TP, and SSTP.
12. **Calling Station ID.** Specifies the Network Access Server telephone number.
13. **Client Friendly Name.** Specifies the RADIUS client name.
14. **Client IPv4 Address.** Specifies the RADIUS client IPv4 address.
15. **Client IPv6 Address.** Specifies the RADIUS client IPv6 address.
16. **Client Vendor.** Specifies the RADIUS client vendor.
17. **MS-RAS Vendor.** Specifies the Network Access Server (NAS) vendor.
18. **Called Station ID.** Specifies the NAS identification string or telephone number.
19. **NAS Identifier.** Specifies the NAS identification string.

20. **NAS IPv4 Address.** Specifies the IPv4 address of NAS.
21. **NAS IPv6 Address.** Specifies the IPv6 address of NAS.
22. **NAS Port Type.** Specifies RADIUS client access type, such as ISDN, VPN, IEEE 802.11 wireless, and so on.

After you have configured the network policy, it should be possible to establish a VPN connection.

LAN ROUTING

You can configure Windows Server 2019 to function as a network router in the same way that you configure traditional hardware devices to perform this role. LAN routing is often used with virtual machines to connect private or internal virtual machine networks with an external network. To perform the LAN routing function, Windows Server 2019 must have two or more network adapters. Windows Server 2019 supports using Key Routing Information Protocol v2 (RIP) for route discovery. You can also use the Routing And Remote Access console to configure static routes.

To configure Windows Server 2019 to function as a router, perform the following steps:

1. From the Tools menu of Server Manager, click Routing And Remote Access.
2. In the Routing And Remote Access console, click the server you want to configure. In the Action menu, click Configure And Enable Routing And Remote Access.
3. On the Welcome page of the Routing And Remote Access Server Setup

Wizard, click Next.

4. On the Configuration page, select Custom Configuration and click Next.
5. On the Custom Configuration page, select LAN Routing and click Next.
6. Click Finish. In the Routing And Remote Access dialog box, click Start Service.

If you want to enable IPv6 routing, perform the following additional steps:

1. In the Routing And Remote Access console, right-click the server and click Properties.
2. On the General tab of the server properties dialog box, select IPv6 Router to enable the server to also route IPv6 traffic.

NETWORK ADDRESS TRANSLATION (NAT)

Network Address Translation (NAT) enables you to share an Internet connection with computers on an internal network. In a typical NAT configuration, the NAT server has two network interfaces, where one network interface is connected to the Internet and the second network interface connects to a network with a private IP address range. Computers on the private IP address range can then establish communication with computers on the Internet. It is also possible to configure port forwarding so all traffic that is sent to a particular port on the NAT server's public interface is directed to a specific IP address/port combination on a host on the private IP address range.

To configure a computer running the Windows Server 2019 operating system with two network adapters to function as a NAT device, ensure one adapter has connectivity to the Internet and perform the following steps:

1. Open the Routing And Remote Access console from the Tools menu in Server Manager.
2. Select the server you want to configure. On the Action menu, click Configure And Enable Routing And Remote Access.
3. On the Welcome To The Routing And Remote Access Server Setup Wizard, click Next.
4. On the Configuration page, select Network Address Translation (NAT) and click Next.
5. On the NAT Internet Connection page, select the network interface that connects to the Internet and click Next.
6. Click Finish to close the Routing And Remote Access Server Setup Wizard.

You can configure NAT properties by right-clicking the NAT node in the Routing And Remote Access console and clicking Properties. Using the NAT Properties dialog box, you can configure address assignments for hosts on the private network, as shown in Figure 15-2. You can use the Name Resolution tab to determine how name resolution works on the private network. It enables clients to communicate using single names or FQDNs rather than IP addresses.

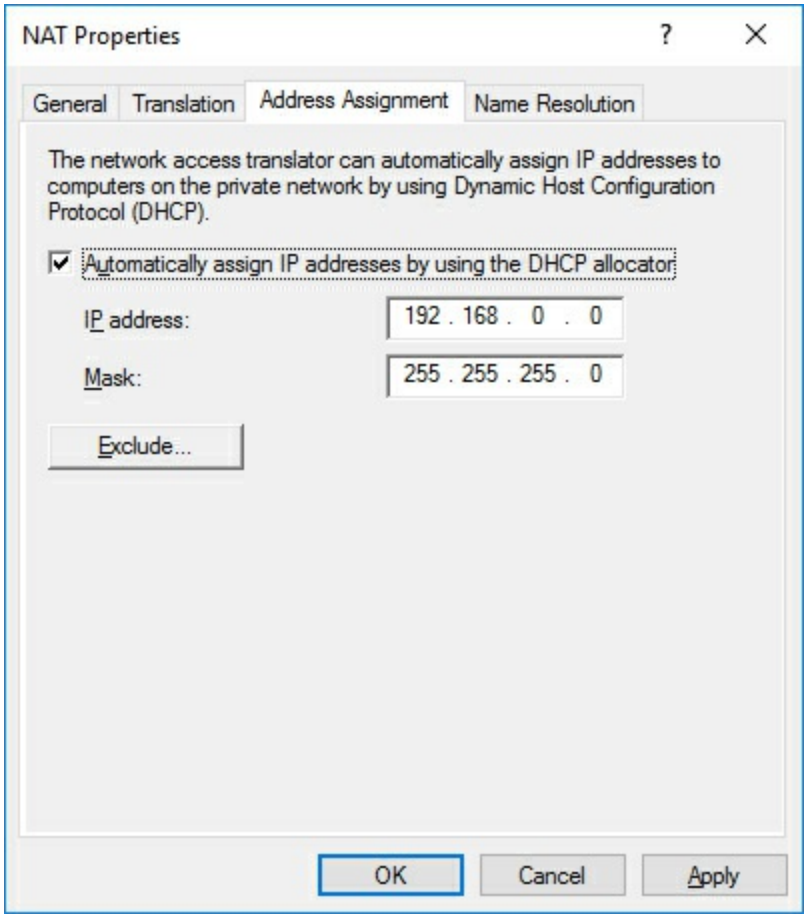


Figure 15-2 NAT address assignment

DIRECTACCESS

DirectAccess is an always-on remote access solution, which means that if a client computer is configured with DirectAccess and connects to the Internet, a persistent connection is established between that computer and the organization's internal network. One big advantage of DirectAccess is that it does not require users to directly authenticate. The client computer determines that it is connected to the Internet by attempting to

connect to a website that is only accessible from the internal network. If the client determines that it is on a network other than the internal network, it automatically initiates a connection to the internal network.

When determining its network connection status, the client first attempts to determine whether it is connected to a native IPv6 network. If the client has been assigned a public IPv6 address, it can make a direct connection to the organization's internal DirectAccess servers. If the client determines that it is not connected to a native IPv6 network, it attempts to create an IPv6 over IPv4 tunnel using the 6to4 and then Teredo technologies. If connections cannot be established using these technologies, it attempts to use IP-HTTPS, which is similar to the SSTP protocol discussed earlier, except that it encapsulates IPv6 traffic within the HTTPS protocol.

When a connection is established, authentication occurs using computer certificates. These computer certificates must be pre-installed on the client, but the benefit of this requirement is that user intervention is not required. Authentication occurs automatically and the connection occurs seamlessly.

Another big advantage of DirectAccess is that it supports remote client management through a feature named

Manage Out. Manage Out enables remote management functionality for DirectAccess clients, allowing administrative tasks, including tasks that in the past could only be performed on computers on a local area network, to be performed on clients.

The primary drawback of DirectAccess is that it requires a domain-joined computer that runs one of the following operating systems:

- Windows 7 Enterprise edition
- Windows 7 Ultimate edition
- Windows 8 Enterprise edition
- Windows 8.1 Enterprise edition
- Windows 10 Enterprise edition

DirectAccess clients are configured through GPOs. The configuration GPO is automatically created through the DirectAccess setup process. This GPO is filtered so it applies only to the security group that you've designated to host the DirectAccess clients.

Unlike the version of DirectAccess that shipped with Windows Server 2008 R2, DirectAccess in Server 2012 and later does not require two consecutive public IPv4 addresses, nor is it necessary to have an Active Directory Certificate Services server deployed on the internal network.

DirectAccess topologies

DirectAccess supports multiple deployment topologies. You don't have to deploy the DirectAccess server with a network adapter directly connected to the Internet. Instead, you can integrate the DirectAccess server with your organization's existing edge topology. During your DirectAccess server deployment, the Remote Access Server Wizard asks you which of the topologies reflects your server configuration. The difference between them is as follows:

- **Edge.** This is the traditional DirectAccess deployment method. The computer hosting the server has two network adapters. The first network adapter is connected directly to the Internet and has been assigned one or more public IPv4 addresses. The second network adapter connects directly to the internal trusted network.
- **Behind An Edge Device (With Two Network Adapters).** In this type of deployment, the DirectAccess server is located behind a dedicated edge firewall. This configuration includes two network adapters. One of the network adapters on the DirectAccess server is connected to the perimeter network behind the edge firewall. The second network adapter connects directly to the internal trusted network.
- **Behind An Edge Device (With A Single Network Adapter).** In this type of deployment, the DirectAccess server has a single network adapter connected to the internal network. The edge firewall passes traffic to the DirectAccess server.

DirectAccess server

The DirectAccess server is a domain-joined computer running Windows Server 2012 or later that accepts connections from DirectAccess clients on untrusted

networks, such as the Internet, and provides access to resources on trusted networks. The DirectAccess server performs the following roles:

- Authenticates DirectAccess clients connecting from untrusted networks
- Functions as an IPsec tunnel mode endpoint for DirectAccess traffic from untrusted networks

Before you can configure a computer running Windows Server 2012 or later to function as a DirectAccess server, you must ensure that it meets the following requirements:

- The server must be a member of an Active Directory Directory Services domain.
- If the server is connected directly to the Internet, it must have two network adapters: one that has a public IP address and one that is connected to the trusted internal network.
- The DirectAccess server can be deployed behind a NAT device, which limits DirectAccess to use IP over HTTPS (IP-HTTPS).
- A server connected to the Internet requires only a single public IPv4 address. Two-Factor Authentication, using either Smart Card or One-Time Password (OTP), however, requires two consecutive public IPv4 addresses.
- The DirectAccess server can also host a VPN server. This functionality was not present in the Windows Server 2008 R2 version of DirectAccess.
- You can configure DirectAccess in a network load-balanced configuration of up to eight nodes.
- The TLS certificate installed on the DirectAccess server must contain an FQDN that resolves through DNS servers on the

Internet to the public IP address assigned to the DirectAccess server or to the gateway through which the DirectAccess server is published.

- The TLS certificate installed on the DirectAccess server must have a Certificate Revocation List (CRL) distribution point that is accessible to clients on the Internet.

> **Inside OUT**
>
> **DirectAccess TLS Certificate**
>
> You should strongly consider obtaining the TLS certificate for your organization's DirectAccess server from a public CA. If you obtain the certificate this way, you don't have to worry about publishing the CRL from your internal certificate services deployment to a location that is accessible to the Internet. Using a trusted third-party CA ensures that the CRL is available to clients on the Internet.

A DirectAccess implementation also relies on the following infrastructure being present:

- **Active Directory domain controller.** DirectAccess clients and servers must be members of an Active Directory domain. By necessity, when you deploy a domain controller, you also deploy a DNS server. Active Directory, due to its basic nature, also makes Group Policy available.
- **Group Policy.** When you configure DirectAccess, the setup wizard creates a set of Group Policy Objects (GPOs) that are configured with settings you choose in the wizard. They apply to DirectAccess clients, DirectAccess servers, and servers that you use to manage DirectAccess.

Prepare DNS servers by removing the ISATAP name from the global query block list. You must take this step on all DNS servers that are hosted on computers running Windows Server operating systems. To

complete this step, remove ISATAP from the GlobalQueryBlockList value on the Computer\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DNS\Parameters hive of the registry, so that it contains only the wpad entry, as shown in Figure 15-3. After making this configuration change, you must restart the DNS server.

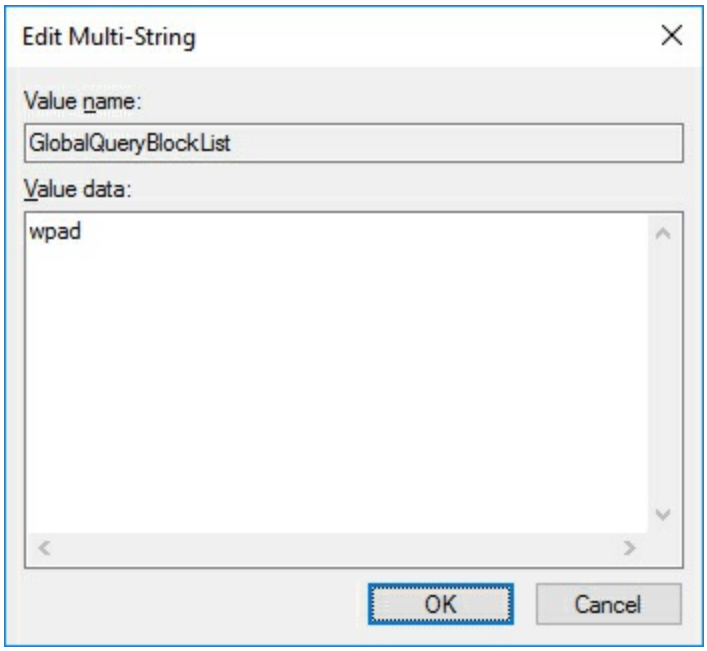


Figure 15-3 GlobalQueryBlockList registry entry

As an alternative to editing the registry, you can also remove ISATAP from the DNS global query block list by issuing the following command on each DNS server and restarting the DNS Server service:

Click here to view code image

```
Dnscmd /config /globalqueryblocklist wpad
```

Network Location Server

The Network Location Server (NLS) is a specially configured server that enables clients to determine whether they are on a trusted network or an untrusted network. The NLS server's only function is to respond to specially crafted HTTPS requests. When the client determines that it has a connection to any network, it sends this specially crafted HTTPS request. If there is a response to this request, the client determines that it is on a trusted network and disables the DirectAccess components. If there is no response to this request, the client assumes that it is connected to an untrusted network and initiates a DirectAccess connection.

DirectAccess clients are informed of the NLS's location through Group Policy. You don't have to configure these policies manually because they are created automatically when you use the DirectAccess Setup Wizard. Any server that hosts a website and has an SSL certificate installed can function as the NLS. You should ensure that the NLS is highly available because if this server fails, it causes all clients that are configured for DirectAccess on the trusted network to assume that they are on an untrusted network.

Configuring DirectAccess

After you understand the infrastructure requirements,

configuring DirectAccess is straightforward. To configure DirectAccess, perform the following steps:

1. Create a security group in Active Directory and add all the computer accounts for computers that are DirectAccess clients. This security group can have any name, but a name such as DirectAccess_Clients makes it easy to remember why you created it.
2. Ensure that you configure DNS with the following information:
    1. The externally resolvable DNS zone needs to have a record mapping the FQDN of the DirectAccess server's external interface to the public IPv4 address of the DirectAccess server.
    2. If you use a certificate issued by your organization's CA, ensure that a DNS record exists for the CRL location.
    3. The internal DNS zone needs a record mapping the name of the NLS to an IP address.
    4. Remove ISATAP from the global query block list on all DNS servers in the organization.
3. If you use your organization's CA, you must configure an appropriate certificate template, as well as deploy a CRL distribution point in a location that can be accessed by clients on the Internet. The certificate template can be a duplicate of the Web Server Certificate Template. You can use this certificate for both the NLS's TLS certificate and the IP-HTTPS certificate for the DirectAccess server. If you don't use your organization's CA to issue certificates, you can use certificates from a public CA for the NLS and the DirectAccess server.
4. Configure firewall rules for all hosts on the trusted network that should be accessible to DirectAccess clients, so they enable inbound and outbound ICMPv6 echo requests. You can configure these rules in a GPO that applies to hosts that DirectAccess clients can access. The rules should have the following properties:
    1. Rule Type: Custom
    2. Protocol Type: ICMPv6

3. Specific ICMP Types: Echo Request

5. Install the Remote Access role on the computer that functions as the DirectAccess server.
6. Open the Remote Access console and then choose between running the Getting Started Wizard and the Remote Access Setup Wizard. The Getting Started Wizard enables administrators to quickly deploy DirectAccess by only requiring a minimal amount of information. The Remote Access Setup Wizard requires more detailed responses but enables administrators to customize their deployment. The rest of this procedure deals with the Remote Access Setup Wizard.
7. The Configure Remote Access page of the Configure Remote Access Wizard enables you to choose between deploying DirectAccess and VPN, deploying DirectAccess only, or deploying VPN only.
8. If you select Deploy DirectAccess Only, you are provided with the Remote Access Setup diagram. This diagram involves a series of steps that enable you to configure the DirectAccess server, clients, and infrastructure. There are four steps, including:

    1. Step 1: Remote Clients
    2. Step 2: Remote Access Server
    3. Step 3: Infrastructure Servers
    4. Step 4: Application Servers

Step 1: Remote clients

The Step 1: Remote Clients section of Remote Access Setup enables you to configure which computers function as DirectAccess clients. When you click the Configure button in the Step 1 area, a three-page wizard appears that enables you to configure the following settings:

1. Choose Deploy Full DirectAccess For Client Access And Remote

Management or Deploy DirectAccess For Remote Management Only. If you choose the first option, the people using DirectAccess clients can access internal network resources when they have an active Internet connection. If you choose the second option, you can perform management tasks on the computer when it's connected on the Internet, but the user can't access internal resources.

2. Select which security groups containing computer accounts you want to enable for DirectAccess. On this page, you can choose Enable DirectAccess For Mobile Computers Only and Use Force Tunneling. When you enable force tunneling, computers designated as DirectAccess clients connect through the remote access server when they connect to both the Internet and the internal trusted network.
3. On the Network Connectivity Assistant (NCA) page, you can configure connectivity information for clients, such as providing the DirectAccess connection name, the help desk's email address, and whether DirectAccess clients use local name resolution.

Step 2: Remote Access Server

The Step 2: Remote Access Server section of the Remote Access Setup diagram has a three-page wizard that enables you to do the following:

1. Configure the network topology and specify the public name or IPv4 address that clients use to connect to DirectAccess. The topology options are Edge, Behind Edge (Two Network Adapters), and Behind Edge (Single Network Adapter), which were discussed in the DirectAccess topology section of this chapter.
2. On the Network Adapters page, verify the network adapter configuration. You can also choose which certificate to use to authenticate IP-HTTPS connections. The certificate you choose should be a typical SSL certificate that uses an FQDN that clients use for connections. You can choose to use a self-signed certificate, although this is not recommended except on test deployments.
3. On the authentication page, choose whether you want to use Active

Directory or Two-Factor Authentication. You can also configure authentication to use computer certificates, but when you do this, you must specify the CA that the computer certificates must be issued from.

Step 3: Infrastructure servers

After you have configured the remote access clients and the remote access server, the next step is to configure infrastructure servers. The Infrastructure Server Setup Wizard takes you through the following steps:

1. On the first page, specify the location of the NLS by using the server's URL, and if you are specifying a separate server, remember to use https rather than http in the address. You also have the option to configure the DirectAccess server as the remote access server and use a self-signed certificate; however, self-signed certificates are more appropriate for tests rather than production deployments.
2. The DNS page enables you to specify the DNS suffixes that should be used with the name resolution and the address of the internal DNS server. On this page, you can also configure how clients should use the DNS server of their local Internet connection, by choosing one of the following options:
    1. Use the DNS server of the local connection if the name isn't resolvable using the DNS server on the trusted network.
    2. Use the DNS server of the local connection if the name isn't resolvable using the DNS server on the trusted network or if the DNS server on the trusted network cannot be contacted.
    3. Use the DNS server of the local connection if any DNS error occurs.
3. The DNS Suffix Search List enables you to configure any DNS suffixes that should be used by the client for any unqualified names. The default settings add the domain name suffix.
4. The Management Servers page enables you to configure the servers used for DirectAccess client management. You can also configure NAP

remediation servers if you are using NAP with DirectAccess.

Step 4: Application servers

Step 4 of the Remote Access Management Console setup enables you to configure the addresses of application servers that require end-to-end authentication when interacting with DirectAccess clients. Unlike the other steps, this step involves configuring only one dialog box, where you specify the security group containing the computer accounts that you want to require end-to-end authentication and encryption for. You can also use this dialog box to limit DirectAccess clients so they can only connect to servers in the listed groups and can't connect to other servers on the trusted network. Use this option in environments with stringent security requirements.

MANAGING REMOTE ACCESS USING POWERSHELL

The Remote Access PowerShell module contains a large number of cmdlets that you can use to manage all aspects of Routing and Remote Access on Windows Server 2019. These cmdlets are described in Table 15-1.

Table 15-1 Remote Access PowerShell cmdlets

Noun	Verbs	Function
BgpCustomRoute	Add, Remove,	View and manage BGP (Border Gateway Protocol) route information.

Technet24

	Get	
BgpPeer	Get, Remove, Set, Start, Add, Stop	View and manage BGP peer configuration.
BgpRouteAggregate	Get, Set, Add, Remove	View and manage aggregate BGP routes.
BgpRouteFlapDampening	Clear, Enable, Get, Set, Disable	Manage the configuration of a BGP route dampening engine.
BgpRouteInformation	Get	View BGP route information.
BgpRouter	Get, Add, Remove, Set	Manage configuration information for BGP routers.
BgpRoutingPolicy	Set, Add, Remove, Get	Manage BGP routing policies.
BgpRoutingPolicyForPeer	Set, Remove, Add	Manage BGP peer routing policies.
BgpStatisti	Get	Retrieves BGP peering-related message and

cs		route advertisement statistics.
DAAppServer	Remove, Get, Add	Manage DirectAccess application server security groups.
DAAppServerConnection	Set	Configures the properties of the connection to application servers and the IPsec security traffic protection policies for the connection.
DAClient	Get, Set, Remove, Add	Manage the client security groups used for DirectAccess client configuration.
DAClientDnsConfiguration	Get, Remove, Set, Add	Manage Name Resolution Policy Table (NRPT) entries and local name resolution properties.
DAEntryPoint	Add, Get, Set, Remove	Manage DirectAccess entry-point settings.
DAEntryPointDC	Set, Get	Configure DirectAccess entry point domain controller settings.
DAMgmtServer	Update, Add, Remove, Get	Manage DirectAccess management servers.
DAMultiSite	Disable, Enable, Get, Set	Manage global settings applied to all DirectAccess entry points in a multisite deployment.

Technet24

DANetworkLocationServer	Get, Set	Manage DirectAccess NLS settings.
DAOtpAuthentication	Disable, Set, Enable, Get	Manage DirectAccess one-time password settings.
DAServer	Get, Set	Manage DirectAccess Server Settings.
RemoteAccess	Get, Install, Uninstall, Set	Manage DirectAccess and VPN configuration.
RemoteAccessAccounting	Get, Set	Configure remote access accounting.
RemoteAccessConfiguration	Get, Set	Manage remote access configuration.
RemoteAccessConnectionStatistics	Get	View information about DirectAccess and VPN connections.
RemoteAccessConnectionStatisticsSumm	Get	View summary information about DirectAccess and VPN connections.

ary		
RemoteAccessHealth	Get	View the current health of a RemoteAccess (RA) deployment.
RemoteAccessInboxAccountingStore	Set, Clear	Manage the inbox of the remote access accounting store.
RemoteAccessIpFilter	Add, Remove, Set, Get	Manage remote access IP filters.
RemoteAccessLoadBalancer	Get, Set	Manage remote access load balancer settings.
RemoteAccessLoadBalancerNode	Remove, Add	Add and remove remote access load balancer nodes.
RemoteAccessRadius	Add, Get, Set	Manage RADIUS servers.
RemoteAccessRoutingDomain	Get, Set, Disable, Enable	Manage remote access routing domains.
RemoteAccessUserActivity	Get	Displays the resources accessed over the active DirectAccess and VPN connections and the resources accessed over historical

Technet24

		DirectAccess and VPN connections.
RoutingProtocolPreference	Set, Get	Manage routing protocol preferences.
VpnAuthProtocol	Set, Get	Manage VPN server authentication parameters.
VpnAuthType	Set	Configure VPN server authentication types.
VpnIPAddressAssignment	Set	Manage VPN server IPv4 address assignment.
VpnIPAddressRange	Remove, Add	Manage VPN server address range.
VpnS2SInterface	Set, Add, Connect, Remove, Disconnect, Get	Manage site-to-site VPN interfaces.
VpnS2SInterfaceStatistics	Clear, Get	View site-to-site VPN interface statistics.
VpnServerConfiguration	Get, Set	Manage VPN server configuration.

VpnSstpProxyRule	Add, Get, New, Set, Remove	Manage Tenant ID to gateway mappings.
VpnTrafficSelector	New	Creates a VPN traffic selector object that configures the IKE traffic selector.
VpnUser	Disconnect	Disconnects a VPN connection originated by a specific user or originating from a specific client computer.

Chapter 16

Remote Desktop Services

Deployment

Remote Desktop Connection Broker

Deployment properties

Remote Desktop Session Host

Remote Desktop Virtualization Host

Remote Desktop Web Access

Remote Desktop licensing

Managing Remote Desktop Services using PowerShell

Remote Desktop Services (RDS) allows an application to run on one computer with the graphical output of that computer displayed on another computer. The benefits of RDS include allowing relatively resource-poor clients to run resource-intensive applications remotely, or to

run applications that are unable to run on one operating system to be accessed by users of another operating system.

There are several ways that you can leverage RDS in Windows Server 2019. The first is using Remote Desktop Session Host. In this form of RDS, people connect to remote desktop sessions that run on a server in a special isolated environment. In versions of Windows Server prior to Windows Server 2008 R2, Remote Desktop Session Host was called Terminal Services. The other primary way to leverage RDS in Windows Server 2019 is through Remote Desktop Virtualization Host. In this form of RDS, users connect to virtual machines to use their desktop environment to run their applications. Remote Desktop Virtualization Host is a form of Virtual Desktop Infrastructure, or VDI.

DEPLOYMENT

Remote Desktop Services deployment differs from other roles and features deployment. With all other roles and features, you select the role-based or feature-based installation option when running the Add Roles and Features Wizard. With Remote Desktop Services, you instead select the Remote Desktop Services installation option as shown in Figure 16-1.

Technet24

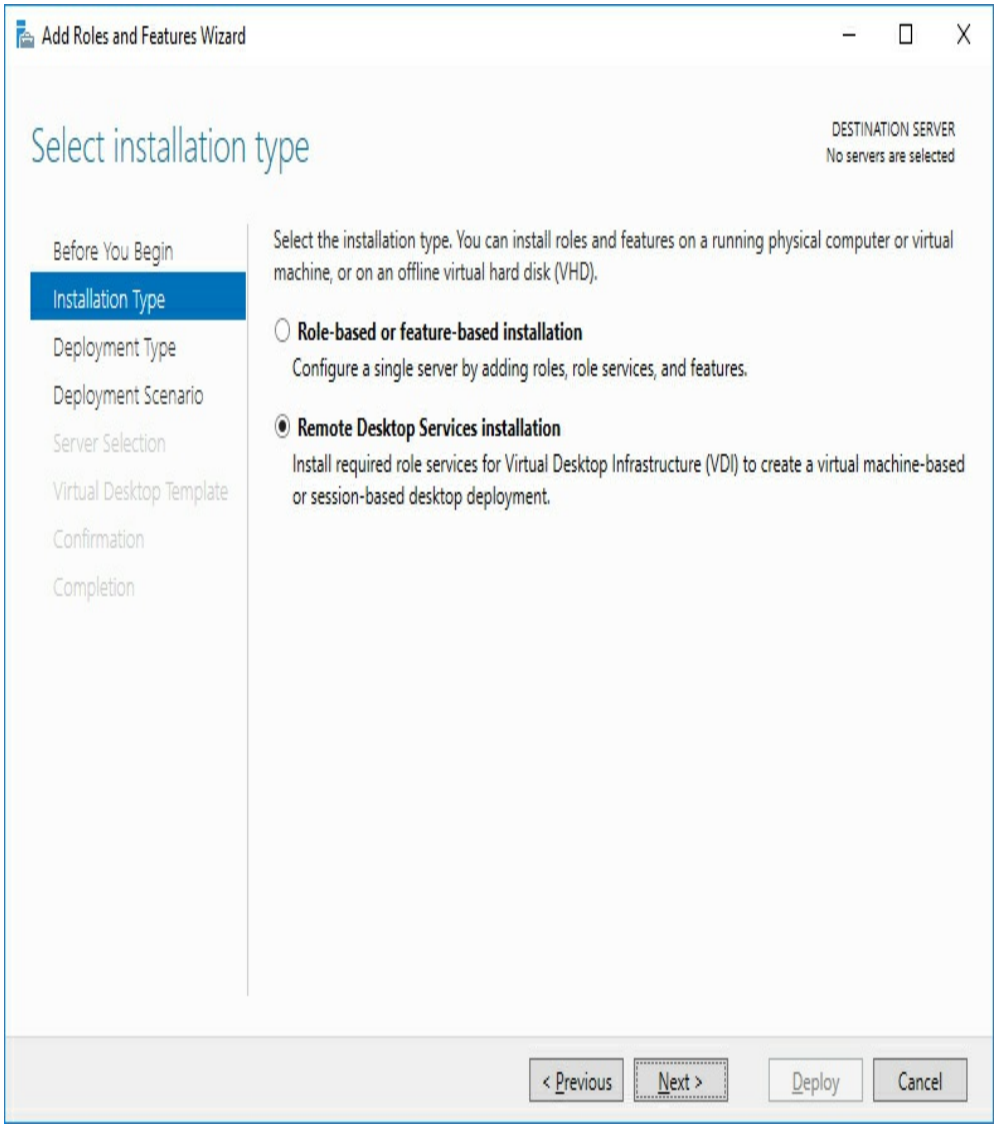


Figure 16-1 Installation type options

Once you've selected Remote Desktop Services installation, the next page in the wizard allows you to choose between the following options:

- **Standard deployment.** Choose this when you are deploying Remote Desktop Services across multiple servers.
- **Quick Start.** Use this option if you are deploying a single Remote

Desktop Services server.

Another option is to use Multipoint Services, which is a Feature On Demand. Use the MultiPoint Services option if you are creating a multipoint server deployment. Multipoint server allows you to run multiple client terminals connected to a single multipoint server. Multipoint services is commonly used in classroom environments.

When you choose the Standard deployment or Quick Start options, you can then select the deployment scenario as shown in Figure 16-2.

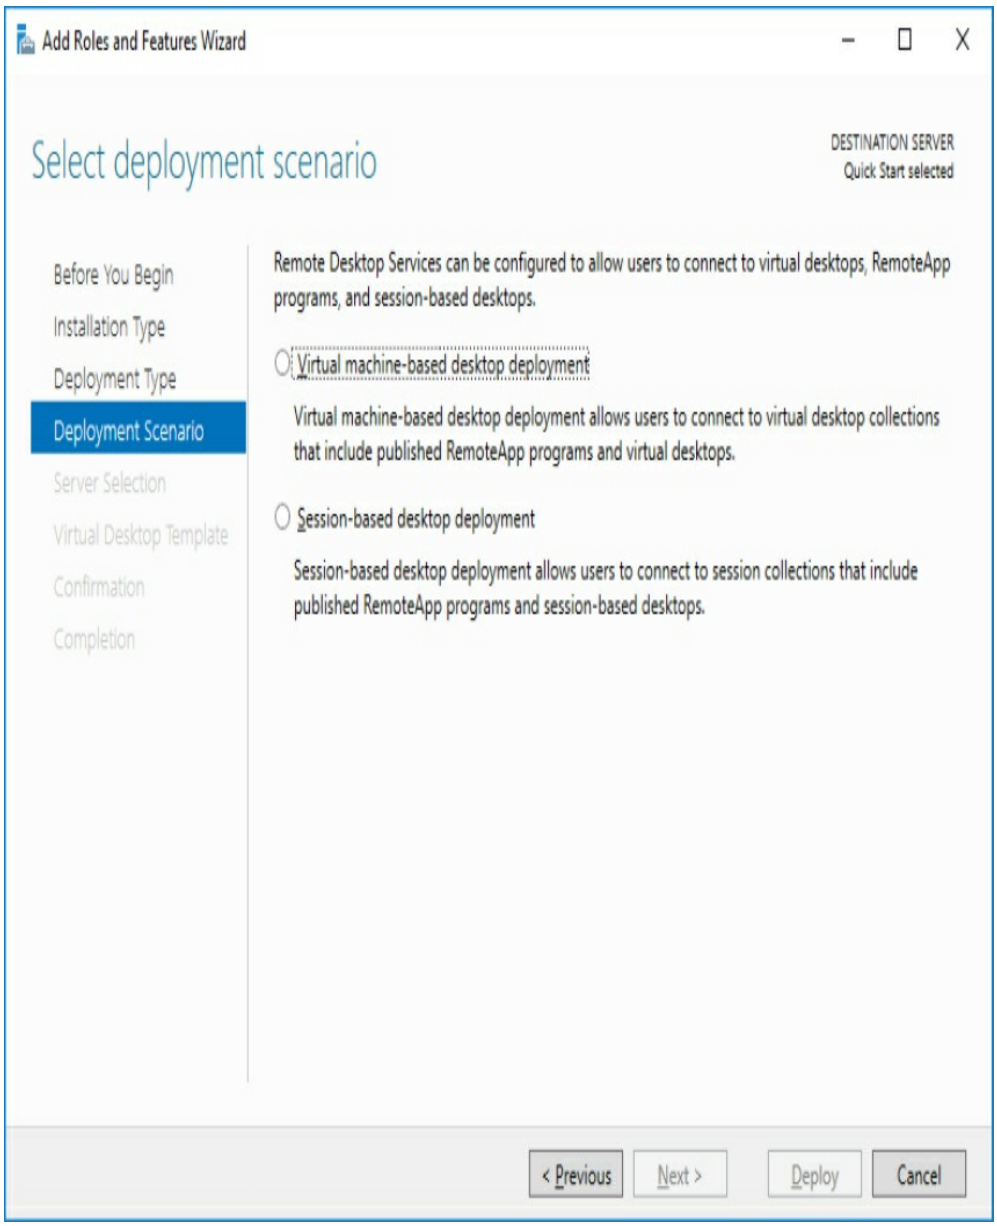


Figure 16-2 Deployment scenario

Session-based deployments are generally less resource intensive than virtualization host deployments because they don't require that the session host server host full virtual machines for each session. Virtualization host

deployments are more likely to provide better application compatibility because there are some applications that don't support being deployed in a Session host-based environment.

Depending on your selection, completing the wizard involves selecting the server that the deployment will occur on, selecting a virtualization template for a virtualization host deployment, installing the roles, and then rebooting the server to complete installation.

REMOTE DESKTOP CONNECTION BROKER

When initiating a connection, either through Remote Desktop Connection client software or the RD Web Access interface, users connect to a session on a Remote Desktop Session Host server or a VDI virtual machine on a Remote Desktop Virtualization Host server through the Remote Desktop Connection Broker. The Remote Desktop Connection Broker role service provides the following functionality:

- Connects incoming remote desktop clients to the appropriate RD Session Host or Virtualization Host server.
- Reconnects disconnected remote desktop clients to the appropriate RD Session Host session or RD Virtualization Host virtual machine.
- Load balances incoming remote desktop client connections according to session collection settings.
- Allows organizations to scale out the RD Session Host or RD

Virtualization Host deployment by adding RD Session Host or RD Virtualization Host servers.

In previous versions of Remote Desktop and Terminal Services, it was possible to directly connect to the session host server. In Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019, all connections are mediated by the Remote Desktop Connection Broker.

DEPLOYMENT PROPERTIES

Deployment properties allow you to configure settings for both RD Session Host and RD Virtualization Host servers. Available through the Collections node of the Remote Desktop Services section of the Server Manager console, you can use Deployment Properties to configure the following settings:

- **RD Gateway.** This allows you to specify RD Gateway settings for your deployment. This includes whether the RD Gateway server is automatically determined, whether one is assigned, or whether to block the use of an RD Gateway server and the sign-in method when using an RD Gateway server.
- **RD Licensing.** This allows you to configure the licensing mode and licensing servers. You can configure multiple RD license servers and the order in which they will be queried to obtain a license for a Remote Desktop Services client.
- **RD Web Access.** The address of the RD Web Access server. Clients can use this to connect to RemoteApp and Remote Desktop sessions using a browser, instead of through the Remote Desktop Connection client.

- **Certificates.** Shown in Figure 16-3, use this section to configure certificates for the Remote Desktop Services deployment.

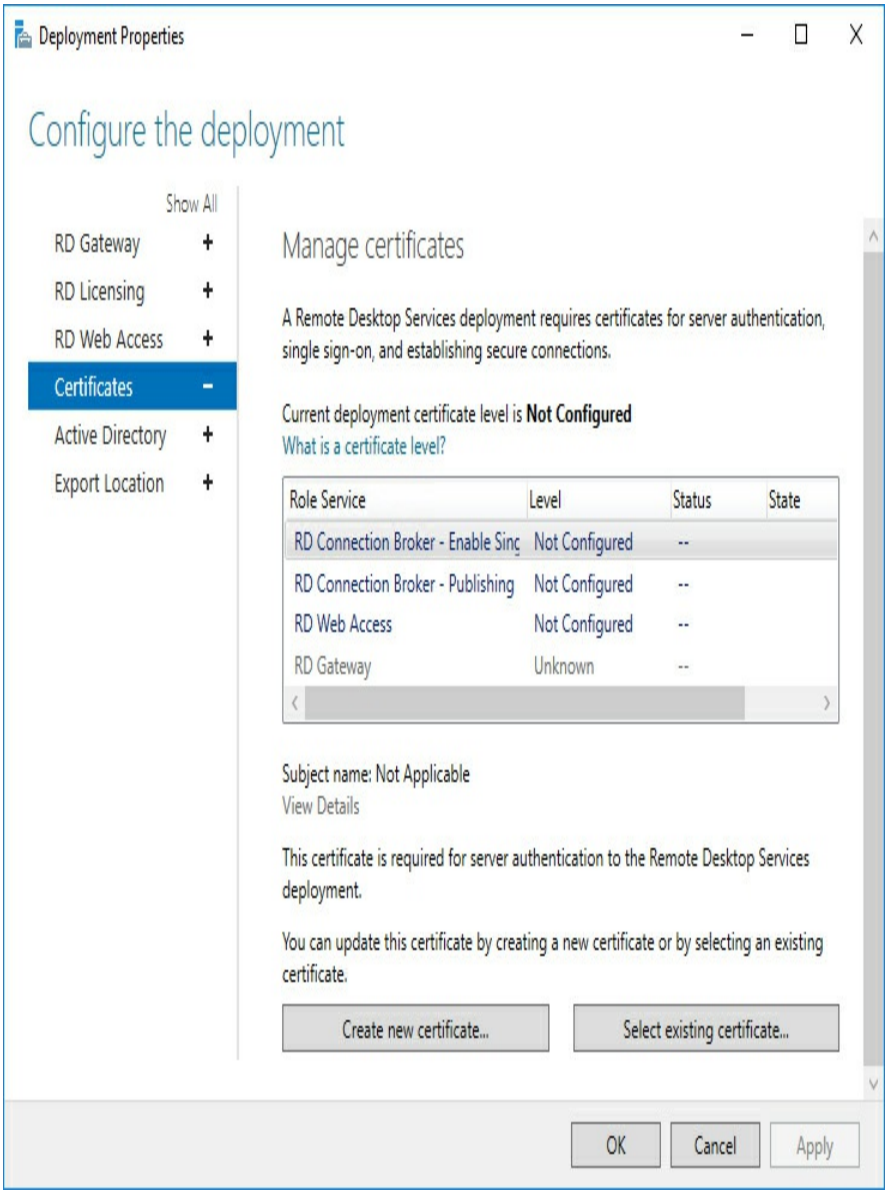


Figure 16-3 Certificates

- **Active Directory.** This is for RD Virtualization Host servers and allows you to configure the domain and organizational unit that computer accounts for VDI virtual machines will be added to once they are joined to the AD DS domain.

Technet24

- **Export Location.** The location where the virtual desktop template is copied to once selected for use when generating further RD Virtualization Host virtual machines.

REMOTE DESKTOP SESSION HOST

Remote Desktop (RD) Session Host provides a full desktop environment to a user. Users connect to the RD Session Host server from a thin client, such as a Windows Thin PC or the Remote Desktop Connection client, which ships with Windows operating systems. An RD Session Host provides a full screen or window that reproduces a Windows desktop, including a taskbar, Start menu, and desktop icons.

The benefits of using RD Session Host include:

- You can deploy an application to many users by installing it on a small number of RD Session Host servers. Users can run applications on RD Session Host servers that might be incompatible with their operating systems.
- Application updates are simplified. You update an application on the RD Session Host server rather than having to push the patch down to individual operating systems.
- It allows you to grant access to a standardized corporate desktop experience for users who have chosen to use their own computer at work in BYOD (Bring Your Own Device) scenarios.
- It allows you to grant access, through a VPN, to a standardized corporate desktop for users at remote locations or for those who telecommute.
- It allows you to grant access to applications for operating systems

that do not support RD Web Access or RD RemoteApp.

The drawbacks to RD Session Host servers is that they require applications that are compatible with the RD Session Host environment. These applications must be installed after the session host role is deployed. You should ensure that applications are installed on all servers that are available to a specific session collection because you won't know which RD Session Host server a user connecting to a session collection will be connected to.

Session collection settings

Session collections determine the properties of RD Session Host servers. You can configure the following RD Session Host settings on a per session collection basis:

- **User Groups.** These settings determine which user groups are associated with a specific session collection. If a user is a member of the user group, they can connect to session host servers that are members of the collection. They can also access RemoteApp applications published to the collection.
- **Session.** Shown in Figure 16-4, these settings determine the amount of time after which a session that is disconnected is automatically ended. The default value is Never. You can be configure the value to be as short as one minute. You can also configure active session limits, how long before an idle session is disconnected, and then determine what happens when a disconnection occurs: should the session be automatically terminated or should reconnection be allowed. You can also use the settings contained in this area to configure whether to use

Technet24

temporary folders and whether those folders are deleted when the session ends.

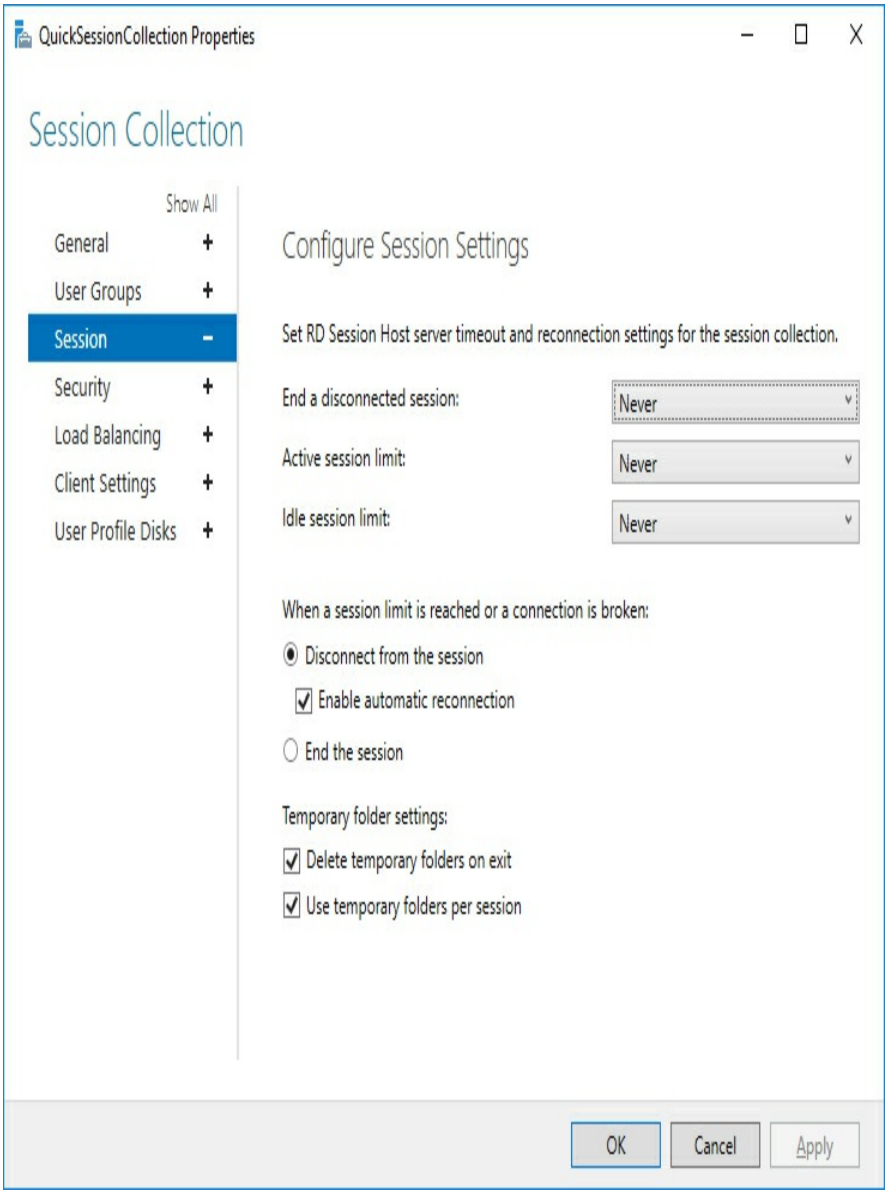


Figure 16-4 Session settings

- **Security.** These settings allow you to configure whether Remote Desktop Connections require RDP Security or TLS security. Encryption options include Low, Client Compatible, High, and FIPS Compliant. You should generally select High unless you are

required to use FIPS Compliant. The drawback of doing so is that some older third-party remote desktop clients may not support this level of encryption.

- **Load Balancing.** This section allows you to configure the load balancing settings for a session collection, including the relative weight the server is assigned when multiple session host servers are available to the session collection and the limit to the number of sessions each session host server can host.
- **Client Settings.** The Client Settings section allows you to configure which redirection options to enable. You can configure redirection for the following:
    - Audio and video playback
    - Audio recording
    - Smart cards
    - Plug-and-play devices
    - Drives
    - Clipboard
    - Printers, including using the default client printing device and the Remote Desktop easy print driver
    - Monitor redirection
- **User Profile Disks.** User profile disks allow users to store settings and folders in a central location. This is useful to allow users to persist data and settings across different session host settings on multiple servers. You can choose the following user profile disk settings:
    - User profile disk location
    - Maximum size
    - Store all user settings and data on the profile disk, with specified folders excluded

- Store only certain folders on the user profile disk

You can also use Session Collection settings to configure which remote app programs will be available to the Session Collection. You can also add and remove session host servers from the session collection. These servers should already have the RD Session Host role installed prior to performing this task.

Personal session desktops

Personal session desktops are a feature of Windows Server 2019 that allow you to create a special type of Session Collection where each user has access to their own personal session host where they enjoy administrative rights. Personal session desktops can be useful for developers who might need administrative rights without giving them those rights over a Remote Desktop Session Host server.

RemoteApp

Rather than open a full session, including a desktop, Start menu, and a taskbar, RemoteApp allows a user to run one or more applications directly. The application runs in such a way that it appears as if it is executing on the local computer. When a client invokes multiple applications from the same host server, the server transmits the applications to the client using the same session.

RemoteApp includes the following benefits:

- Any application that can be run on the RD Session Host server can be run as a RemoteApp application.
- You can associate RemoteApp applications with local file extensions. This means that the RemoteApp application starts when the user attempts to open an associated document.
- Users are often unaware that an application is running remotely.
- RemoteApp applications can be used with RD Gateway, allowing access to clients on remote networks such as the Internet.

You publish a RemoteApp program to a Session Collection once you've installed it on a RD Session Host server through the session collection area of the Server Manager console.

Group Policy configuration

The Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host server node of Group Policy contains several policies that you can use to configure Session Host servers. There are many settings that can only be configured through group policy and PowerShell that are not available in the Remote Desktop Services area of the Server Manager Console. The coverage here is brief, as there are a vast number of policies, and you've already learned about the functionality that most of them implement.

- **Application Compatibility.** Through the Application Compatibility node, you can access the Turn Off Windows Installer

RDS Compatibility, Turn On Remote Desktop IP Virtualization, Select The Network Adaptor To Be Used For Remote Desktop IP Virtualization, and Do Not Use Remote Desktop Session Host Server IP Address When Virtual IP Address Is Not Available.

- **Connections.** Through the Connections node, you can access the following policies: Automatic Reconnection, Allow Users To Connect Remotely Using Remote Desktop Services, Deny Logoff Of An Administrator Logged In To The Console Session, Configure Keep-Alive Connection Interval, Limit Number Of Connections, Suspend User Sign-In To Complete App Registration, Set Rules For Remote Control Of Remote Desktop Services User Sessions, Select RDP Transport Protocols, Select Network Detection On The Server, Restrict Remote Desktop Services Users To A Single Remote Desktop Services Session, Allow Remote Start Of Unlisted Programs, and Turn Off Fair Share CPU Scheduling.
- **Device and Resource Redirection.** Through the Device and Resource Redirection node, you can configure the following policies: Allow Audio And Video Playback Redirection, Allow Audio Recording Redirection, Limit Audio Playback Quality, Do Not Allow Clipboard Redirection, Do Not Allow COM Port Redirection, Do Not Allow Drive Redirection, Do Not Allow LPT Port Redirection, Do Not Allow Supported Plug And Play Device Redirection, Do Not Allow Smart Card Device Redirection, and Allow Time Zone Redirection.
- **Licensing.** You can access the Licensing node to configure the following policies: Use The Specified Remote Desktop License Servers, Hide Notifications About RD Licensing Problems That Affect The RD Session Host Server, and Set The Remote Desktop Services Licensing Mode.
- **Printer Redirection.** Through the Printer Redirection node, you can configure the following policies: Do Not Set Default Client Printer To Be Default Printer In A Session, Do Not Allow Client Printer Redirection, Use Remote Desktop Easy Print Printer Driver First, Specify RD Session Host Server Fallback Printer Driver

Behavior, and Redirect Only The Default Client Printer.

- **Profiles.** You can find the following policies in the Profiles node: Limit The Size Of The Entire Roaming User Profile Cache, Set Remote Desktop Services User Home Directory, Use Mandatory Profiles On The RD Session Host Server, and Set Path For Remote Desktop Services Roaming User Profile.
- **RD Connection Broker.** You can find the following policies in the RD Connection Broker node: Join RD Connection Broker, Configure RD Connection Broker Farm Name, Use IP Address Redirection, Configure RD Connection Broker Server Name, and Use RD Connection Broker Load Balancing.
- **Remote Session Environment.** You can use the Remote Session Environment node to locate the following policies: Limit Maximum Color Depth, Enforce Removal Of Remote Desktop Wallpaper, Use The Hardware Default Graphics Adapter For All Remote Desktop Services Sessions, Limit Maximum Display Resolution, Limit Number Of Monitors, Remove Disconnect Option From Shut Down Dialog, Remove Windows Security Item From Start Menu, Use Advanced RemoteFX Graphics For RemoteApp, Prioritize H.264/AVC 444 Graphics Mode For Remote Desktop Connections, Configure H.264/AVC Hardware Encoding For Remote Desktop Connections, Configure Compression For RemoteFX Data, Configure Image Quality For RemoteFX Adaptive Graphics, Enable RemoteFX Encoding For RemoteFX Clients Designed For Windows Server 2008 R2 SP1, Configure RemoteFX Adaptive Graphics, Start A Program On Connection, Always Show Desktop On Connection, Allow Desktop Composition For Remote Desktop Sessions, and Do Not Allow Font Smoothing.
- **RemoteFX for Windows Server 2008 R2.** This node contains the following policies: Configure Remote FX, Optimize Visual Experience When Using RemoteFX and Optimize Visual Experience For Remote Desktop Service Sessions.
- **Security.** The Security node contains the following policies: Server Authentication Certificate Template, Set Client Connection

Encryption Level, Always Prompt For Password Upon Connection, Require Secure RPC Communication, Require Use Of Specific Security Layer for Remote (RDP) Connections, Do Not Allow Local Administrators To Customize Permissions, and Require User Authentication For Remote Connections By Using Network Level Authentication.

- **Session Time Limits.** You can find the following policies in the Session Time Limits node: Set Time Limit For Disconnected Sessions, Set Time Limit For Active But Idle Remote Desktop Services Sessions, Set Time Limit For Active Remote Desktop Services Sessions, End Session When Time Limits Are Reached, and Set Time Limit For Logoff Of RemoteApp Sessions.
- **Temporary Folders.** The Temporary Folders node contains the following policies: Do Not Delete Temp Folder Upon Exit and Do Not Use Temporary Folders Per Session.

Inside OUT

Useful policies

The most useful policy groups are in the Connections and Session Time Limits nodes. These contain policies that control how many clients are connected to each RD Session Host server and how long they remain connected, affecting the RD Session Host server's capacity.

REMOTE DESKTOP VIRTUALIZATION HOST

The Remote Desktop Virtualization Host allows you to configure Remote Desktop technologies to support Virtual Desktop Infrastructure (VDI). You can use RD Virtualization Host to allow clients to connect using Remote Desktop to personal virtual machines running on Hyper-V hosts. You can configure RD Virtualization

Host to assign users to a specific unique virtual machine or you can configure RD Virtualization Host to direct users to a pool of shared virtual machines. The RD Virtualization Host role service must be installed on a computer that hosts the Hyper-V role and the client virtual machines.

This type of virtualization takes substantially more resources than RD Session Host sessions but also allows greater customization of the user experience. Virtual machines can be treated just like desktop machines, with users using a thin client to access the RD Virtualization Host virtual machines. RD Virtualization Host integrates with RD Session Broker to ensure that users are redirected to their existing virtual machine in the event that their session is disconnected.

RD Virtualization Host computers have the following requirements:

- Must be available remotely via Windows PowerShell.
- Must be domain joined.
- Must be running Hyper-V.
- Server must support hardware-assisted virtualization. This means that you cannot run an RD Virtualization Host server within Hyper-V nested virtualization. Nested virtualization is supported by Windows Server 2016 and Windows Server 2019.
- Server must have a network adapter that supports the creation of a Hyper-V virtual switch.
- RD Connection Broker server must have the right to join virtual

desktops to the AD DS domain.

- If performing a quick start deployment, a prepared template computer desktop computer virtual hard disk as shown in Figure 16-5.

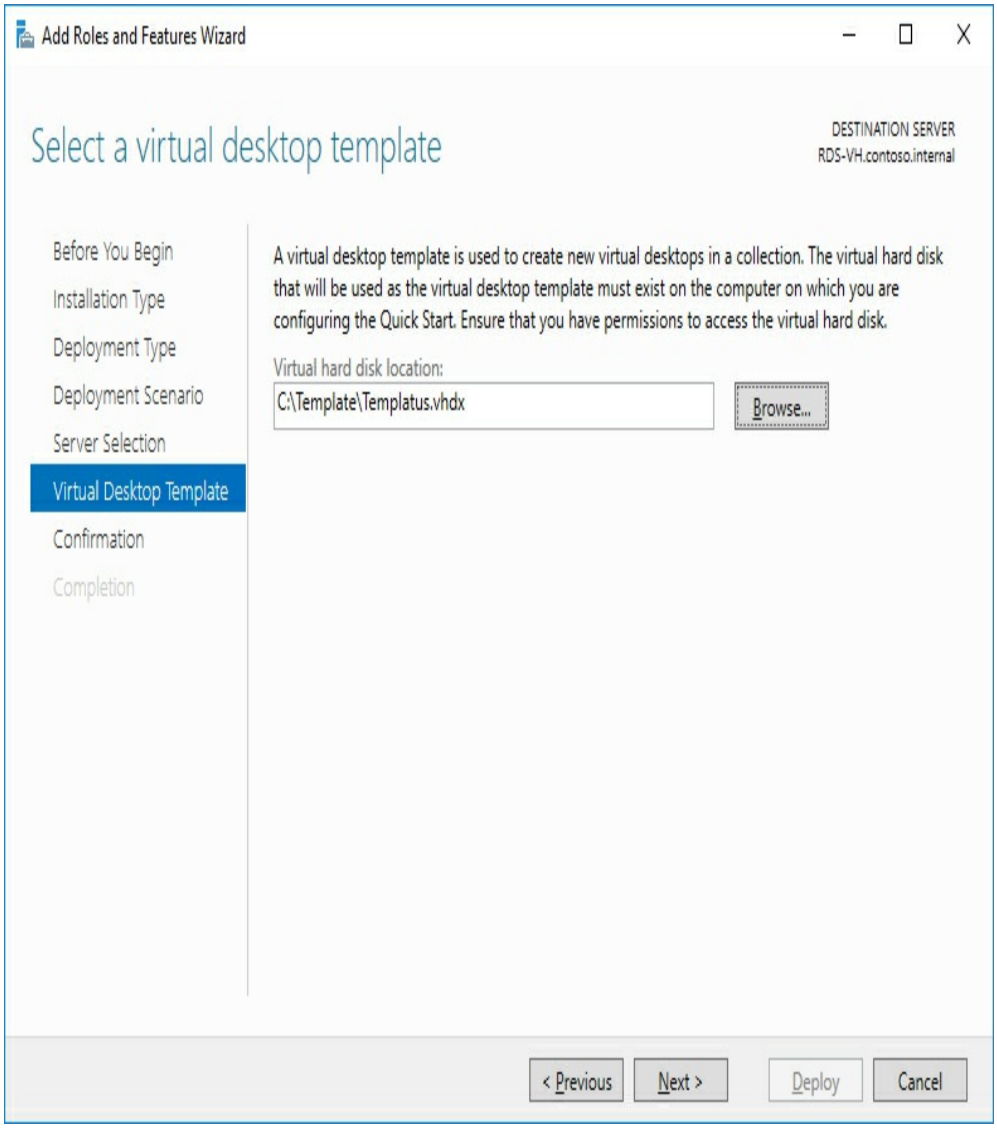


Figure 16-5 Specify the template for RD Virtualization Host server

Virtual machine preparation

To prepare a virtual machine for use as a template in an RD Virtualization Host deployment, you should perform the following tasks:

- Deploy the virtual machine and install all the applications that will be available to users of the RD Virtualization Host VMs. Perform this deployment on one of the Hyper-V servers that will function as an RD Virtualization Host. You can build this virtual machine on another computer and then import it on the RD Virtualization Host. Although not required, you should configure this virtual machine as a Generation 2 VM because this will give you the greatest flexibility in terms of virtual machine configuration.
- Once the applications are installed, sysprep the virtual machine, generalize and set it to boot using the Out Of Box Experience setting. Once sysprepped, the virtual machine will be shut down.

Virtual desktop collections

Virtual desktop collections function in a manner similar to a session collection. You configure who can connect and select the resources they have access to on a per-collection basis. When creating the virtual desktop collection, you need to provide the following information:

- **Collection Name.** Name of the virtual desktop collection.
- **Collection Type.** This can be a pooled virtual desktop collection or a personal virtual desktop collection.
- **Virtual Desktop Template.** The prepared template virtual machine.
- **Virtual Desktop Settings.** You can specify a Sysprep answer file or provide unattended installation settings. This information is used to join the virtual machine to the domain. This requires that the

computer account of the host computer has permission to perform the AD DS domain join.

- **Unattended Settings.** Provide the time zone, AD DS domain, and organizational unit used in the domain join operation.
- **Users and User Groups.** Shown in Figure 16-6, these settings allow you to specify which users and groups have access to the collection. They also allow you to configure the number of virtual desktops to be created in the collection, and the name prefix and suffix for the computer accounts of RDS VDI virtual machines.

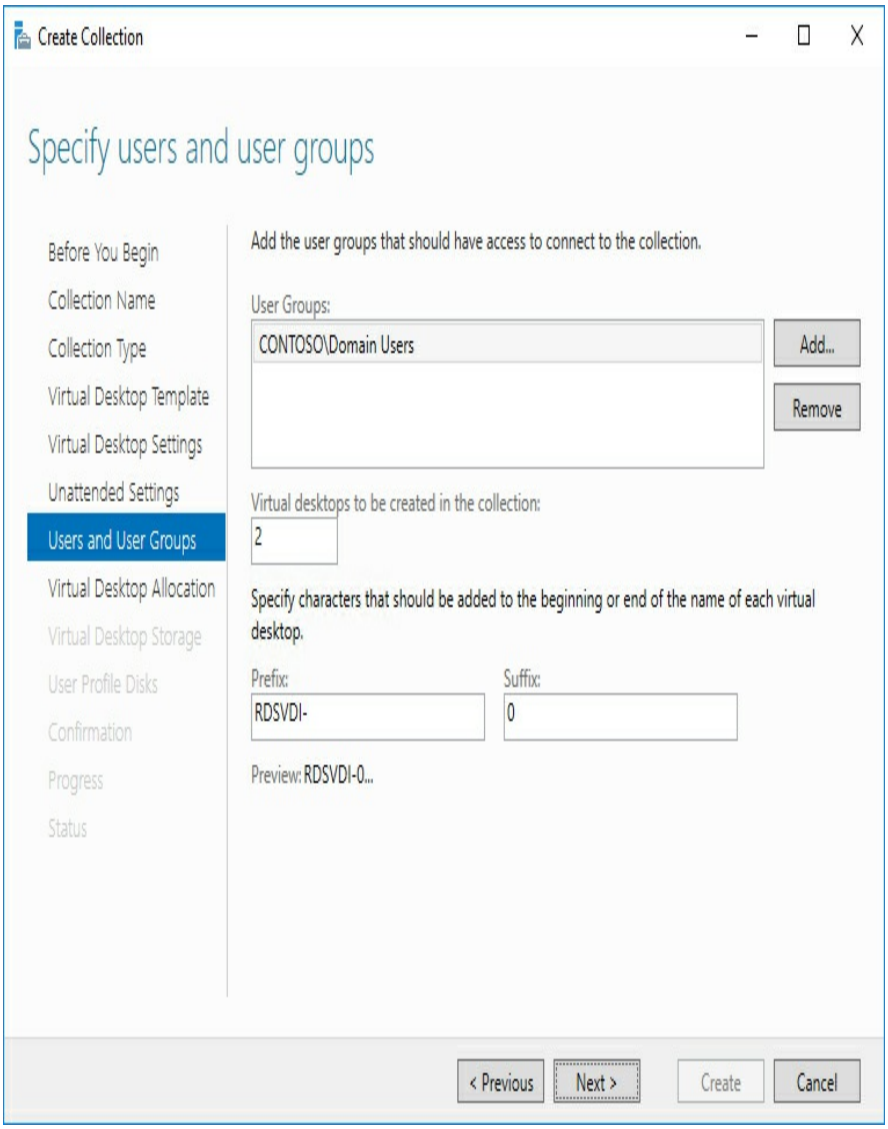


Figure 16-6 RDS VDI User and User Groups settings

- **Virtual Desktop Allocation.** Allows you to specify the allocation of VDI virtual machines in the collection to RD Virtualization Host servers. For example, you can have 10 virtual machines on one host server and 20 on another.
- **Virtual Desktop Storage.** Allows you to configure the type of storage used for the virtual desktops.
- **User Profile Disks.** Allows you to configure and enable user profile disk settings and the location.

Once the virtual desktop collection is created, you can also configure client settings for the collection. This determines whether functionality such as audio and video playback, audio recording, smart cards, plug and play devices, drives, the clipboard, and printers are redirected from the Remote Desktop client to the RD virtual machine.

Pooled virtual desktops

Pooled virtual desktops are collections of identically configured virtual machines that users connecting a Remote Desktop Virtualization Host through a Remote Desktop Connection Broker are assigned. A user will be connected to the first available pooled virtual desktop on first sign in, and will only be redirected to the same pooled virtual desktop in the event that their connection is disrupted and they haven't signed out. If they sign out and reconnect, they may be assigned to a new virtual desktop within the pool. Pooled virtual desktops do not

Technet24

retain session information. When a user signs out from a pooled virtual desktop, the virtual machine reverts to its original pre sign-in state.

Personal virtual desktops

You use personal virtual desktops when you want to assign users persistent VDI virtual machines. This means that users can make customizations and store data on the virtual machines that persist across sessions. The virtual machines are associated with the user, and they are automatically connected to that virtual machine when they initiate a session through the Remote Desktop Connection Broker. Personal virtual desktops retain session information across connections and are not reverted to their original pre sign-in state when the session ends.

DDA and RemoteFX

Discrete Device Assignment (DDA) and RemoteFX allow virtual machines in a VDI deployment to leverage the graphical processing power of compatible graphics cards. In turn, this allows workloads that need to leverage substantive graphical processing power, such as those used in Computer Aided Design (CAD). RemoteFX has been deprecated in Windows Server 2019 in favor of Discrete Device Assignment (DDA). Chapter 6, "Hyper-V," includes more information on DDA. In Windows

Server 2016 deployments, graphics cards used for DDA must be installed on the Remote Desktop Virtualization Host, and special graphics adapter drivers will need to be configured for each virtual machine that takes advantage of DDA or RemoteFX.

REMOTE DESKTOP WEB ACCESS

You can use Remote Desktop (RD) Web Access to allow clients to connect to an RD Session Host or Virtualization Host session through a webpage link rather than through Remote Desktop Connection client software or by running a RemoteApp application. To use RD Web Access, clients need to run a supported version of the Windows Client or Windows Server operating system. You can also use Remote Desktop Web Access to publish applications directly to the Start menu of computers running the Windows 8.1 and Windows 10 operating systems.

REMOTE DESKTOP LICENSING

Each client or user that accesses a Remote Desktop server, either through an RD Session Host session, a RemoteApp application, or an RD Web Access application, must have a license. This special type of license is not included with the licenses you get when you purchase Windows Server or Windows client operating systems such as Windows 10. This license is

called a Remote Desktop Services Client Access License (RDS CAL). You install RDS CALs on an RD license server. When a user, computer, or thin client initiates a session to a Remote Desktop Services server, the server performs a check to verify that a valid license has been issued or is available.

You configure licensing settings for a Remote Desktop server through Deployment Properties as shown in Figure 16-7.

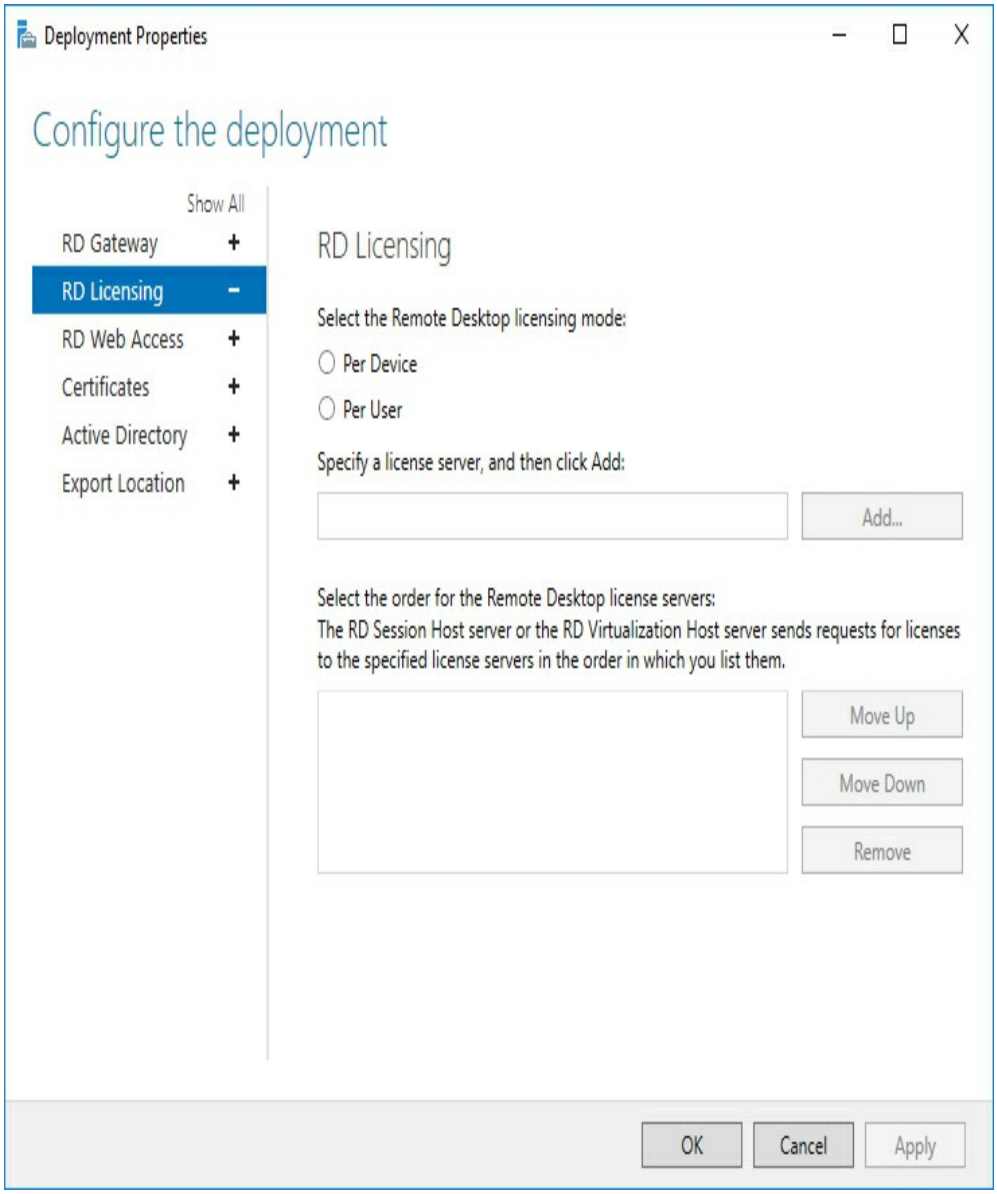


Figure 16-7 RD Licensing settings

Installing RDS CALs

You can install several types of licenses on an RD license server. The different license types are as follows:

- **RDS Per User CAL.** An RDS Per User CAL gives a specific user

account the ability to access any RD Session Host server in an organization from any computer or device. This type of CAL is a good idea when you have the same person accessing Remote Desktop from different devices.

- **RDS Per Device CAL.** The RDS Per Device CAL licenses a specific device, such as a normal computer or thin client, with the ability to connect to an RD Session Host server. RDS Per Device CALs are reset automatically by the RD license server after a random period between 52 and 89 days. This means that if a computer or thin client is decommissioned, the license assigned to it is reclaimed automatically and can be used again. If you are decommissioning a large number of computers, you can revoke 20% of the RDS Per Device CALs.
- **RDS External Connector.** Use this license type to allow people who are external to your organization to access a single Remote Desktop server. A separate RDS External Connector license must be purchased for each Remote Desktop server accessible to external users.
- **Service Providers License Agreement.** This is a special type of license agreement for service providers. It is useful when providing Remote Desktop Services to multiple organizations.

You can mix and match RDS Per User CALs and RDS Per Device CALS as necessary. When using Remote Desktop Virtualization Hosts, you also need to purchase a Virtual Enterprise Centralized Desktop licensing agreement for the VDI operating system. You can use the Manage RDS CALs Wizard in the RD Licensing Manager console to:

- Migrate RDS Cals to different RD Licensing Servers
- Rebuild the RD Licensing Server database

Inside OUT

Licensing is complicated

Unless you are certain, ensure that you check with a Microsoft licensing specialist about your licensing configuration.

Activating a License Server

You must activate an RD license server before it is able to issue CALs. The activation process is reasonably similar to the one that you need to go through when you perform product activation for a Windows operating system or for an application such as Microsoft Office. The method you use depends on whether the RD license server has a connection to the Internet. If the server does have a connection to the Internet, you can perform automatic activation. If the server does not have a direct connection to the Internet, you can perform an activation using a telephone, or you can navigate to a website on a computer or mobile device that does have a connection to the Internet and perform activation by filling out a form and entering the activation code on the RD license server.

MANAGING REMOTE DESKTOP SERVICES USING POWERSHELL

The Remote Desktop PowerShell module contains a large number of cmdlets that you can use to manage all aspects of Remote Desktop Services on Windows Server

2019. These cmdlets are described in Table 16-1.

Table 16-1 Remote Access PowerShell cmdlets

Noun	Verbs	Function
RDActiveManagementServer	Set	Configure active RD Connection Broker server
RDAvailableApp	Get	View applications that can be published from a collection
RDCertificate	Get, New, Set	Manage Remote Desktop Services certificate configuration
RDClientAccessName	Set	Manage DNS name that clients use to connect to RDS deployment
RDConnectionBrokerHighAvailability	Get, Set	Configure connection broker high availability
RDDatabaseConnectionString	Remove, Set	Configure RD database connection strings
RDDeploymentGatewayConfiguration	Get, Set	Manage RD deployment gateway configuration

RDFileTypeAssociation	Get, Set	Manage RD filetype associations
RDLicenseConfiguration	Get, Set	Configure RD licensing
RDPersonalSessionDesktopAssignment	Export, Get, Import, Remove, Set	Manage association between personal session and a user
RDPersonalVirtualDesktopAssignment	Export, Get, Import, Remove, Set	Manage association between virtual desktop and a user
RDPersonalVirtualDesktopPatchSchedule	Get, New, Remove, Set	Configure personal virtual desktop patch schedule
RDRemoteApp	Get, New, Remove, Set	Manage RD RemoteApp
RDRemoteDesktop	Get, Set	Publish a remote desktop to a collection

Technet24

RDServer	Add, Get, Remove	Manage server membership of a Remote Desktop deployment
RDSession Collection	Get, New, Remove	Manage RD session collections
RDSession Collection Configuration	Get, Set	Manage RD session collection configurations
RDSession Deployment	New	Create new RD session deployment
RDSession Host	Add, Get, Remove, Set	Manage RD Session Host session collection server membership
RDUser	Disconnect	Disconnect a connected RD user
RDUserLogoff	Invoke	Sign out a connected RD user
RDUserSession	Get	View RD user session information
RDUserMessage	Send	Send a message to an RD user

RDVirtualDesktop	Get, Move	Manage RD virtual desktops
RDVirtualDesktopADMachineAccountReuse	Disable, Enable, Test	Manage whether RD Connection Broker server reuses existing AD computer accounts for pooled virtual desktops created using a template in a managed collection
RDVirtualDesktopCollection	Add, Get, New, Remove, Set, Update	Manage virtual desktop membership of virtual desktop collections
RDVirtualDesktopFromCollection	Remove	Remove a virtual desktop from a collection
RDVirtualDesktopCollectionConfiguration	Get	View virtual desktop collection configuration
RDVirtualDesktopCollectionJob	Stop	End a virtual desktop collection job
RDVirtualDesktopCo	Get	View virtual desktop connection job status

Technet24

llectionJob Status		
RDVirtual DesktopCo ncurrency	Get, Set	Manage virtual desktop concurrency
RDVirtual DesktopDe ployment	New	Create a new virtual desktop deployment
RDVirtual DesktopIdl eCount	Get, Set	Manage virtual desktop idle settings
RDVirtual DesktopTe mplateExp ortPath	Get, Set	Configure virtual desktop template paths
RDWorksp ace	Get, Set	Manage workspace name settings for Remote Desktop deployment
RDOUAcc ess	Grant, Test	Configures OU access for the RD Connection Broker server

Chapter 17

Azure IaaS and hybrid services

Windows Server IaaS VMs

Azure Active Directory

Azure hybrid cloud services

In the first part of this chapter, you'll learn about deploying and managing Azure Infrastructure as a Service (IaaS) VMs running Windows Server. You'll then learn about Azure Active Directory, how you can synchronize your on-premises identities, and how you can domain join Azure IaaS VMs running Windows Server to this cloud-based identity service. In the final part of the chapter, you'll learn about the hybrid technologies available in Azure that allow you to extend your on-premises environment into the cloud.

WINDOWS SERVER IAAS VMS

Technet24

IaaS allows you to host your IT infrastructure in the cloud. Rather than deploy the virtual machine in an on-premises hypervisor, which you must pay for and maintain yourself, IaaS allows you to run your virtual machines in a cloud-provider's infrastructure, which the cloud provider is responsible for maintaining.

As shown here, there are several benefits to running Windows Server 2019 in an Azure virtual machine over running a Windows Server 2019 VM on-premises:

- **Efficient costing.** You will be charged on a per-minute (and in some cases, on a per-second) basis for the hosted virtual machine. You aren't charged for deallocated virtual machines. In certain circumstances, it might be cheaper to run the virtual machine in the cloud than paying for the infrastructure to host the virtual machine in your own datacenter.
- **Efficient use of licensing.** When you deploy Windows Server 2019 on IaaS, the operating system license is included in your charges. If you have an existing software assurance license, you can use Bring Your Own License (BYOL) functionality.
- **Up-to-date deployment images.** When deployed, the operating system will be fully patched with the latest software updates. This means that the operating system will be secure at deployment. When you maintain your own images, you need to either make sure that they are up to date or spend time after deployment ensuring that they are up to date. Gallery images also include the Azure virtual machine agent. The Azure virtual machine agent enables the VM extensions. The VM extensions enable post VM deployment configuration and VM recovery features, and they can reset the default VM admin account password.

The default gallery options are the datacenter editions of

Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019. The expanded images available in the gallery are shown in Figure 17-1.

Technet24

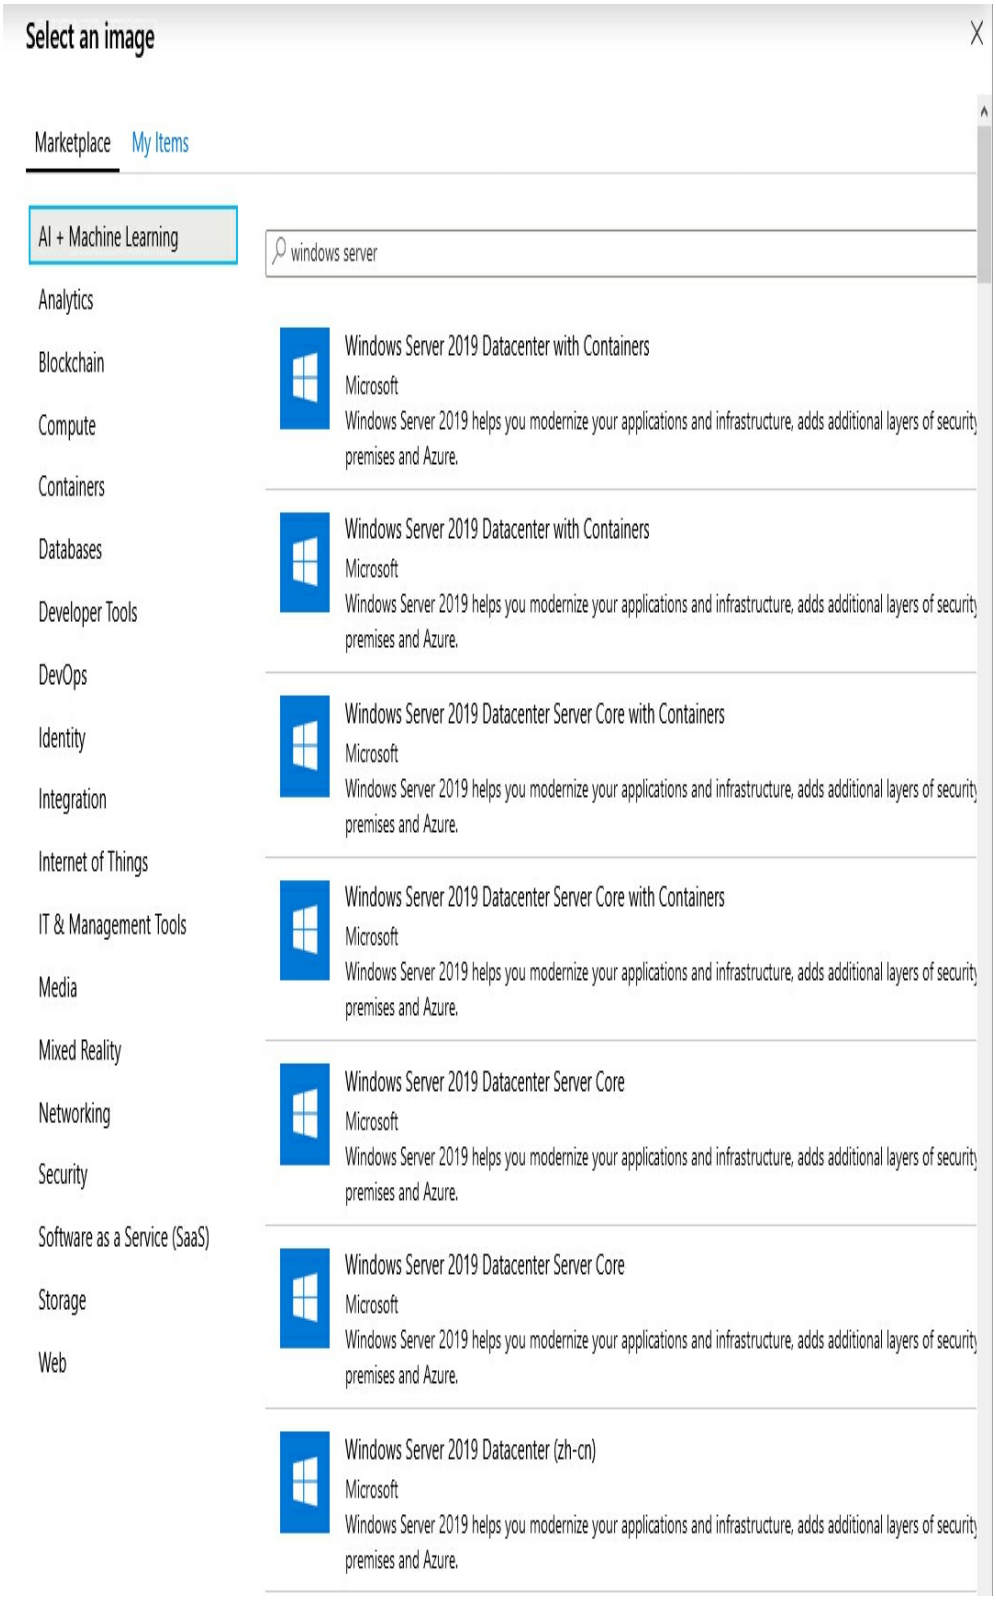

Select an image
Marketplace
My Items
AI + Machine Learning
Analytics
Blockchain
Compute
Containers
Databases
Developer Tools
DevOps
Identity
Integration
Internet of Things
IT & Management Tools
Media
Mixed Reality
Networking
Security
Software as a Service (SaaS)
Storage
Web
windows server
Windows Server 2019 Datacenter with Containers
Microsoft
Windows Server 2019 helps you modernize your applications and infrastructure, adds additional layers of security
premises and Azure.
Windows Server 2019 Datacenter with Containers
Microsoft
Windows Server 2019 helps you modernize your applications and infrastructure, adds additional layers of security
premises and Azure.
Windows Server 2019 Datacenter Server Core with Containers
Microsoft
Windows Server 2019 helps you modernize your applications and infrastructure, adds additional layers of security
premises and Azure.
Windows Server 2019 Datacenter Server Core with Containers
Microsoft
Windows Server 2019 helps you modernize your applications and infrastructure, adds additional layers of security
premises and Azure.
Windows Server 2019 Datacenter Server Core
Microsoft
Windows Server 2019 helps you modernize your applications and infrastructure, adds additional layers of security
premises and Azure.
Windows Server 2019 Datacenter Server Core
Microsoft
Windows Server 2019 helps you modernize your applications and infrastructure, adds additional layers of security
premises and Azure.
Windows Server 2019 Datacenter (zh-cn)
Microsoft
Windows Server 2019 helps you modernize your applications and infrastructure, adds additional layers of security
premises and Azure.

Figure 17-1 Windows Server Azure IaaS VM images

Creating Azure IaaS VMs

To create an Azure IaaS VM using the Azure portal, perform the following steps:

1. On the Azure portal home page, click Create A Resource, click Compute, and then click Virtual Machine.
2. On the Create A Virtual Machine page, configure the following:
    1. **Subscription.** Specify which subscription the VM will belong to.
    2. **Resource Group:** Specify the resource group in which to place the VM. You can create a new resource group or use an existing resource group. Resource groups should host resources that share a common lifecycle, such as VMs that function as tiers of a multitier application.
    3. **Virtual Machine Name.** This is a name that identifies the virtual machine.
    4. **Region:** This is where you choose the Azure region that will host the virtual machine. Choose a region that is proximate with where you want the application to run.
    5. **Availability Options.** This is where you choose whether you want to use replicated VMs in Availability Zones or Availability Sets. To get the strongest Azure SLA, you will need to choose an availability option.
    6. **Image.** This is where you choose the image you want to deploy. Figure 17-2 shows the selection of Windows Server 2019 Datacenter Edition.

Technet24

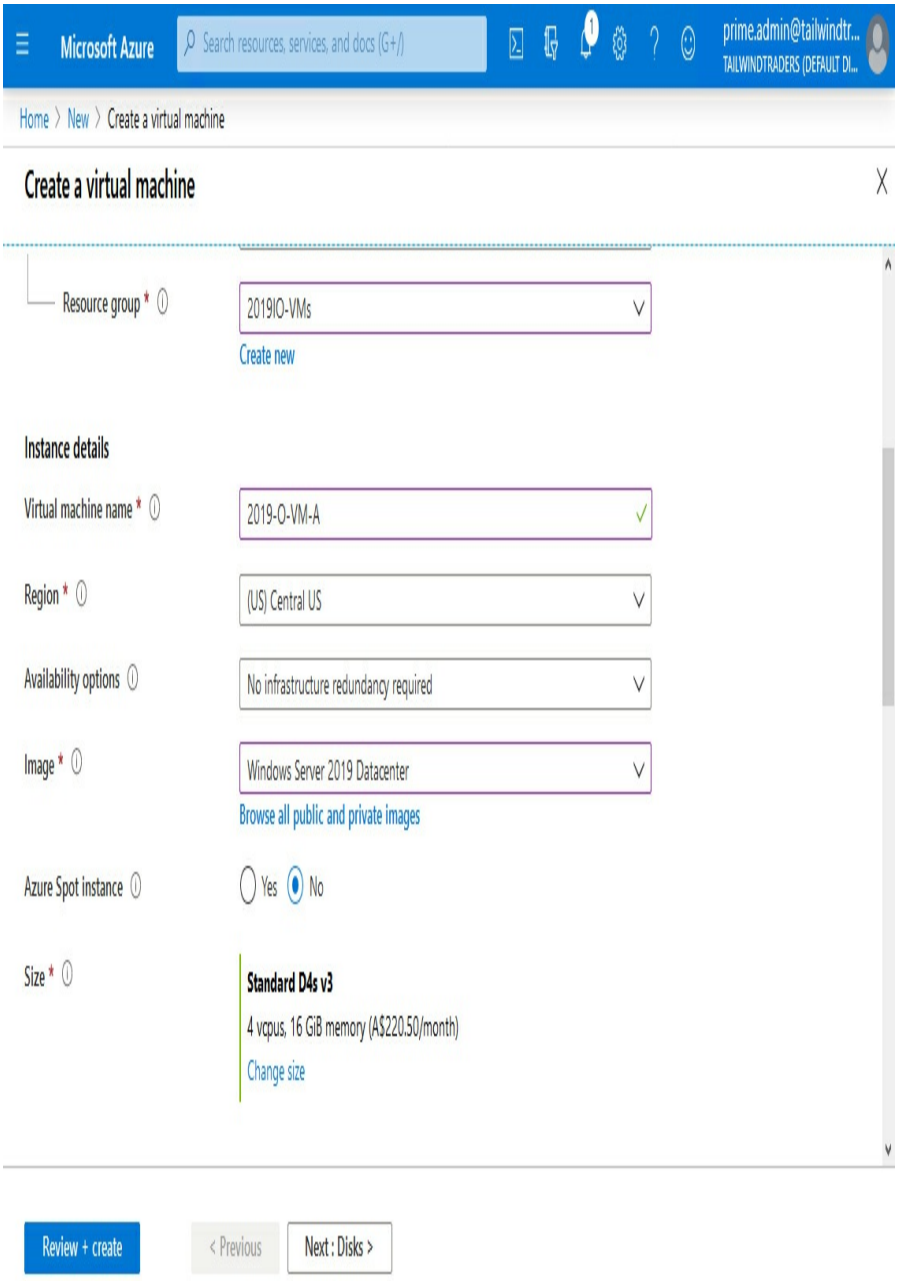


Figure 17-2 Creating a virtual machine

7. **Azure Spot Instance.** Choose this option for a discounted rate, though your VM may be shut down if Azure needs to recall capacity for traditional pay-as-you-go workloads.
8. **Size.** This is where you specify virtual CPU, memory, data disks,

maximum IOPS, and whether the VM supports premium disks. A commonly chosen size at the time this book is published is a D4s_v3, which provides four virtual CPUs, 16 GB RAM, supports eight data disks, and supports the use of premium disks.

9. **Administrator Account Username.** This cannot be set to *Administrator*, *Admin*, or *Root*. Make this name something consistent that you will remember, such as *Prime* or *Chancellor*.
10. **Password.** Create a complex password that is a minimum of 12 and a maximum of 123 characters in length.
11. **Public Inbound Ports.** Choose whether you want to have the VM accessible on specific ports to hosts on the Internet. There is no need to open these ports unless the workload is serving content to hosts on the Internet, such as having it host a public web server. You can also configure which ports are open after deployment.
12. **Specify Whether You Have A Windows Server License.** If your organization has Software Assurance, you can reduce your licensing costs by specifying that you have an existing Windows Server license.

From this point, you can click Review And Create to create the VM, or you can keep navigating through the deployment wizard and specify the properties of the operating system disk and extra disks. Also, you can specify which virtual network the VM should connect to, whether it should have a public IP address, and whether a basic network security group should be created for the VM, as shown in Figure 17-3. You should only allocate a VM a public IP address if the VM needs to serve content to hosts on the public Internet. Technologies such as Azure Bastion or VPNs, covered later in this chapter, allow you to make administrative connections to VMs

Technet24

that only have private IP addresses.

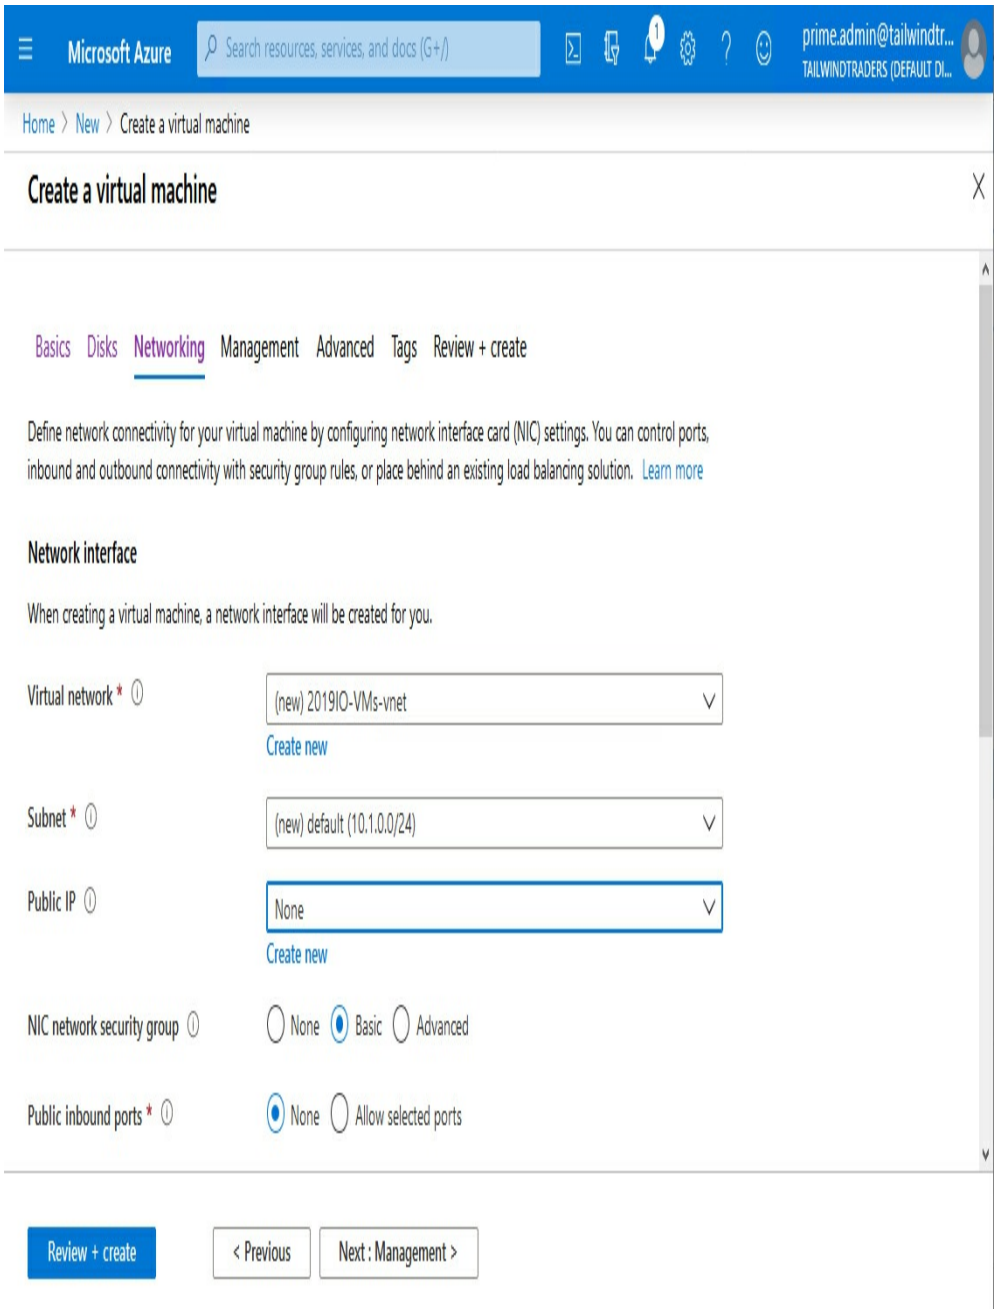


Figure 17-3 IaaS VM network configuration

On the Management page, shown in Figure 17-4, you

can choose whether you want to have the VM monitored by Azure Security center. Also, you must specify whether you want to enable boot diagnostics and/or operating system guest diagnostics and specify a storage account to store diagnostic data. Boot diagnostics allow you to make an EMS connection and view screenshots of the VM running on the Azure host. Guest OS diagnostics allow you to get up-to-date performance counter information from the VM displayed in the Azure portal.

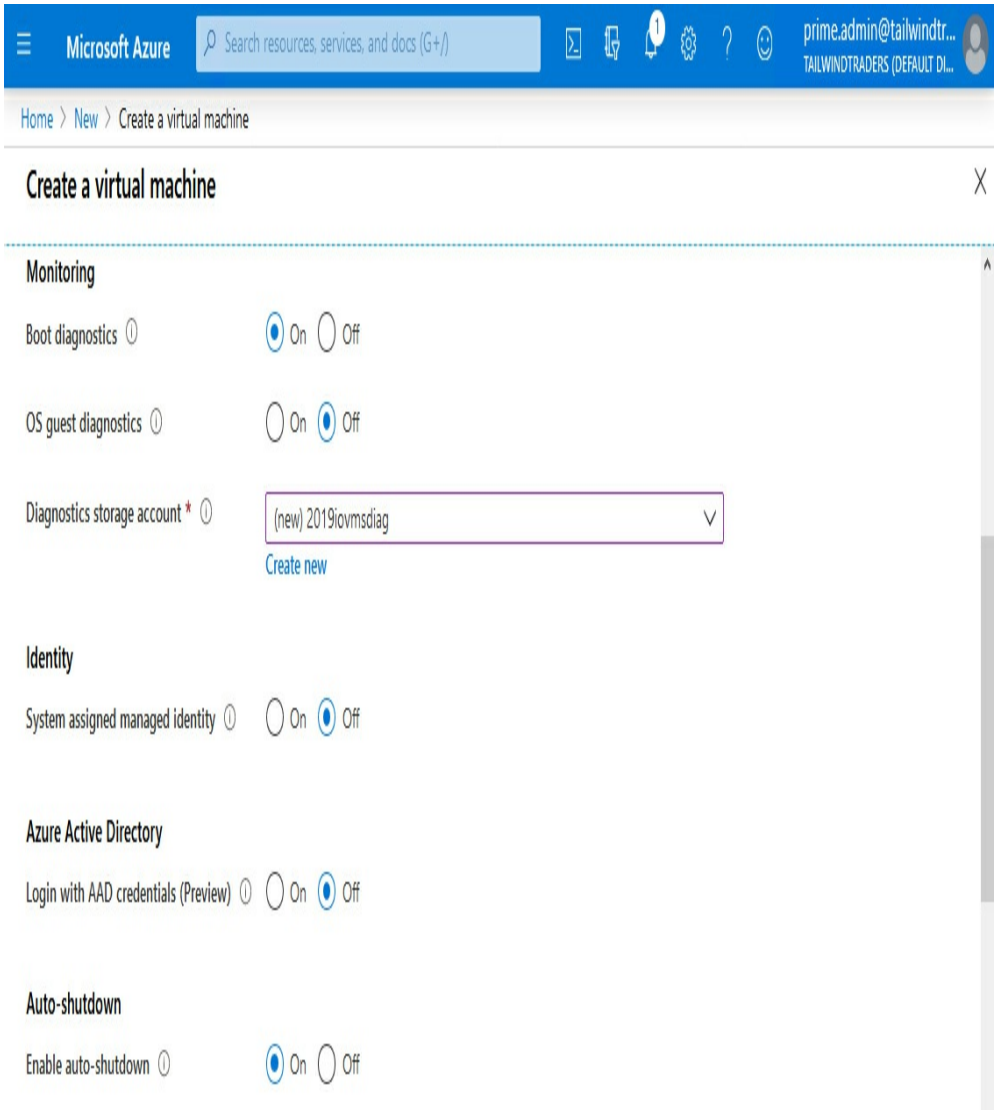


Figure 17-4 IaaS VM diagnostics configuration

You can also configure a system-assigned managed identity, which is similar to an Active Directory Computer account and allows the VM to access cloud resources without having to store credentials. You can configure whether you want to be able to sign in to the VM using Azure Active Directory credentials and

whether the VM automatically shuts down at a specific time each day if it has been left running. You can also enable a backup for the VM during creation. You'll learn about configuring Azure Backup for a VM later in the chapter.

On the Advanced page of the Create A Virtual Machine deployment wizard, you can specify whether you want to run the VM on a private Azure host, whether you want the VM to be in proximity with other resources deployed in the same region, and whether you want the VM to be a Generation 1 or a Generation 2 VM. The choice of generations provides similar benefits to choosing between Generation 1 or Generation 2 VMs when deploying in Hyper-V.

Run the following set of commands to deploy a Windows Server 2019 VM named *2019IO-C* using Cloud Shell to the resource group *2019IO-RG*, where the username is set to *prime* and the administrator password is set to *&&Tailwind&&Tailwind*, and to connect the VM to the subnet *Alpha* on virtual network *VNet1* :

Click here to view code image

```
adminPassword=&&Tailwind&&Tailwind

az vm create --resource-group 2019IO-RG --name 20190IO-C --image
win2019datacenter
--vnet-name VNet1 --subnet Alpha --admin-username prime --admin-
password $adminPassword
```

Later in the chapter, you'll learn how to create an Azure IaaS VM running Windows Server using Windows Admin Center.

> **More Info**
>
> **Windows VMs in Azure**
>
> You can learn more about Windows VMs in Azure at https://docs.microsoft.com/en-us/azure/virtual-machines/windows/.

IaaS VM networking

IaaS VMs connect to virtual networks. An IaaS virtual network is a collection of subnets that share a common private IP address space in the RFC 1918 range. For example, you might create a virtual network that uses the *192.168.0.0/16* address space and create subnets, such as *192.168.10.0/24*. Azure IaaS virtual machines connect to the subnets in an Azure virtual network.

IaaS virtual networks

Azure IaaS VMs can only use virtual networks that are in the same location as the IaaS virtual machine. For example, if you are deploying an IaaS VM to Australia South East, you'll only be able to connect the IaaS VM directly to a virtual network in Australia South East.

To create a new network, perform the following steps:

1. In the Azure console, click Create A Resource, and in the Search bar, type **Network**. In the results, click Virtual Network, and then click Create.

2. On the Create Virtual Network blade shown in Figure 17-5, provide the following information:

    1. **Name.** This is the name of the virtual network and must be unique to the subscription.
    2. **Address space.** This is the IPv4 address space that will be used by hosts on the virtual network. You have the option to add an IPv6 address space.
    3. **Subscription.** Choose which Azure subscription the virtual network will be associated with.
    4. **Resource Group.** Choose an existing resource group within the subscription or create a new resource group.
    5. **Subnet Name.** This is the first subnet within the address space. Because you might want to add a gateway subnet for a VPN at some point in the future, it is a good idea to avoid creating a subnet that starts right at the beginning of the address space. For example, if your address space is *192.168.0.0/16*, create the first subnet as *192.168.1.0/24* or *192.168.10.0/24* rather than *192.168.0.0/24*. This ensures that you always have extra space for special subnet types, such as gateway subnets or Azure bastion subnets.
    6. **Subnet Address Range.** This is the address range of the first subnet within the address space. This range can consume the entire address space.
    7. **DDoS Protection.** Select between the basic Distributed Denial Of Service Protection or enable the Standard Distributed Denial Of Service Protection. The basic Distributed Denial Of Service Protection level incurs no extra fee, while Standard Distributed Denial Of Service Protection level incurs extra fees.
    8. **Service Endpoints.** This setting allows you to configure access to Azure IaaS service endpoints, such as Azure Key Vault from the virtual network. You'll learn about using Azure Key Vault to enable BitLocker later in this chapter.

Technet24

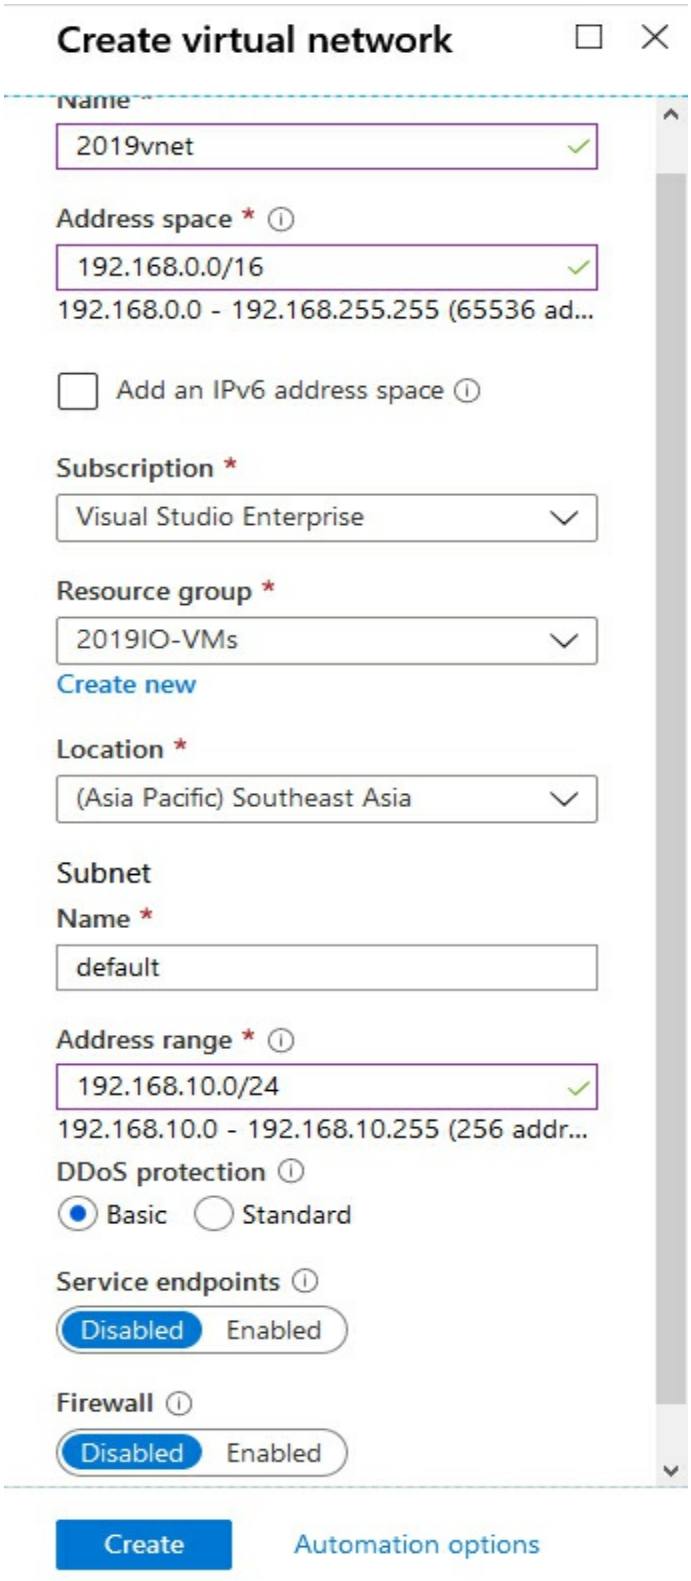

Create virtual network
Name *
2019vnet
Address space * ⓘ
192.168.0.0/16
192.168.0.0 - 192.168.255.255 (65536 ad...
Add an IPv6 address space ⓘ
Subscription *
Visual Studio Enterprise
Resource group *
2019IO-VMs
Create new
Location *
(Asia Pacific) Southeast Asia
Subnet
Name *
default
Address range * ⓘ
192.168.10.0/24
192.168.10.0 - 192.168.10.255 (256 addr...
DDoS protection ⓘ
Basic
Standard
Service endpoints ⓘ
Disabled
Enabled
Firewall ⓘ
Disabled
Enabled
Create
Automation options

Figure 17-5 Virtual network creation

Once you have deployed the virtual network, you can configure the following properties for the virtual network using the Azure console:

- **Address Space.** This property allows you to add additional address spaces to the Azure virtual network. You partition these address spaces using subnets.
- **Connected Devices.** This property lists the current devices that are connected to the Azure virtual network. This property includes a list of VM network adapters and the internal IP address to which those adapters are assigned.
- **Subnets.** This property shows subnets that you create within the address space and allows you to put different Azure virtual machines on separate subnets within the same virtual network.
- **DNS Servers.** This property allows you to configure the IP addresses of the DNS servers assigned by the DHCP server used by the Azure virtual network. Use this to configure DNS settings when you deploy a domain controller on an Azure IaaS VM or when you deploy Azure AD DS for an Azure virtual network.
- **Properties.** This property allows you to change which subscription the Azure virtual network is associated with.
- **Locks.** This setting allows you to apply configuration locks, which blocks changes being made to the settings of the resource unless the lock is first removed.
- **Automation Script.** This setting allows you to access the JSON template file that you can use to reproduce and redeploy virtual networks with similar settings.

By default, hosts located on one subnet in a virtual network will automatically be able to communicate with

Technet24

other subnets on a virtual network. You can modify this behavior by configuring user-defined routes, network security groups, Azure Firewall, or Network Virtual Appliance; doing so allows you to configure the subnets within a virtual network in a similar manner to the way that you might segment traffic on an on-premises network.

More Info

Azure Virtual Networks

You can learn more about Azure Virtual Networks at https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-overview.

IP addressing

A virtual machine on an Azure network will have an internally assigned IP address in the range specified by the virtual network it is associated with. You can configure this assignment to be static or dynamic, as shown in Figure 17-6. It is important to remember that you assign an IP address as dynamic or static on the network adapter object within Azure; you don't manage IP address configuration from within the IaaS virtual machine operating system.

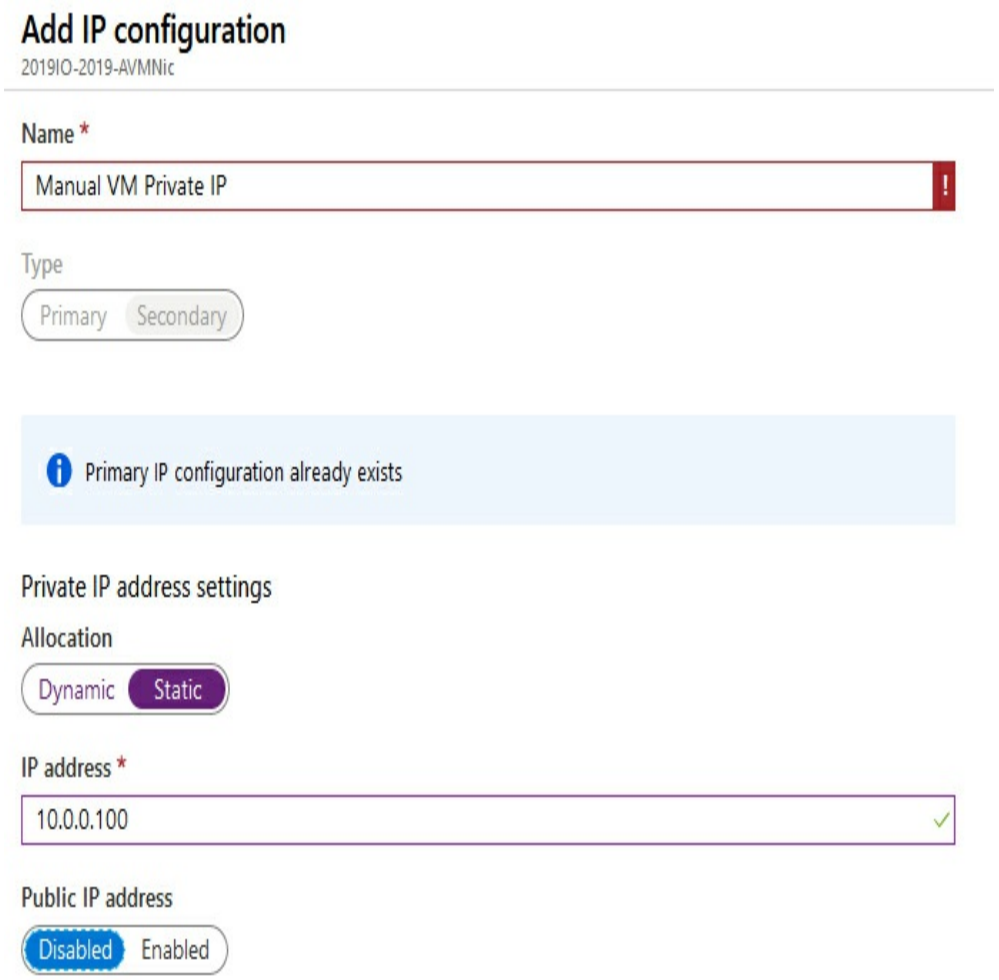


Figure 17-6 Add IP configuration

Inside OUT

Setting Azure IaaS VM addresses

You never configure the IP address for an Azure IaaS VM through virtual network adapter settings inside the VM. Instead, you configure the IP address for the IaaS VM on the network adapter properties in either the Azure portal or other management tools, such as Cloud Shell. An IaaS VM should always be configured to use dynamic IP addressing. If you mess this up and manually assign an IP address within the VM, you'll almost always (depending on what setting you configured) be unable to communicate with the VM. You can fix this by removing the existing virtual network adapter in Azure and replacing it with a new one. The new one will be configured for DHCP, and your networking problem will be fixed.

You may also assign a public IP address to the VM, but you should only do this in the event that the VM needs to be directly accessible to hosts on the Internet. A VM with an internally assigned IP address can perform outbound communication to hosts on the Internet without a public IP address, so a public address is only necessary if hosts on the Internet need to directly communicate with the IaaS VM.

Even VMs that do need to be accessible to hosts on the Internet can avoid having a public IP address if they are sitting behind a load balancer, web application, or network virtual appliance. A network virtual appliance (NVA) is similar to a traditional perimeter firewall or application gateway device and mediates traffic flow through to a web application or VM running in the cloud.

You can determine which IP addresses are assigned to an Azure virtual machine on the Network Interfaces blade of the VM's properties, as shown in Figure 17-7.

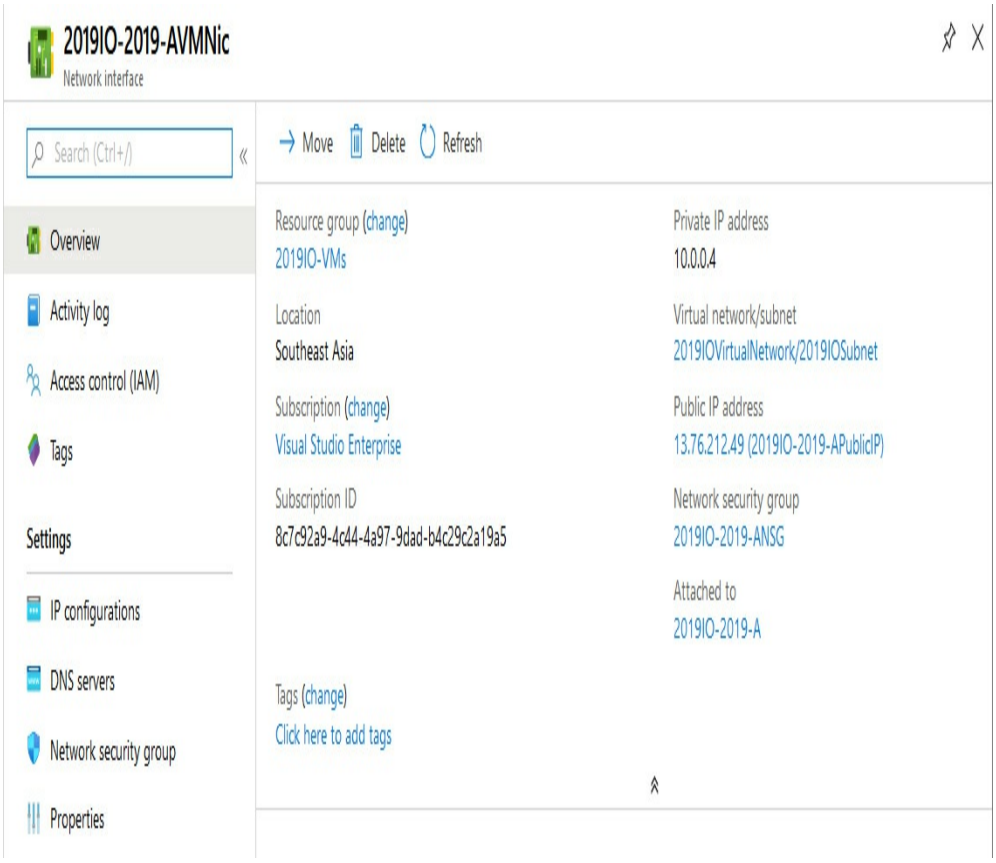


Figure 17-7 IaaS VM network interface properties

You can also apply a region-specific DNS name to the network interface so that you can connect to the VM using a FQDN rather than IP address, as shown in Figure 17-8. If you have a DNS server or you have configured an Azure DNS zone, you can then create a CNAME record that points to the region-specific DNS name for future connections. This saves you from always connecting to a *cloudapp.azure.com* address, and it will also remain pointing to the VM if the VM changes its IP address.

Technet24

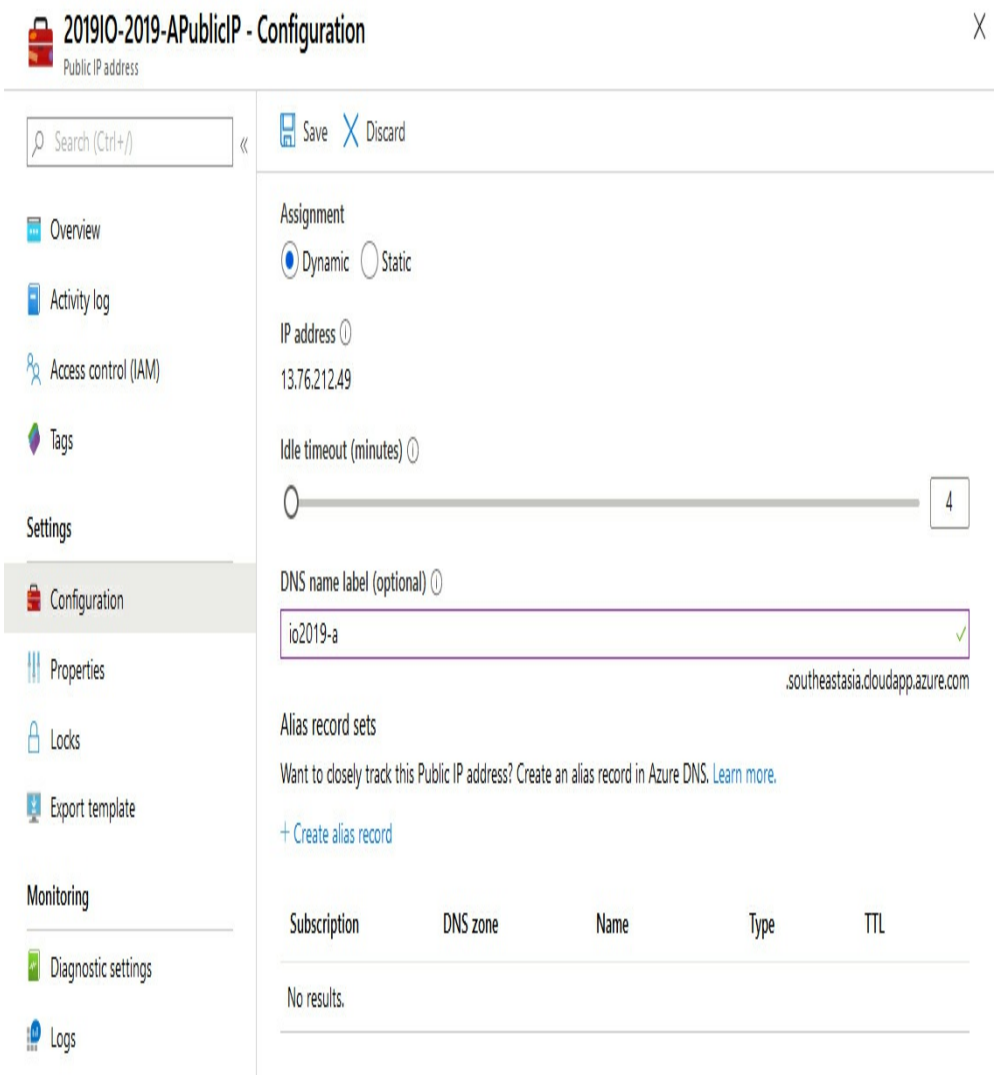


Figure 17-8 DNS name label

More Info

IP addressing

You can learn more about IP addressing in Azure at https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-ip-addresses-overview-arm.

Network Security Groups

A Network Security Group (NSG) is a packet filter for mediating traffic at the virtual network subnet level and also at the level of an IaaS VM's network adapter, as shown in Figure 17-9. When you create a virtual machine, an NSG is automatically applied to the VM's network adapter interface, and you can choose whether you want to allow traffic on management ports, such as TCP port 3389 and through to the VM.

Technet24

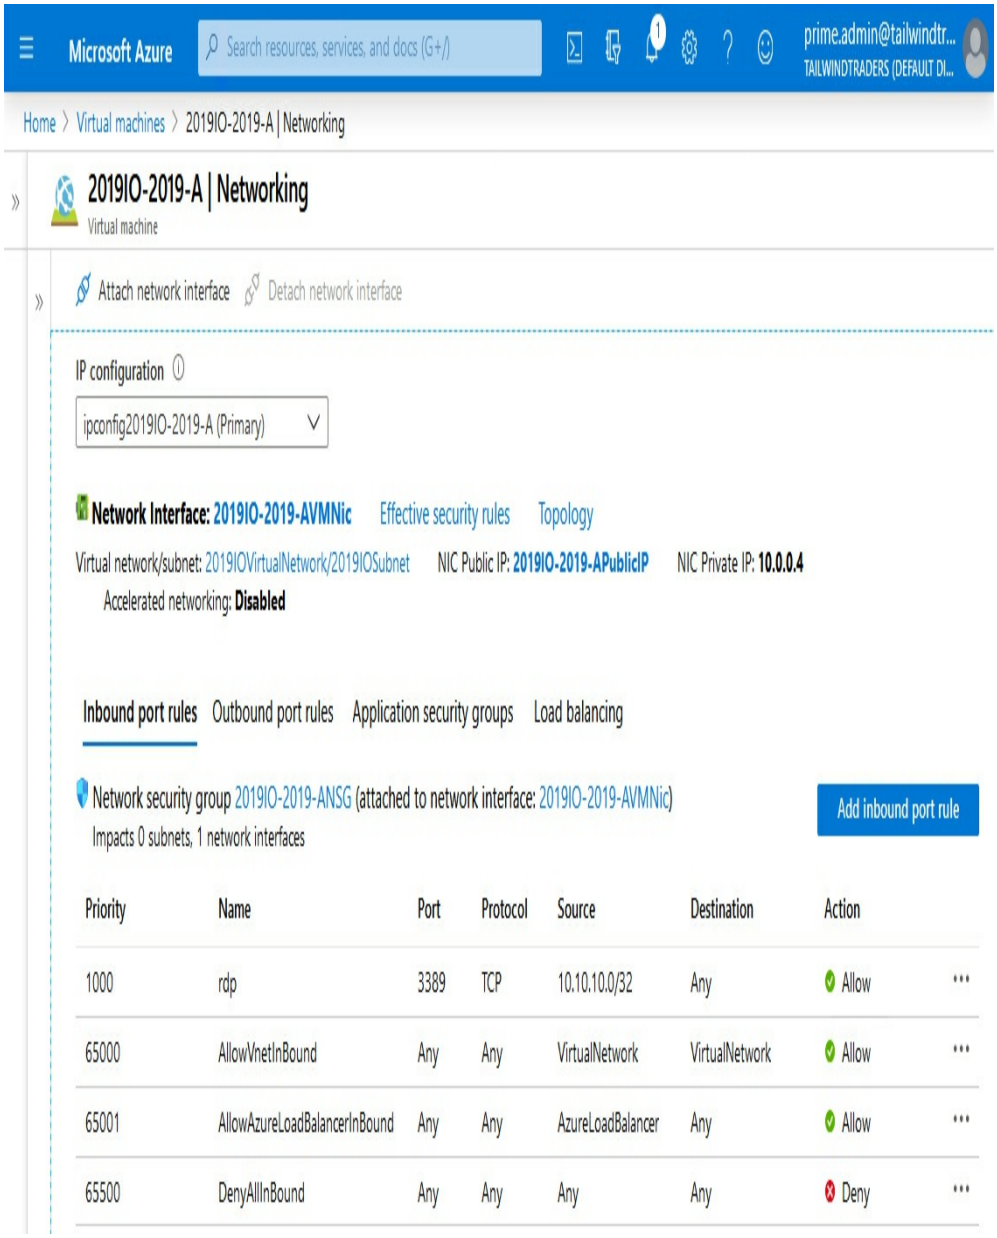


Figure 17-9 Network Security Group

An NSG rule has the following elements:

- **Priority.** NSG rules are processed in order, with lower numbers processed before higher ones. The moment traffic meets the conditions of a rule, that rule is enforced, and no further rules are

processed.

- **Source.** This specifies the source address or subnet of traffic. It can also include a service tag, which allows you to specify a particular Azure service or an application security group (a way of identifying the network identity of a series of workloads that make up an application).
- **Destination.** This specifies the destination address or subnet of traffic. It can also include a service tag, which allows you to specify a particular Azure service or an application security group (a way of identifying the network identity of a series of workloads that make up an application).
- **Protocol.** This can be configured to TCP, UDP, ICMP, or Any.
- **Port Range.** This allows you to specify either an individual port or a range of ports.
- **Direction.** This specifies whether the rule applies to inbound or outbound traffic.
- **Action.** This allows you to specify whether you want to allow or deny the traffic.

Network security groups are fairly basic in that they only work on the basis of IP address information and cannot be configured on the basis of a fully qualified domain name. Azure offers more advanced ways of mediating traffic flow, including Azure Firewall and Network Virtual Appliances, which can be used in conjunction with NSGs; however, these topics are beyond the scope of this text.

More Info

NSGs

You can learn more about NSGs at https://docs.microsoft.com/en-us/azure/virtual-

network/security-overview.

VPNs and IaaS virtual networks

You can configure IaaS virtual networks to support VPN connections by configuring a VPN gateway. IaaS virtual network VPN gateways support site-to-site connections. A site-to-site connection allows you to connect an IaaS virtual network to an existing network, just as you might connect a branch office network to a head office network in your on-premises environment. IaaS virtual network VPN gateways also support point-to-site VPN connections. This allows you to connect individual host computers to IaaS virtual networks. Windows Server 2019 even includes simplified set up of a point-to-site VPN connection through deployment of Azure Network Adapter, which is covered later in this chapter.

More Info

Azure VPN gateways

You can learn more about Azure VPN gateways at https://docs.microsoft.com/en-us/azure/vpn-gateway/.

IaaS VM administration

Once you've deployed an IaaS VM, there are a variety of tasks that you need to perform, from configuring how users will connect to the virtual machine to ensuring the VM is regularly backed up.

In your on-premises environment, you would not allow most users of a VM to have access to a VM in Hyper-V manager and would only allow them to connect directly to the VM using RDP or PowerShell. You should do the same thing in Azure because most people who use VMs in Azure don't need to interact with them in the Azure console.

Azure IaaS VMs are only visible to users in the Azure portal if they have a role that grants them that right. The default Azure IaaS VM Role-Based Access Control (RBAC) roles are:

- **Virtual Machine Contributor.** Users who hold this role can manage virtual machines through the Azure console and perform operations such as restarting and deleting the virtual machine. Membership in this role does not provide the user with access to the virtual machine itself. It also does not provide access to the virtual network or storage account to which the virtual machine is connected.
- **Virtual Machine Administrator Login.** If the VM is configured to allow log in using Azure AD accounts, assigning this role grants the user local administrator privileges in the virtual machine. Users who hold this role can view the details of a virtual machine in the portal but not change the properties of the virtual machine.
- **Virtual Machine User Login.** Users who hold this role can view the details of a virtual machine in the Azure portal and can log in using their Azure AD account with user permissions. Users who hold this role cannot change the properties of the virtual machine.

Users who aren't members of the Virtual Machine Administrator Login or Virtual Machine User Login

roles can still access virtual machines if they have a local account on the virtual machine. Put another way, just because you can't see it in the portal doesn't mean you can't RDP to it if you have the address and the correct ports are open.

Remote PowerShell

You can initiate a remote PowerShell session from hosts on the Internet. Another option is to run a Cloud Shell session in a browser and perform PowerShell remote administration in this manner. Cloud Shell is a browser-based CLI, and it is a lot simpler to use than adding the Azure CLI to your local computer. There is a Cloud Shell icon on the top panel of the Azure Console.

You can enable Remote PowerShell on an Azure IaaS Windows VM by performing the following steps from Cloud Shell:

1. Ensure that Cloud Shell has PowerShell enabled by running the *pwsh* command.
2. From the PowerShell prompt in Cloud Shell, type the following command to enter local Administrator credentials for the Azure IaaS Windows VM:

   Click here to view code image

   ```
   $cred=get-credential
   ```

3. From the PowerShell prompt in Cloud Shell, type the following command to enable PowerShell remoting on the Azure IaaS Windows VM, where the VM name is *2019-IO-A* and the resource group that

hosts the VM is *2019-IO-RG*:

Click here to view code image

```
Enable-AzVMPSRemoting -Name 2019-IO-A -
ResourceGroupName 2019-IO-RG -Protocol
https -OSType Windows
```

4. Once this command has completed executing, you can use the *Enter-AzVM* cmdlet to establish a remote PowerShell session. For example, run this command to connect to the VM named *2019-IO-A* in resource group *2019-IO-RG*:

   Click here to view code image

```
Enter-AzVM -name 2019-IO-A -ResourceGroupName
2019-IO-RG -Credential $cred
```

More Info

PowerShell remoting from Cloud Shell

You can learn more about PowerShell remoting from Cloud Shell at https://techcommunity.microsoft.com/t5/itops-talk-blog/powershell-basics-connecting-to-vms-with-azure-psremoting/ba-p/428403.

Azure Bastion

Azure Bastion allows you to establish an RDP session to a Windows Server IaaS VM through a standards-compliant browser such as Microsoft Edge, Mozilla Firefox, or Google Chrome rather than having to use a remote desktop client. You can think of Azure Bastion as "Jump box as a Service" because it allows access to IaaS VMs that do not have a public IP address. Prior to the

release of Azure Bastion, the only way to gain access to an IaaS VM that didn't have a public IP address was either through a VPN to the virtual network that hosted the VM or by deploying a jump box VM with a public IP address from which you then create a secondary connection through to the target VM. If you have configured an SSH server on the IaaS VM, Bastion also supports creating SSH connections to Linux IaaS VMs or Windows Server.

Prior to deploying Azure Bastion, you need to create a special subnet named *AzureBastionSubnet* on the virtual network that hosts your IaaS VMs. Once you deploy Azure Bastion, the service will manage the network security group configuration to allow you to successfully make connections.

More Info

Azure Bastion

You can learn more about Azure Bastion at https://docs.microsoft.com/en-us/azure/bastion/bastion-overview.

Just In Time VM Access

Rather than have management ports, such as the port used for Remote Desktop Protocol, TCP port 3389, open to hosts on the Internet all the time, Just In Time (JIT) VM Access allows you to open a specific management port for a limited duration of time and only open that

port to a small range of IP addresses. You only need to use JIT if you require management port access to an Azure IaaS VM from a host on the Internet. If the VM only has a private IP address or you can get by using a browser-based RDP session, then Azure Bastion is likely a better option.

JIT VM Access requires that you use Azure Security Center. This requires an extra monthly cost per VM, so keep that in mind if you are thinking about configuring JIT.

To configure a VM for JIT, perform the following steps:

1. Open Azure Security Center in the Azure portal.
2. Under Advanced Cloud Defense, click Just in Time VM Access.
3. On the Recommended tab, review the list of VMs for which JIT is recommended, as shown in Figure 17-10. Select the VM for which you wish to enable JIT and then click Enable JIT On 1 VMs (this dialog box will change depending on the number of VMs you select).

Technet24

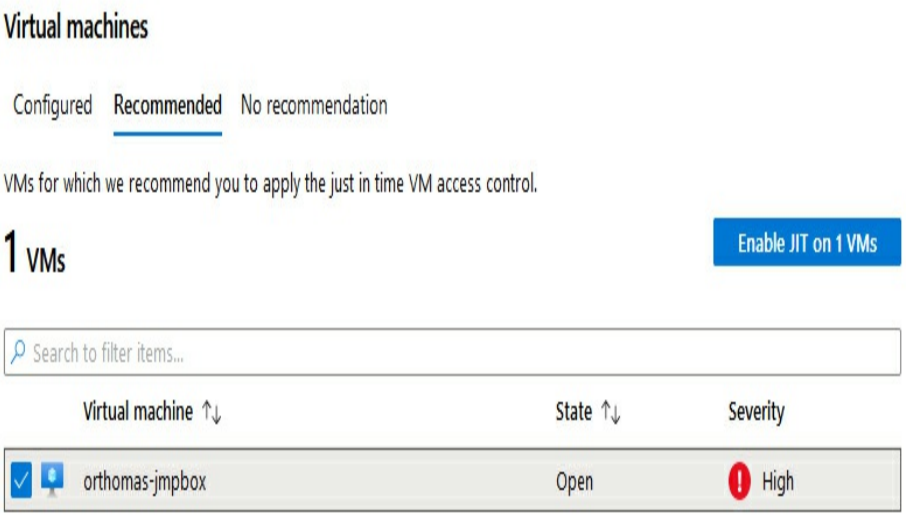


Figure 17-10 Enable JIT

4. On the JIT Access Configuration page, click the management port that you want to configure.
5. On the Add Port Configuration page, shown in Figure 17-11, under Allowed Source IPs, specify whether you want to allow source IPs to be configured on a Per Request basis or per CIDR Block. If you choose the Per Request option, when a JIT request is made, the IP address used to access the portal will be the one granted access to the management port. If you choose the CIDR Block option, then a specific pre-configured IP address range will be granted access. You can also choose a Max Request Time on this page. This determines the length of the window for new requests on this management port. It will not terminate an existing connection, so even if the maximum request time is three hours, if you establish an RDP session to the host, it will stay open as long as you want and will not close automatically once the Max Request Time period has expired.



Figure 17-11 Add Port Configuration

To gain access via JIT once it is configured for a VM, perform the following steps:

1. In Azure Security Center, under Advanced Cloud Defense, click Just In Time VM Access.
2. On the Configured tab, select the VM for which JIT is configured and click Request Access.
3. On the Request Access page, select the port for which you wish to request access and set the toggle to On. Verify that you have selected either My IP or IP Range. You may also provide a message to be appended to the log entry to the request detailing why access was requested. Finally, click Open Ports. After a few moments, the port will be open, and you'll be able to establish a remote management session using Remote Desktop Connection software.

More Info

Just in Time VM Access

You can learn more about Just in Time VM Access at https://docs.microsoft.com/en-us/azure/security-center/security-center-just-in-time.

Resizing a VM

Azure allows you to resize an IaaS VM. Resizing a VM allows you to change the IaaS VM's processor and memory allocation. Resizing a VM might also alter the number of disks that can be associated with a VM; cheaper SKUs will be limited to less storage, and in some cases, they will be limited to standard rather than premium storage types. If the Azure IaaS VM is running, the only drawback is that the VM will need to restart for the resize to occur.

To resize an Azure IaaS VM, perform the following steps:

1. Open the VM's page in the Azure portal under Virtual Machines.
2. In the left menu under Settings, click Size.
3. Select the new size that you wish to apply to the virtual machine.
4. Click Resize.

The virtual machine will then restart, and the new size will apply.

More Info

Resizing an Azure IaaS VM

You can learn more about resizing an Azure IaaS VM at https://docs.microsoft.com/en-us/azure/virtual-machines/windows/resize-vm.

IaaS VM backup

An advantage to running an IaaS VM in Azure is that it is far simpler to set up a daily backup than it is when you have to manage all the backup and recovery infrastructure yourself. This is because Azure IaaS VM Backup is a feature of Azure. Essentially, all you need to do is enable Azure IaaS VM Backup, and it looks after itself.

To enable backup and recovery of an Azure IaaS VM, perform the following steps:

1. In Cloud Shell, run the following command to create a recovery services vault to host the backup data:

   Click here to view code image

   ```
   az backup vault create --resource-group
   VMResourceGroup --name
   IaaSVMRecoveryServicesVault --location
   southeasteasia
   ```

2. Run the following command to enable backup for an Azure VM:

   Click here to view code image

   ```
   az backup protection enable-for-vm --resource-
   group VMResourceGroup --vault-name
   IaaSVMRecoveryServicesVault --vm NameOfVM --
   policy-name DefaultPolicy
   ```

3. Run the following command to trigger a manual backup for an Azure VM:

Click here to view code image

```
az backup protection backup-now  --resource-
group VMResourceGroup --vault-name
IaaSVMRecoveryServicesVault
--container-name NameOfVM --item-name NameOfVM
--retain-until 12-12-2020
```

More Info

Azure IaaS VM Backup

You can learn more about Azure IaaS VM Backup at https://docs.microsoft.com/en-us/azure/backup/quick-backup-vm-cli.

IaaS VM Encryption

To support IaaS VM disk encryption, the Windows Server IaaS VM must be able to do the following. (By default, IaaS VM can do this unless you remove the default network security group rules.)

- The server must be able to connect to the key vault endpoint so that the Windows Server IaaS VM can store and retrieve encryption keys.
- The server must be able to connect to an Azure Active Directory endpoint at *login.microsoftonline.com* so that it can retrieve a token that allows it to connect to the key vault that holds the encryption keys.
- The server must be able to connect to an Azure storage endpoint that hosts the Azure extension repository and the Azure storage account that stores the VM's virtual hard disks.
- If the Windows Server IaaS VM is domain joined, do not configure BitLocker-related Group Policies other than the Configure User

Storage Of BitLocker Recovery Information: Allow 256-Bit Recovery Key policy. This policy is usually configured automatically during the encryption process, and Azure Disk Encryption will fail if a Group Policy conflict in any TPM- or BitLocker-related policies exists.

Perform the following steps to encrypt the hard disk drives of a Windows Server IaaS VM:

1. Create an Azure Key Vault to store the encryption key used to encrypt the hard disk of the VM. You can do this with the following code:

   Click here to view code image

   ```
   az keyvault create --name "KeyVaultName" --
   resource-group "VMResourceGroup"
   --location AzureRegion --enabled-for-disk-
   encryption
   ```

2. Once the key vault is created, issue the following command to encrypt the VM.

   Click here to view code image

   ```
   az vm encryption enable -g VMResourceGroup --
   name VMName --disk-encryption-
   keyvault KeyVaultName
   ```

More Info

Encrypt Azure IaaS VM disks

You can learn more about encrypting Azure IaaS VM disks at https://docs.microsoft.com/en-us/azure/security/fundamentals/azure-disk-encryption-vms-vmss.

AZURE ACTIVE DIRECTORY

Azure Active Directory is Microsoft's identity and access management service hosted in Azure. It is mostly used as the identity service for products such as Office 365, but you can also integrate it with Azure services such as Azure AD DS, which allows you to domain join IaaS VMs and replicate some of the functionality of an on-premises Active Directory instance. You can also configure synchronization between an on-premises Active Directory instance and an Azure AD instance, which allows the same username and password to be used on-premises and in Azure. You can also configure self-service password reset, which allows your on-premises users to navigate to a web page hosted in Azure to reset their on-premises passwords without having to contact your organization's service desk or to deploy a product such as Microsoft Identity Manager.

When you create an Azure subscription, the subscription tenancy is automatically assigned a custom *onmicrosoft.com* domain. The tenant name is in the format *name.onmicrosoft.com*, where *name* is a name based on the email address that you used to sign up to the Azure account with. This name is unique, and no two organizations share the same tenant name. While you're unlikely to actually use the *onmicrosoft.com* domain name once you have fully configured your organization's tenancy, it's important to note that the tenant name cannot be changed after you configure your

Azure subscription. This name remains with the subscription over the course of the subscription's existence and cannot be removed.

You can assign a domain name that you own to the Azure AD tenancy so that you don't have to use the tenant name on a regular basis. Once you set this default domain name, new Azure AD accounts will be assigned this name in their UPN suffixes. You'll also be able to convert existing accounts to use this UPN suffix. Figure 17-12 shows *tailwindtraders.net* assigned as the primary domain name to an Azure Active Directory instance.

Azure AD supports the addition of up to 900 domains to a single subscription. You can use separate domain names with a subscription, such as *contoso.com* or *tailwindtraders.com*. You can also associate subdomains of a domain name such as *partners.tailwindtraders.com* or *australia.contoso.com*.

Technet24

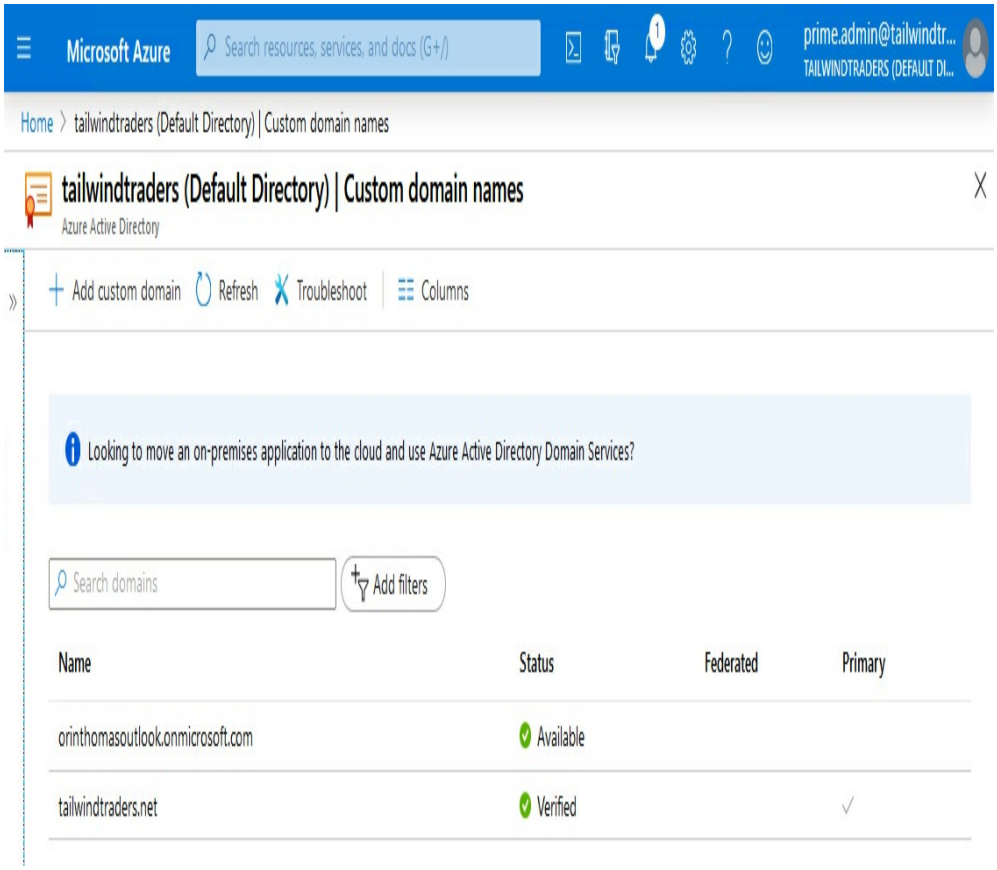


Figure 17-12 Azure AD custom domain names

More Info

Azure Active Directory

You can learn more about Azure AD at https://docs.microsoft.com/en-us/azure/active-directory/.

Inside OUT

A lot more than identity synchronization

Azure AD does far more than what is mentioned here. It can function as the identity service for applications, supports conditional access to web applications, and can be easily integrated with third-party identity providers. For the purposes of this book, we are just exploring how you might use Azure AD to support IaaS VMs and hybrid services.

Azure Active Directory Connect

Azure AD Connect allows you to connect your on-premises Active Directory accounts with an Azure AD instance. This is useful not only for applications running in Azure, but it allows you to implement single sign-on if your organization is using Microsoft 365 or Office 365. Single sign-on allows you to use one identity to access on-premises and cloud resources. In many scenarios, the user won't even be required to reauthenticate.

Azure AD Connect supports a variety of sign-in options. You configure which one you want to use when setting up Azure AD Connect. The default method, Password Synchronization, is appropriate for the majority of organizations that will use Azure AD Connect to synchronize identities to the cloud.

Password synchronization

Hashes of on-premises Active Directory user passwords synchronize to Azure AD, and changed passwords immediately synchronize to Azure AD. Actual passwords are never sent to Azure AD and are not stored in Azure AD. This allows for single sign-on for users of computers that are joined to an Active Directory domain that synchronizes to Azure AD. Also, password synchronization allows you to enable password write-back for self-service password reset functionality

through Azure AD.

Pass-through authentication

When authenticating to Azure AD, the user's password is validated against an on-premises Active Directory domain controller. Passwords and password hashes are not present in Azure AD. Pass-through authentication allows for on-premises password policies to apply. Pass-through authentication requires that Azure AD Connect have an agent on a computer joined to the domain that hosts the Active Directory instance that contains the relevant user accounts. Pass-through authentication also allows single sign-on for users of domain joined machines.

With pass-through authentication, the user's password is validated against the on-premises Active Directory controller. The password doesn't need to be present in Azure AD in any form. This allows for on-premises policies, such as sign-in hour restrictions, to be evaluated during authentication to cloud services.

Pass-through authentication uses a simple agent on a Windows Server 2012 R2, Windows Server 2016, or Windows Server 2019 domain-joined machine in the on-premises environment. This agent listens for password validation requests. It doesn't require any inbound ports to be open to the Internet.

In addition, you can also enable single sign-on for users on domain-joined machines that are on the corporate network. With single sign-on, enabled users only need to enter a username to help them securely access cloud resources.

Active Directory Federation

This allows users to authenticate to Azure AD resources using on-premises credentials. When you choose the Federation With AD FS option, Active Directory Federation Services is installed and configured; also, a Web Application Proxy server to facilitate communication between the on-premises AD FS deployment and Microsoft Azure Active Directory is installed. This is the most complicated identity synchronization configuration, and it is only likely to be implemented in environments with complicated identity configurations.

More Info

Azure AD Connect sign-in options

You can learn more about sign-in options by consulting the following article: https://docs.microsoft.com/en-us/azure/active-directory/connect/active-directory-aadconnect-user-signin.

Azure AD Connect server requirements

Azure AD Connect is software that you install on a computer that manages the process of synchronizing

objects between the on-premises Active Directory and the Azure Active Directory instance. You can install Azure AD Connect on computers running the Windows Server 2012 or later operating systems:

Azure AD Connect has the following requirements:

- It must be installed on a Windows Server instance that has the GUI version of the operating system installed. You cannot install Azure AD connect on a computer running the Server Core operating system.
- You can deploy Azure AD Connect on a computer that is either a domain controller or a member server. If you use the custom options, a standalone server can be used.
- The server hosting Azure AD Connect requires .NET Framework 4.5.1 or later.
- The server hosting Azure AD Connect requires Microsoft PowerShell 3.0 or later.
- The server hosting Azure AD Connect must not have PowerShell Transcription enabled through Group Policy.
- If you are deploying Azure AD Connect with Active Directory Federation Services, you must use Windows Server 2012 R2 or later for the Web Application Proxy, and Windows remote management must be enabled on the servers that will host AD FS roles.
- If global administrators will have multifactor authentication (MFA) enabled, then the URL *https://secure.aadcdn.microsoftonline-p.com* must be configured as a trusted site.

Inside OUT

Forest functional level

Azure AD Connect requires that the on-premises Active Directory environment be configured at the Windows Server 2003 forest functional level or higher. Windows

Server 2019 domain controllers support the Windows Server 2016 forest functional level.

Connectivity requirements

The computer with Azure AD Connect installed must be a member of a domain in the forest that you want to synchronize, and it must have connectivity to a writable domain controller in each domain of the forest you wish to synchronize on the following ports:

- **DNS.** TCP/UDP Port 53
- **Kerberos.** TCP/UDP Port 88
- **RPC.** TCP Port 135
- **LDAP.** TCP/UDP Port 389
- **SSL.** TCP Port 443
- **SMB.** TCP 445

The computer with Azure AD Connect installed must be able to establish communication with the Microsoft Azure servers on the Internet over TCP port 443. The computer with Azure AD Connect installed can be located on an internal network as long as it can initiate communication on TCP port 443. The computer hosting Azure AD Connect does not need a publicly routable IP address. The computer hosting Azure AD Connect always initiates synchronization communication to Microsoft Azure. Microsoft Azure Active Directory does not initiate synchronization communication to the

computer hosting Azure AD Connect on the on-premises network.

Microsoft recommends that you do not install Azure AD Connect on a domain controller. If you are going to be replicating more than 50,000 objects, Microsoft recommends that you deploy SQL Server on a computer that is separate from the computer that will host Azure AD Connect. If you plan to host the SQL Server instance on a separate computer, ensure that communication is possible between the computer hosting Azure AD Connect and the computer hosting the SQL Instance on TCP port 1433.

If you are going to use a separate SQL Server instance, ensure that the account used to install and configure Azure AD Connect has “systems administrator” rights on the SQL instance and that the service account used for Azure AD Connect has “public” permissions on the Azure AD Connect database.

SQL Server requirements

When you deploy Azure AD connect, you have the option of having Azure AD Connect install a SQL Server Express instance, or you can choose to have Azure AD Connect leverage a full instance of SQL Server. SQL Server Express is limited to a maximum database size of 10 GB. In terms of Azure AD Connect, this means that Azure AD Connect is only able to manage 100,000

objects. This is likely to be adequate for all but the largest environments.

For environments that require Azure AD Connect to manage more than 100,000 objects, you'll need to have Azure AD Connect leverage a full instance of SQL Server. Azure AD Connect can use all versions of Microsoft SQL Server, from Microsoft SQL Server 2012 with the most recent service pack to SQL Server 2019. It is important to note that SQL Azure is not supported as a database for Azure AD Connect. If you are deploying a full instance of SQL Server to support Azure AD Connect, ensure that the following prerequisites are met:

- **Use a case-insensitive SQL collation.** Case-insensitive collations have the *_CI_* identifier included in their names. Case-sensitive collations (those that use the *_CS_* designation) are not supported for use with Azure AD Connect.
- **You can only use one sync engine per SQL instance.** If you have an additional Azure AD Connect sync engine or if you are using Microsoft Identity Manager in your environment, each sync engine requires its own separate SQL instance.

Requirements for deployment accounts

You use two accounts when configuring Azure AD Connect. One account must have specific Azure AD permissions; the other account must have specific on-premises Active Directory permissions. The accounts that you use to install and configure Azure AD Connect

Technet24

have the following requirements:

- The account used to configure Azure AD Connect must have Global Administrator privileges in the Azure AD tenancy. You should create a separate account for this task and configure the account with a complex password that does not expire. This account is used for the synchronization process between on-premises AD and Azure AD.
- The account used to install and configure Azure AD Connect must have Enterprise Administrator permissions within the on-premises Active Directory forest if you will be using Express installation settings. This account is only required during installation and configuration. Once Azure AD Connect is installed and configured, this account no longer needs Enterprise Administrator permissions. The best practice is to create a separate account for Azure AD Connect installation and configuration and to temporarily add this account to the Enterprise Admins group during the installation and configuration process. Once Azure AD Connect is installed and configured, this account can be removed from the Enterprise Admins group. You should not attempt to change the account used after Azure AD Connect is set up and configured because Azure AD Connect always attempts to run using the original account.
- The account used to install and configure Azure AD Connect must be a member of the local Administrators group on the computer on which Azure AD Connect is installed.

Installing Azure AD Connect

Installing Azure AD Connect with Express settings is appropriate if your organization has a single Active Directory forest and you wish to use password synchronization for authentication. The Azure AD Connect Express settings are appropriate for most organizations. You can download the Azure AD Connect

installation files from Microsoft's download center website.

To install Azure AD Connect with Express settings, perform the following steps:

1. Double-click the *AzureADConnect.msi* file that you've downloaded from the Microsoft download center. You will be prompted with a security warning. After clicking Run, Azure AD Connect will be installed on your computer. When the installation is complete, you will be presented with the splash screen detailing license terms and displaying a privacy notice. You'll need to agree to these terms before clicking Continue.
2. If your organization has an internal non-routable domain, it will be necessary for you to use custom settings. The best practice is to use domain synchronization when your on-premises Active Directory instance and your Azure Active Directory instance use the same routable domain name. Click Continue.
3. On the Install Required Components page, shown in Figure 17-13, choose between the following options:
    1. **Specify A Custom Installation Location.** Choose this option if you want to install Azure AD Connect in a separate location, such as on another volume.
    2. **Specify An Existing SQL Server.** Choose this option if you want to specify an alternate SQL server instance. By default, Azure AD Connect will install a SQL Server Express instance.
    3. **Use An Existing Service Account.** You can configure Azure AD Connect to use an existing service account. By default, Azure AD Connect will create a service account. You can configure Azure AD Connect to use a Group Managed Service account. You'll need to use an existing service account if you are using Azure AD Connect with a remote SQL Server instance or if communication with Azure will occur through a proxy server that requires authentication.
    4. **Specify Custom Sync Groups.** When you deploy Azure AD

Technet24

Connect, it will create four local groups on the server that hosts the Azure AD Connect instance. These groups are the Administrators group, Operators group, Password Reset group, and the Browse group. If you want to use your own set of groups, you can specify them here. These groups must be local to the host server and not a member of the domain.

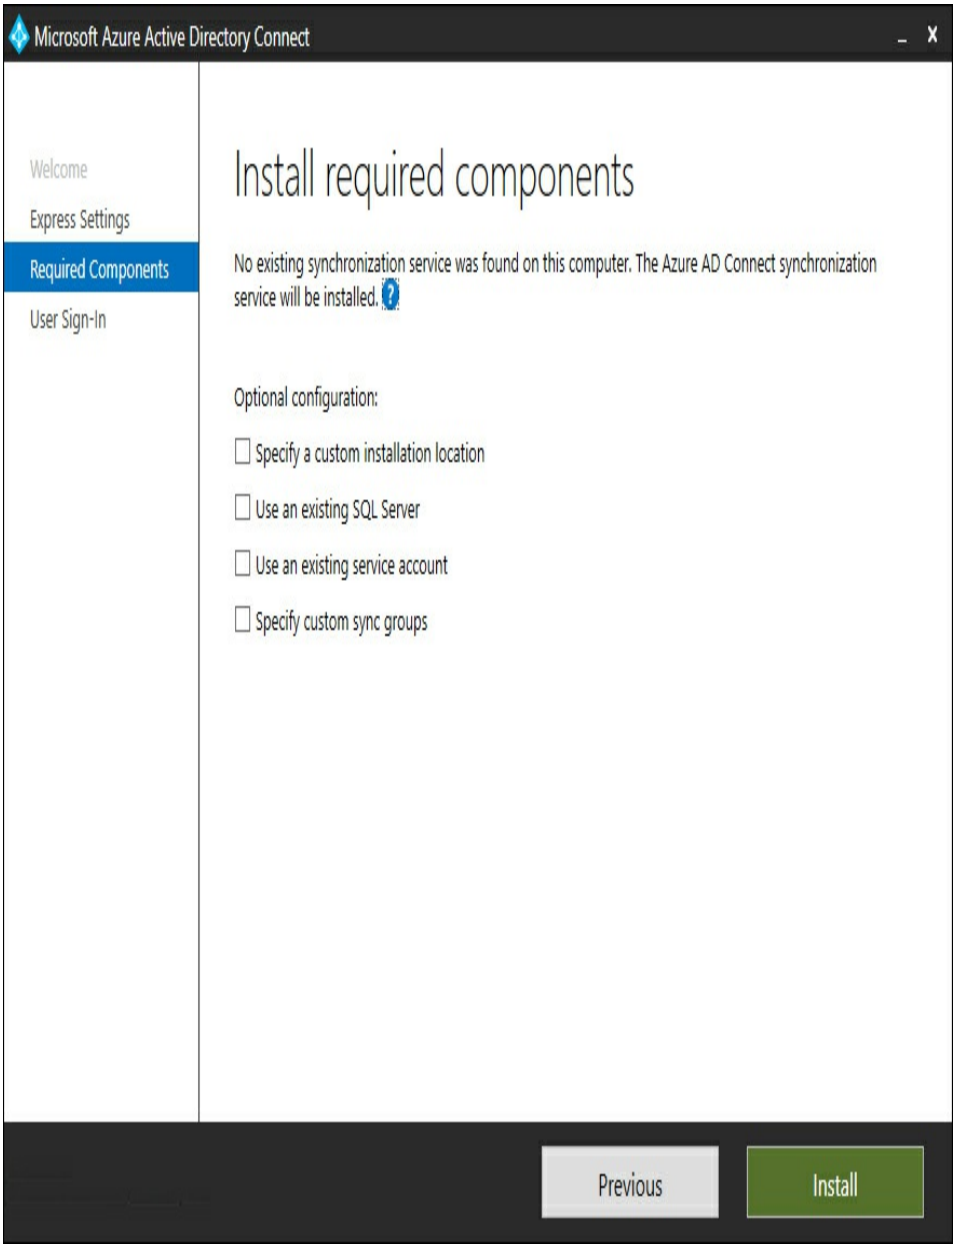


Figure 17-13 Install Required Components

4. Once you have specified which custom options you require—you are not required to select any options—click Install.
5. On the User Sign-In page, shown in Figure 17-14, specify what type of sign on you want to allow. You can choose between the following options, the details of which were covered earlier in this chapter. Most organizations will choose Password Synchronization because this is the most straightforward option:
    1. Password Synchronization
    2. Pass-Through Authentication
    3. Federation With AD FS
    4. Federation With PingFederate
    5. Do Not Configure
    6. Enable Single Sign-On

Technet24

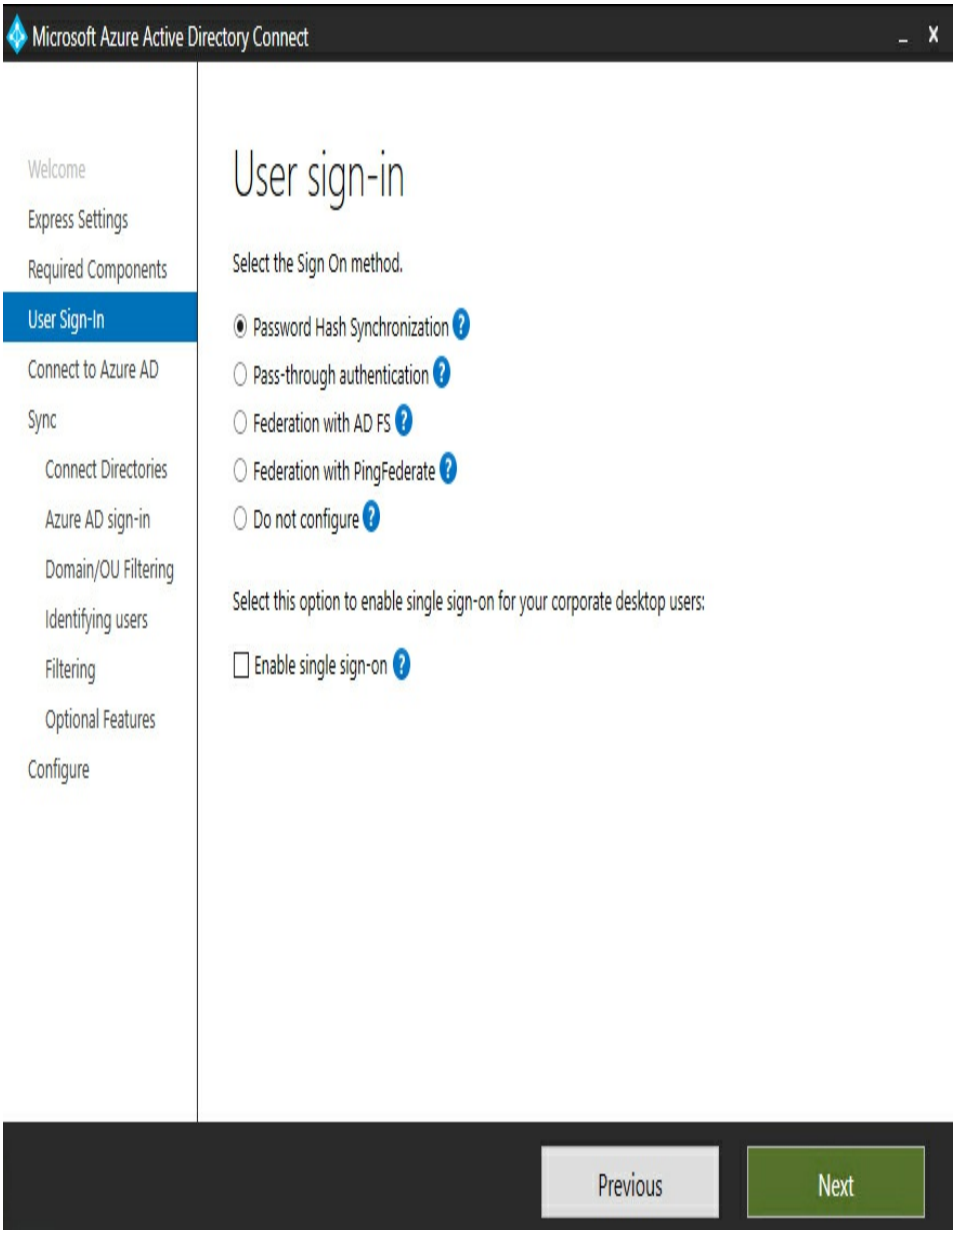


Figure 17-14 User Sign-In options

6. On the Connect To Azure AD page, provide the credentials of an account with Global Administrator privileges in Azure AD. Microsoft recommends you use an account in the default *onmicrosoft.com* domain associated with the Azure AD instance to which you will be connecting. If you choose the Federation With AD FS option, ensure that you do not sign in using an account in a domain that you will

enable for federation. Figure 17-15 shows a sign-in with a Password Synchronization scenario.

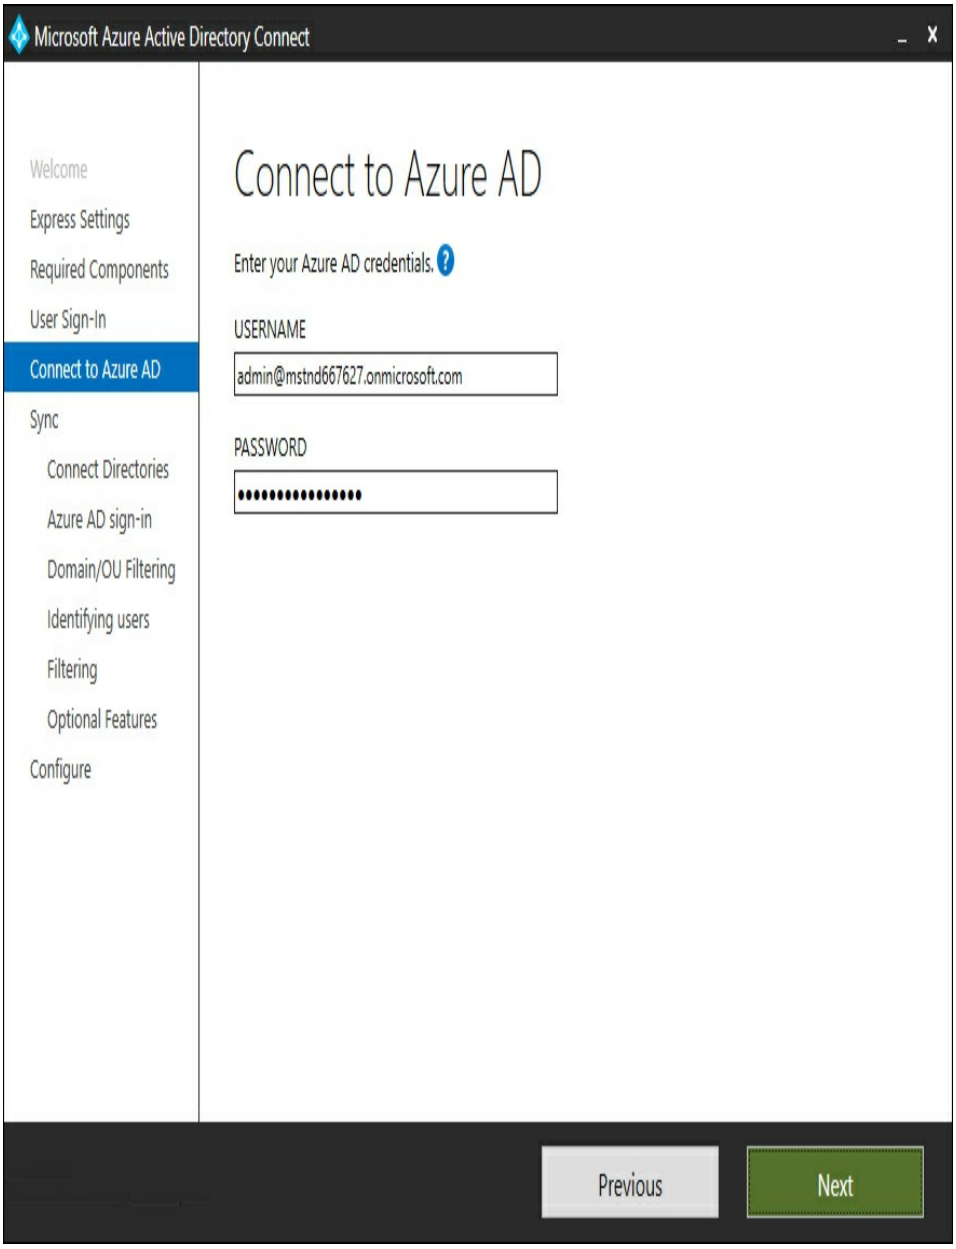


Figure 17-15 Connect to Azure AD

7. Once Azure AD Connect has connected to Azure AD, you will be able to specify the directory type to synchronize as well as the forest. Click Add Directory to add a forest. When you add a forest by clicking Add

Directory, you will need to specify the credentials of an account that will perform periodic synchronization. Unless you are certain that you have applied the minimum necessary privileges to an account, you should provide Enterprise Administrator credentials and allow Azure AD Connect to create the account, as shown in Figure 17-16. This will ensure that the account is only assigned the privileges necessary to perform synchronization tasks.

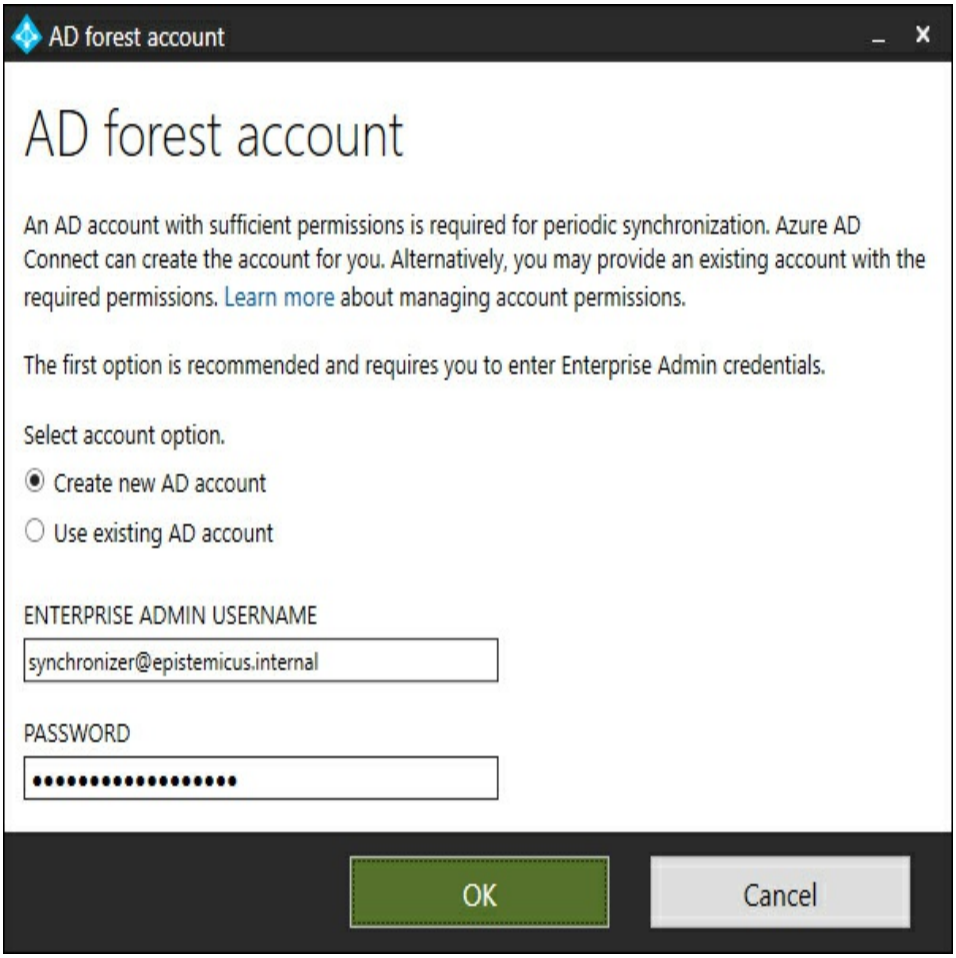


Figure 17-16 AD Forest Account

8. Once the credentials have been verified, as shown in Figure 17-17, click Next.

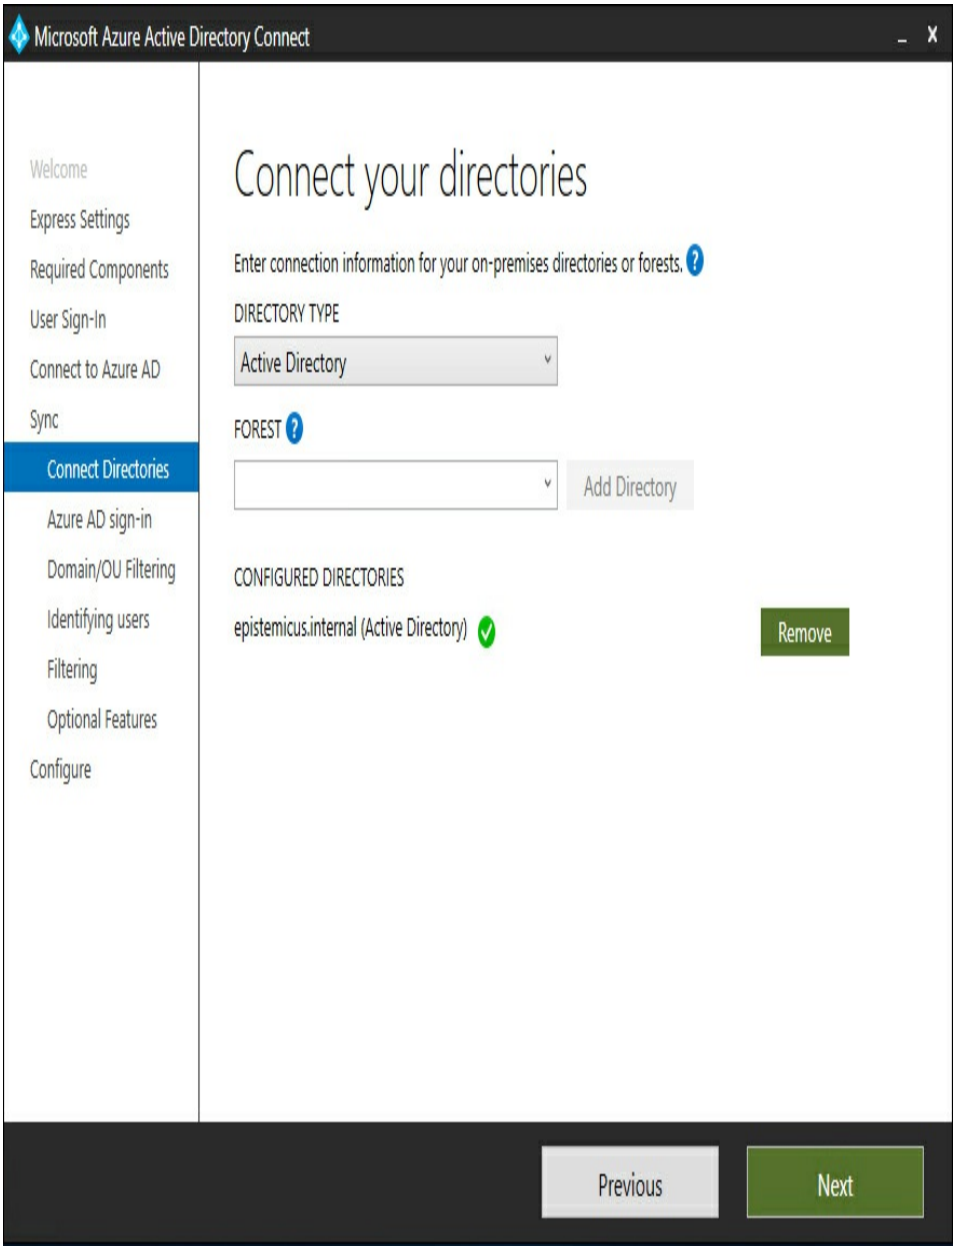


Figure 17-17 Connect Your Directories

9. On the Azure AD Sign-In Configuration page, shown in Figure 17-18, review the UPN suffix and then inspect the on-premises attribute to use as the Azure AD username. You'll need to ensure that accounts use a routable Azure AD username.

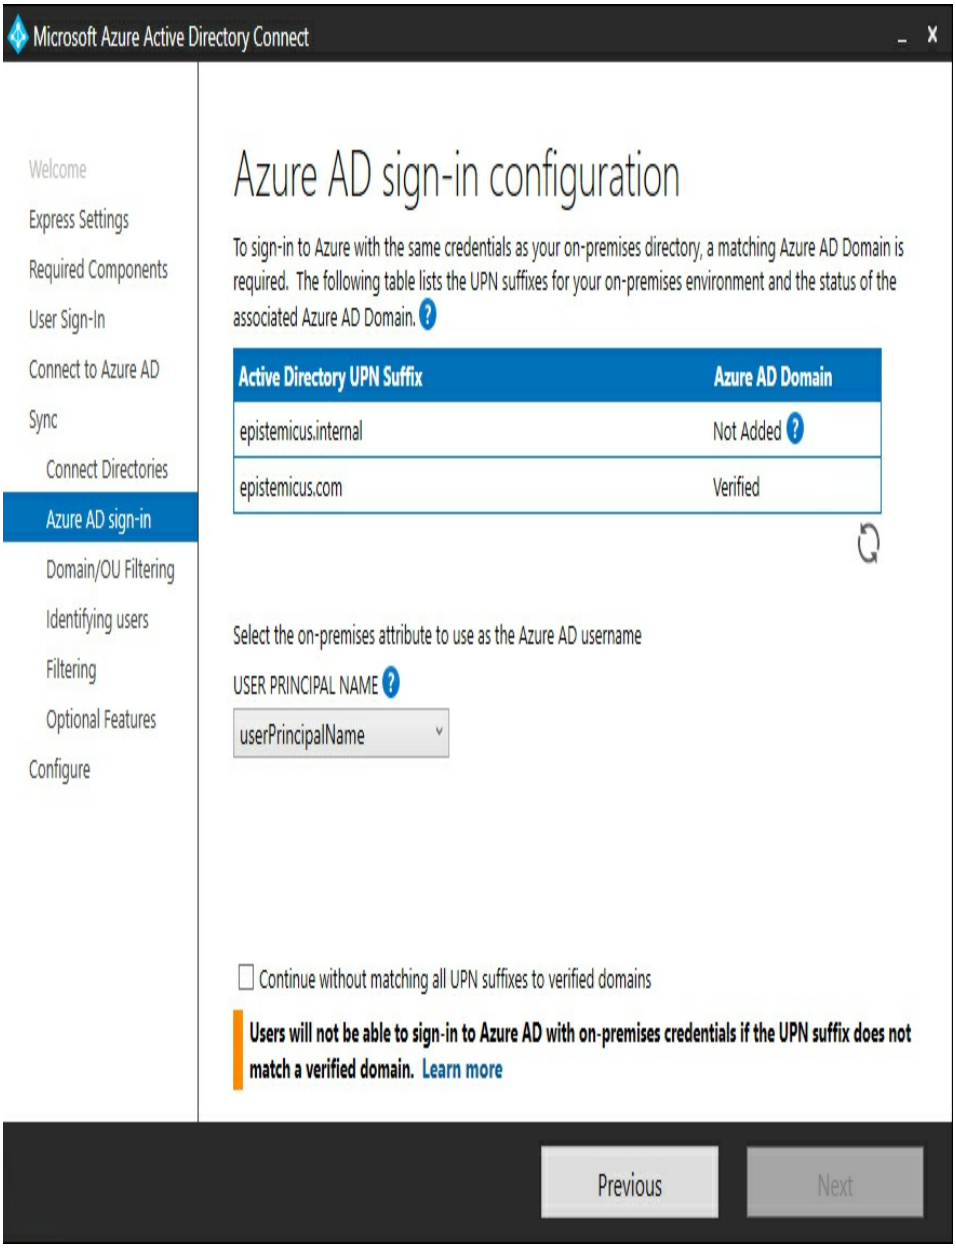


Figure 17-18 Azure AD Sign-In Configuration

10. On the Domain And OU Filtering page, select whether you want to sync all objects or just objects in specific domains and OUs.
11. On the Uniquely Identifying Your Users page, shown in Figure 17-19, specify how users are to be identified. By default, users should only have one representation across all directories. If users exist in multiple

directories, you can have matches identified by a specific active directory attribute, with the default being the Mail Attribute.

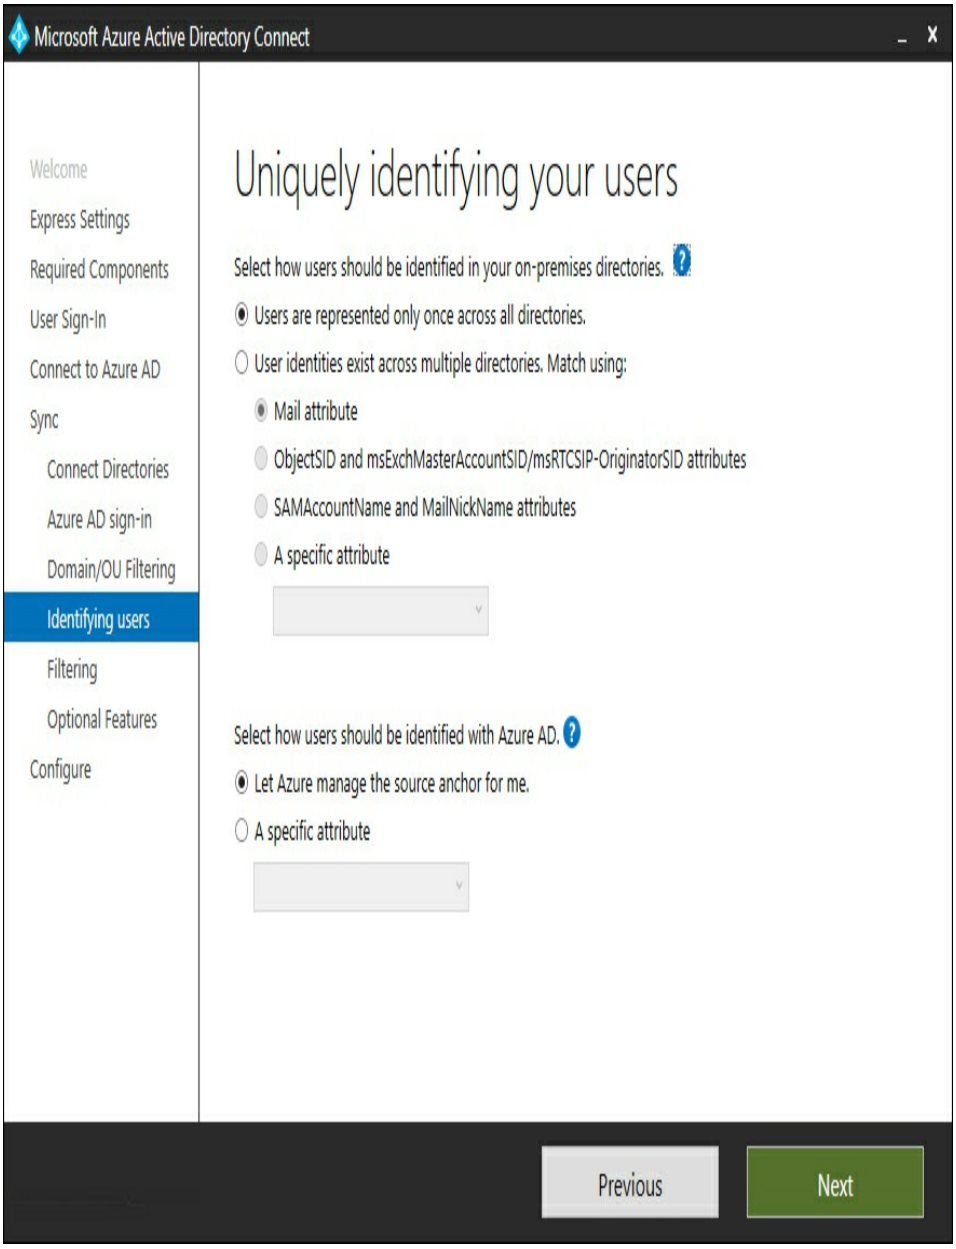


Figure 17-19 Uniquely Identifying Your Users

12. On the Filter Users And Devices page, specify whether you want to synchronize all users and devices or only members of a specific group. Figure 17-20 shows members of the Microsoft 365-Pilot-Users group

being configured so that their accounts will be synchronized with Azure.

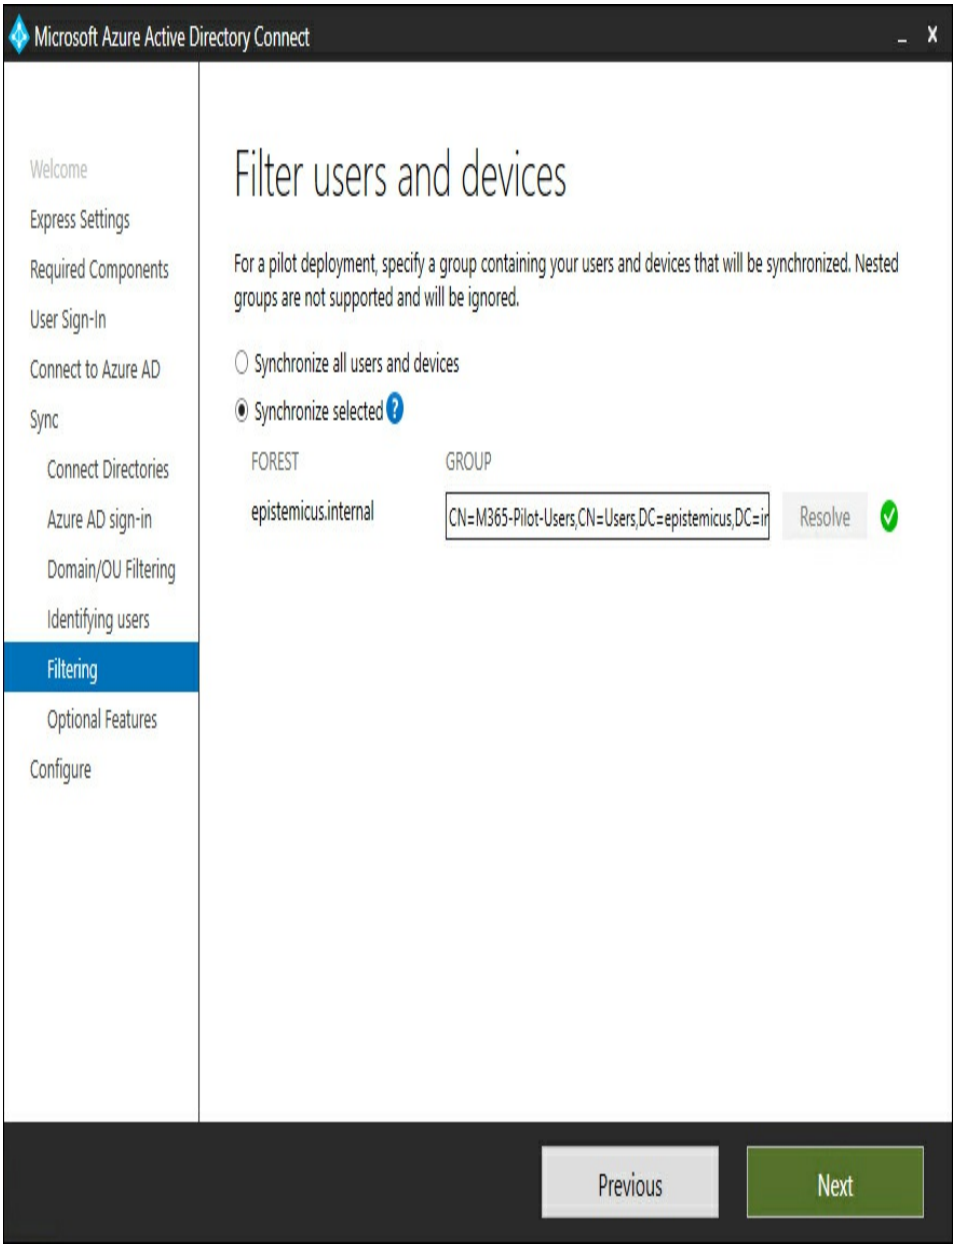


Figure 17-20 Filter Users And Devices

13. On the Optional Features page, shown in Figure 17-21, select any optional features that you want to configure. These features include:

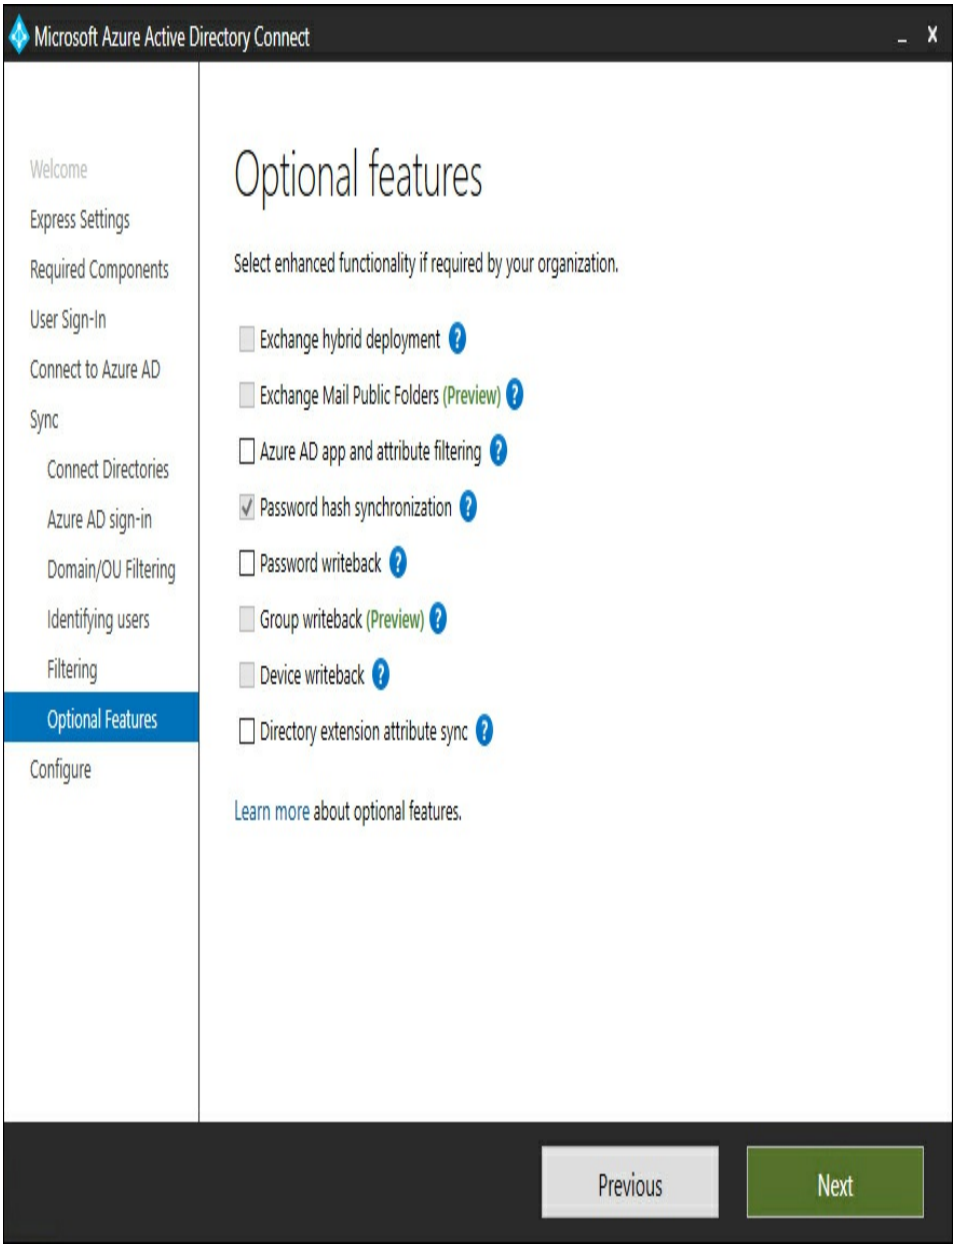


Figure 17-21 Optional Features

1. **Exchange Hybrid Deployment.** This option is suitable for organizations that have an Office 365 deployment and where there are mailboxes hosted both on-premises and in the cloud.
2. **Exchange Mail Public Folders.** This feature allows organizations to synchronize mail-enabled public folder objects from an on-

premises Active Directory environment to Microsoft Office 365.

3. **Azure AD App And Attribute Filtering.** Selecting this option gives you the ability to be more selective about which attributes are synchronized between the on-premises environment and Azure AD.
4. **Password Hash Synchronization.** This option synchronizes a hash of the user's on-premises password Azure AD. When the user authenticates to Azure AD, the submitted password is hashed using the same process, and if the hashes match, the user is authenticated. Each time a user updates his or her password on-premises, the updated password hash synchronizes to Azure AD.
5. **Password Writeback.** Password writeback allows users to change their passwords in the cloud and have the changed password written back to the on-premises Active Directory instance.
6. **Group Writeback.** Changes made to groups in Azure AD are written back to the on-premises AD instance.
7. **Device Writeback.** Information about devices registered by the user in Azure AD is written back to the on-premises AD instance.
8. **Directory Extension Attribute Sync.** Allows you to extend Azure AD schema based on extensions made to your organization's on-premises Active Directory instance.

14. On the Ready To Configure page, you can choose to start synchronization or to enable staging mode. When you configure staging mode, Azure AD Connect will prepare the synchronization process, but it will not synchronize any data with Azure AD.

Using UPN suffixes and non-routable domains

Prior to performing synchronization between an on-premises Active Directory environment and an Azure Active Directory instance, you must ensure that all user account objects in the on-premises Active Directory

environment are configured with a value for the UPN suffix that can function for both the on-premises environment and any application that you want to use it with in the cloud.

This is not a problem when an organization's internal Active Directory domain suffix is a publicly routable domain. For example, a domain name, such as *contoso.com* or *adatum.com*, which is resolvable by public DNS servers, will suffice. Things become more complicated when the organization's internal Active Directory domain suffix is not publicly routable.

If a domain is non-routable, the default Azure AD instance domain, such as *adatum2020.onmicrosoft.com*, should be used for the UPN suffix. This requires modifying the UPN suffix of accounts stored in the on-premises Active Directory instance. Modification of UPN after the initial synchronization has occurred is not supported. So, you need to ensure that on-premises Active Directory UPNs are properly configured prior to performing initial synchronization using Azure AD Connect.

Perform the following steps to add a UPN suffix to the on-premises Active Directory if the Active Directory domain uses a non-routable namespace:

1. Open the Active Directory Domains And Trusts console and select Active Directory Domains And Trusts.

2. On the Action menu, click Properties.
3. On the UPN Suffixes tab, enter the UPN suffix to be used with Azure Active Directory. Figure 17-22 shows the UPN suffix of *epistemicus.com*.

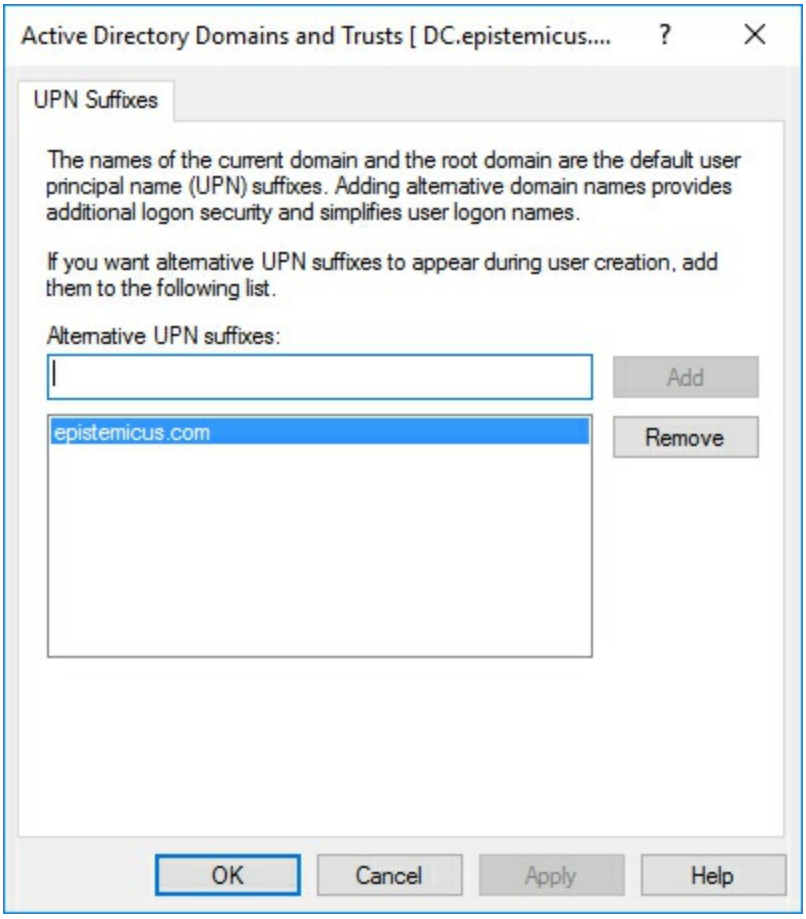


Figure 17-22 Configure a UPN suffix for routable domain

4. Once the UPN suffix has been added in Active Directory Domains And Trusts, you can assign the UPN suffix to user accounts. You can do this manually, as shown in Figure 17-23, by using the Account tab of the user's Properties dialog box in Active Directory Users And Computers.

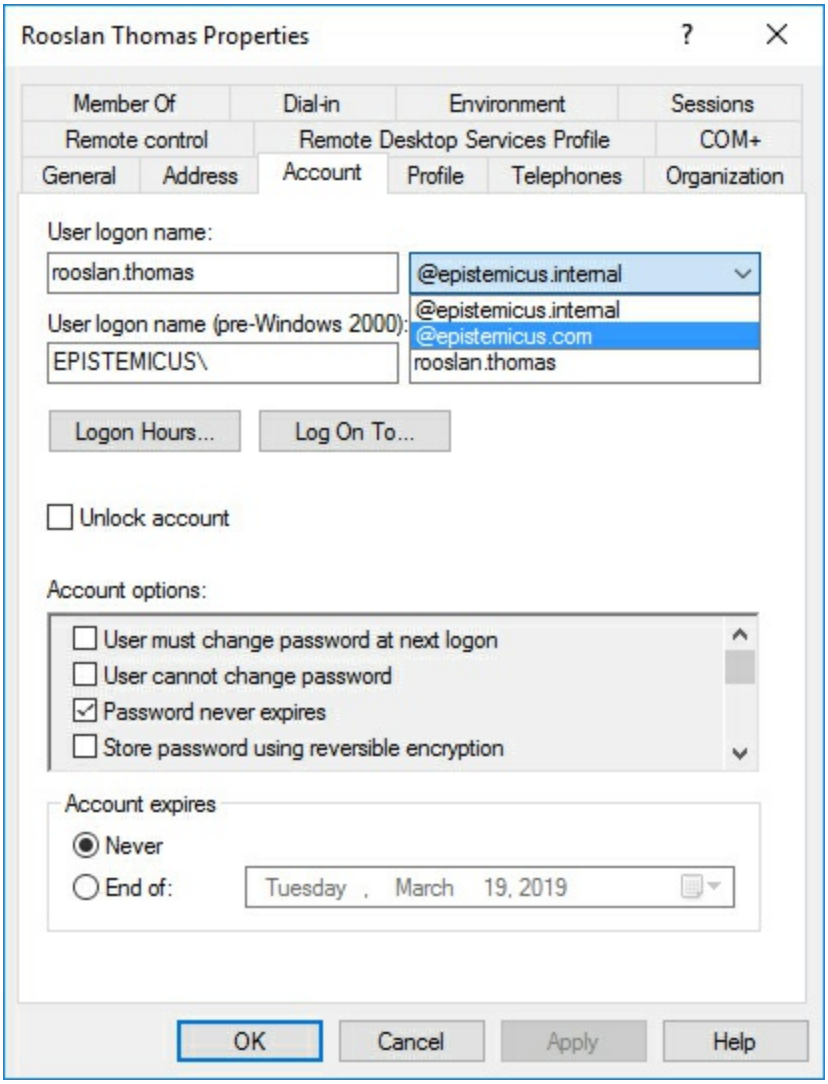


Figure 17-23 Configure UPN

5. You can use tools like ADModify.NET to reset the UPNs of multiple accounts, as shown in Figure 17-24.

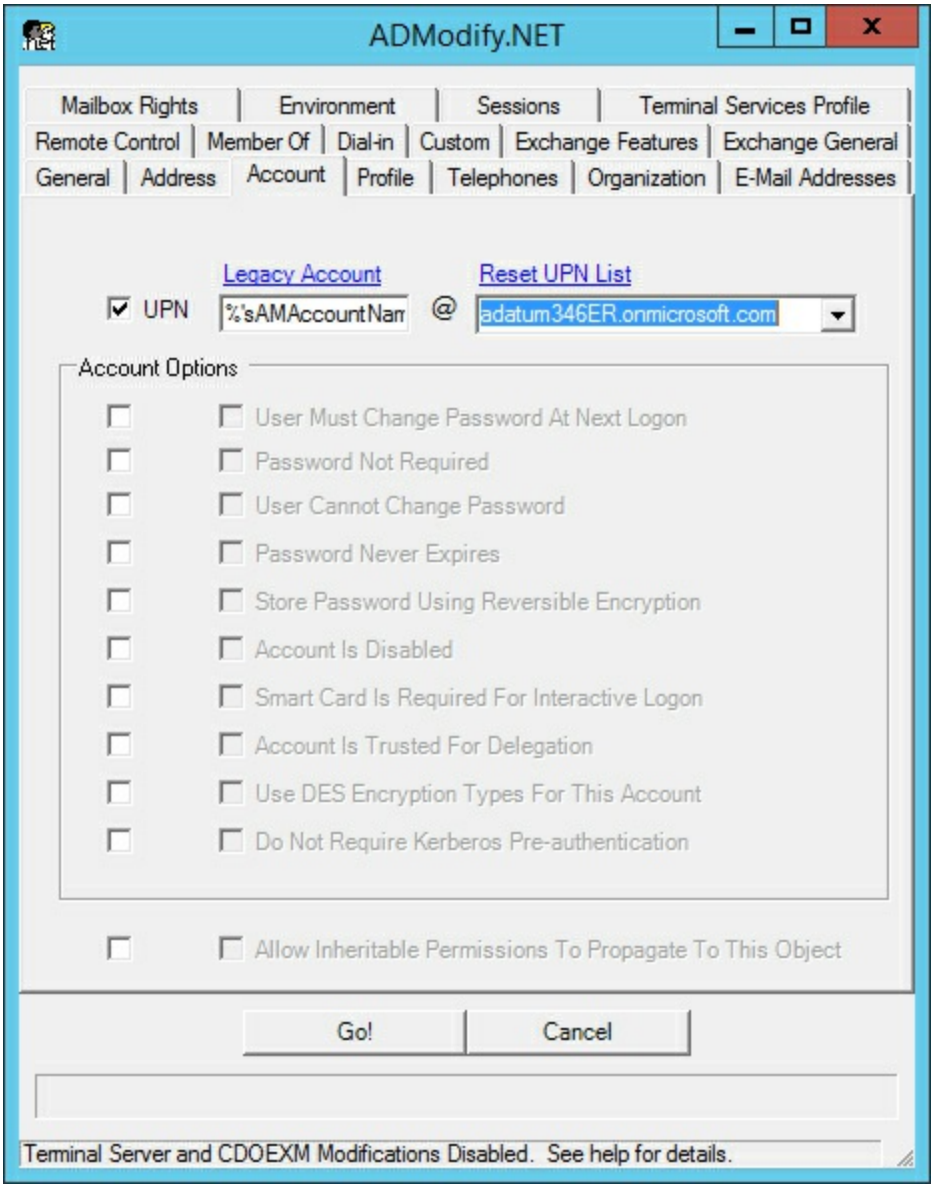


Figure 17-24 ADModify.NET

6. You can also use Microsoft PowerShell scripts to reset the UPNs of multiple user accounts. For example, the following script resets UPN suffixes of all user accounts in the *epistemicus.internal* domain to *epistemicus.onmicrosoft.com*.

   Click here to view code image

```
Get-ADUser -Filter {UserPrincipalName -like
```

```
"*@epistemicus.internal"} -SearchBase
"DC=epistemicus,DC=internal" |
ForEach-Object {
$UPN =
$_.UserPrincipalName.Replace("epistemicus.inter
nal","epistemicus.onmicrosoft.com")
Set-ADUser $_ -UserPrincipalName $UPN
}
```

Monitor Azure AD Connect Health

Azure AD Connect Health is a tool available in the Azure Active Directory Admin Center, shown in Figure 17-25, that allows you to monitor the health of synchronization between your organization's on-premises directory and Azure Active Directory.

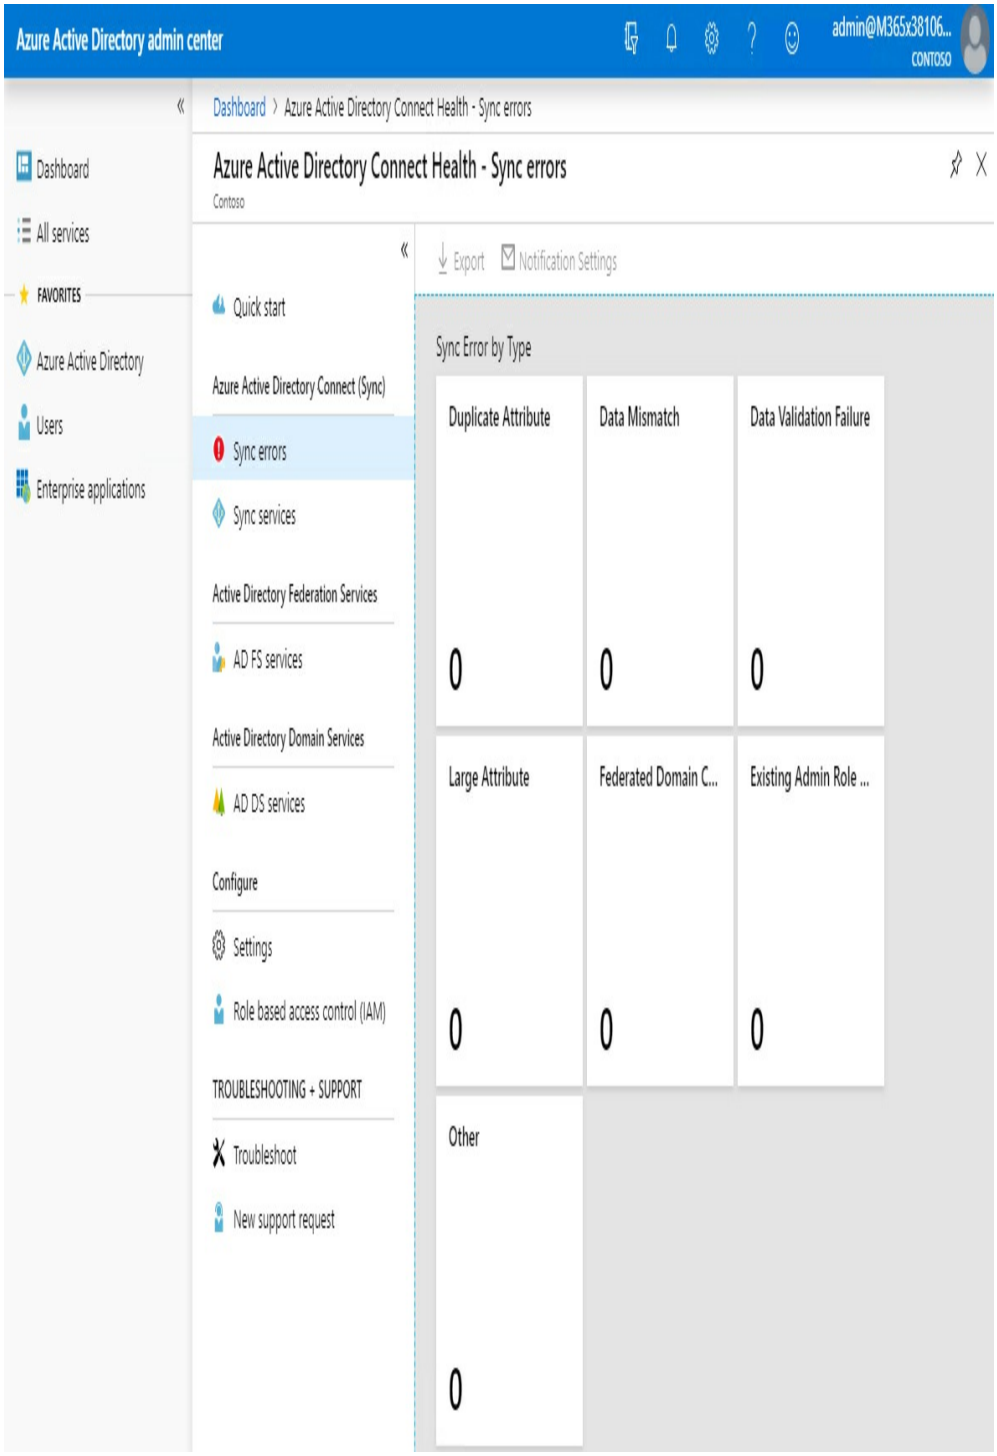


Figure 17-25 Azure AD Connect Health

You can use Azure AD Connect health to view information about:

- **Sync Errors.** This will display errors including Duplicate Attribute, Data Mismatch, Data Validation Failure, Large Attribute, Federated Domain Change, and Existing Admin Role Conflicts.
- **Sync Services.** This handles information about which services are synchronizing with Azure Active Directory.
- **AD FS Services.** This displays information about AD FS when Azure AD Connect is configured for federation. Includes information about errors and issues.
- **AD DS Services.** This displays information about domains and forests connected to Azure Active Directory.

More Info

Azure AD Connect Health

You can learn more about Azure AD Connect Health at: https://docs.microsoft.com/en-us/azure/active-directory/hybrid/whatis-azure-ad-connect.

Forcing synchronization

By default, synchronization occurs between the on-premises directory and Azure every 30 minutes. In some cases, you'll make a change to a user account or create a collection of user accounts and want to get those changes or new accounts up into the Azure Active Directory instance that supports the Office 365 tenancy as fast as possible. You can force synchronization by running the Azure AD Connect wizard again, or you can use the Synchronization Service Manager.

To perform a full synchronization using Synchronization Service Manager, perform the following steps:

1. Open the Synchronization Service Manager, either by clicking on Synchronization Service from the Start menu, or by running *miisclient.exe* located in the *C:\Program Files\Microsoft Azure AD Sync\UIShell* folder.
2. Click the Connectors tab.
3. On the Connectors tab, click the name of your Active Directory domain service, as shown in Figure 17-26.

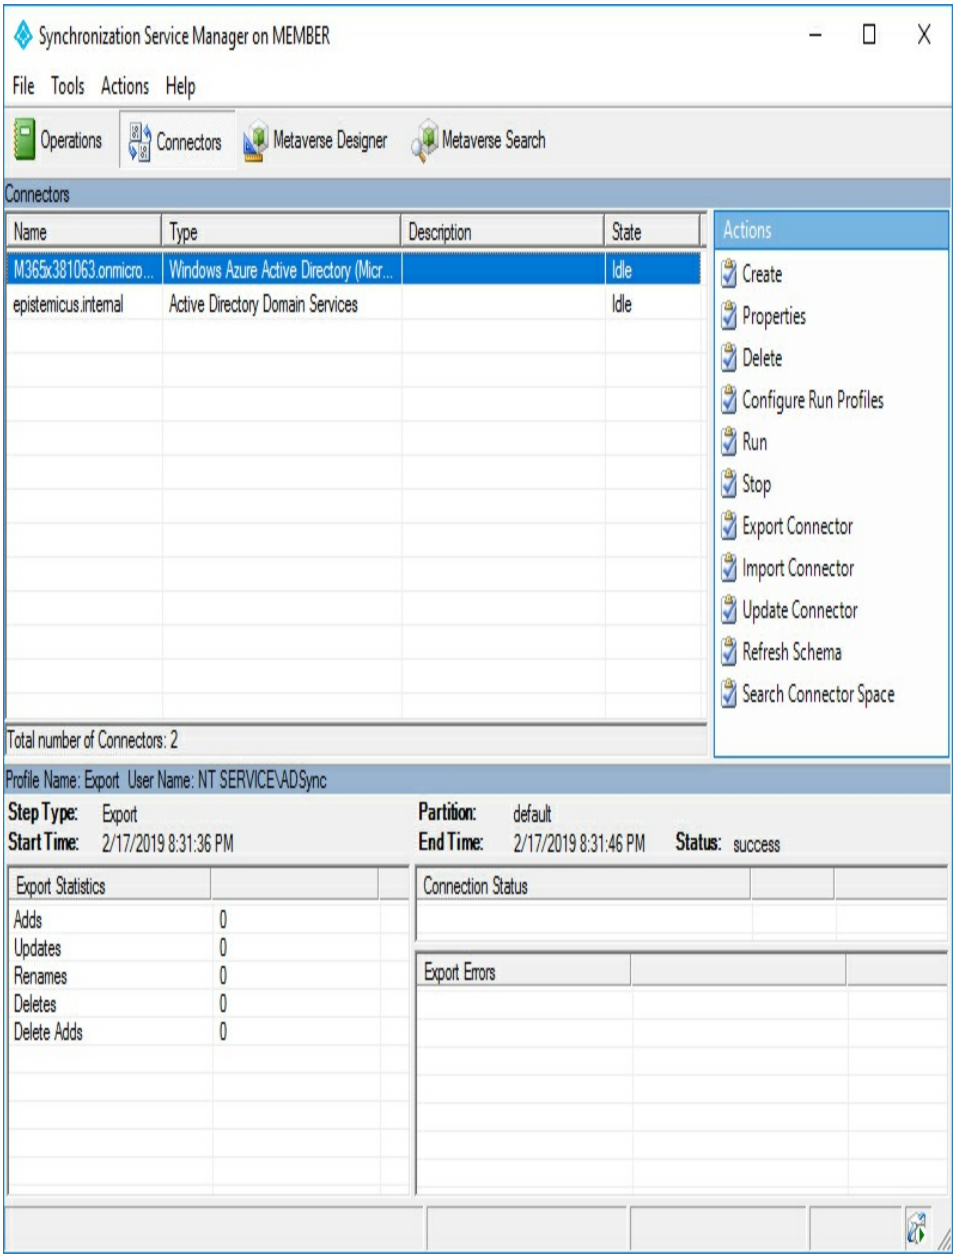


Figure 17-26 Synchronization Service Manager

4. On the Actions pane, click Run.
5. In the Run Connector dialog box, select Full Synchronization, as shown in Figure 17-27.

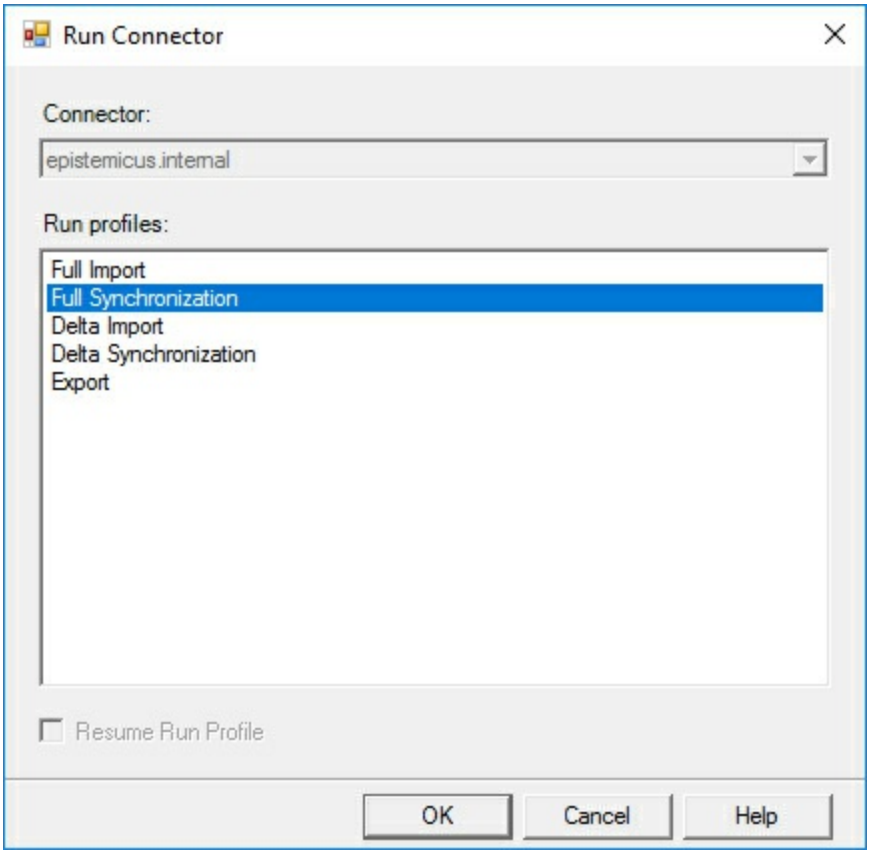


Figure 17-27 Choosing Full Synchronization

Rather than performing a Full Synchronization, you can trigger one of the following types of synchronization using the Synchronization Service Manager:

- **Full Synchronization.** Performs a full synchronization.
- **Delta Import.** Imports changed schema and objects.
- **Delta Synchronization.** Synchronizes only objects changed since the last sync.
- **Export.** Writes data from the Azure instance to the on-premises instance.
- **Full Import.** A full import and full sync is suitable for initiating the first full synchronization or the first full synchronization after

you have changed the filtering parameters.

You can also use the Synchronization Service Manager to configure extensive filtering options, though for tasks such as configuring OU-based filtering, Microsoft recommends that you first attempt configuring filtering using the Azure AD Connect setup wizard and only rely on a tool such as Synchronization Service Manager if problems arise.

> **More Info**
>
> **Synchronization Service Manager**
>
> You can learn more about Synchronization Service Manager at: https://docs.microsoft.com/en-us/azure/active-directory/connect/active-directory-aadconnectsync-service-manager-ui

You can also force synchronization by using the *Start-ADSyncCycle* cmdlet. You can use this cmdlet to trigger either a Delta Synchronization or a Full Synchronization. To force a delta sync cycle, run the following command:

Click here to view code image

```
Start-ADSyncSyncCycle -PolicyType Delta
```

To trigger a full sync cycle, run this command:

Click here to view code image

```
Start-ADSyncSyncCycle -PolicyType Initial
```

Configure object filters

When you use Azure AD Connect to synchronize on-premises Active Directory to an Azure Active Directory instance, the default setting is to have all user accounts, group accounts, and mail-enabled contact objects synchronized up to the cloud. For some organizations, synchronizing everything is exactly what they want. Other organizations want to be more selective about which objects are synchronized from the on-premises Active Directory environment to the Azure Active Directory instance that supports the Office 365 tenancy.

With Azure AD Connect, you can choose to filter based on the following options, as shown in Figure 17-28:

- **Domain based.** In a forest with multiple domains, you can configure filtering so that only objects from some domains (and not others) are filtered.
- **Organizational unit (OU) based.** With this filtering type, you choose which objects are filtered based on their location within specific organizational units.

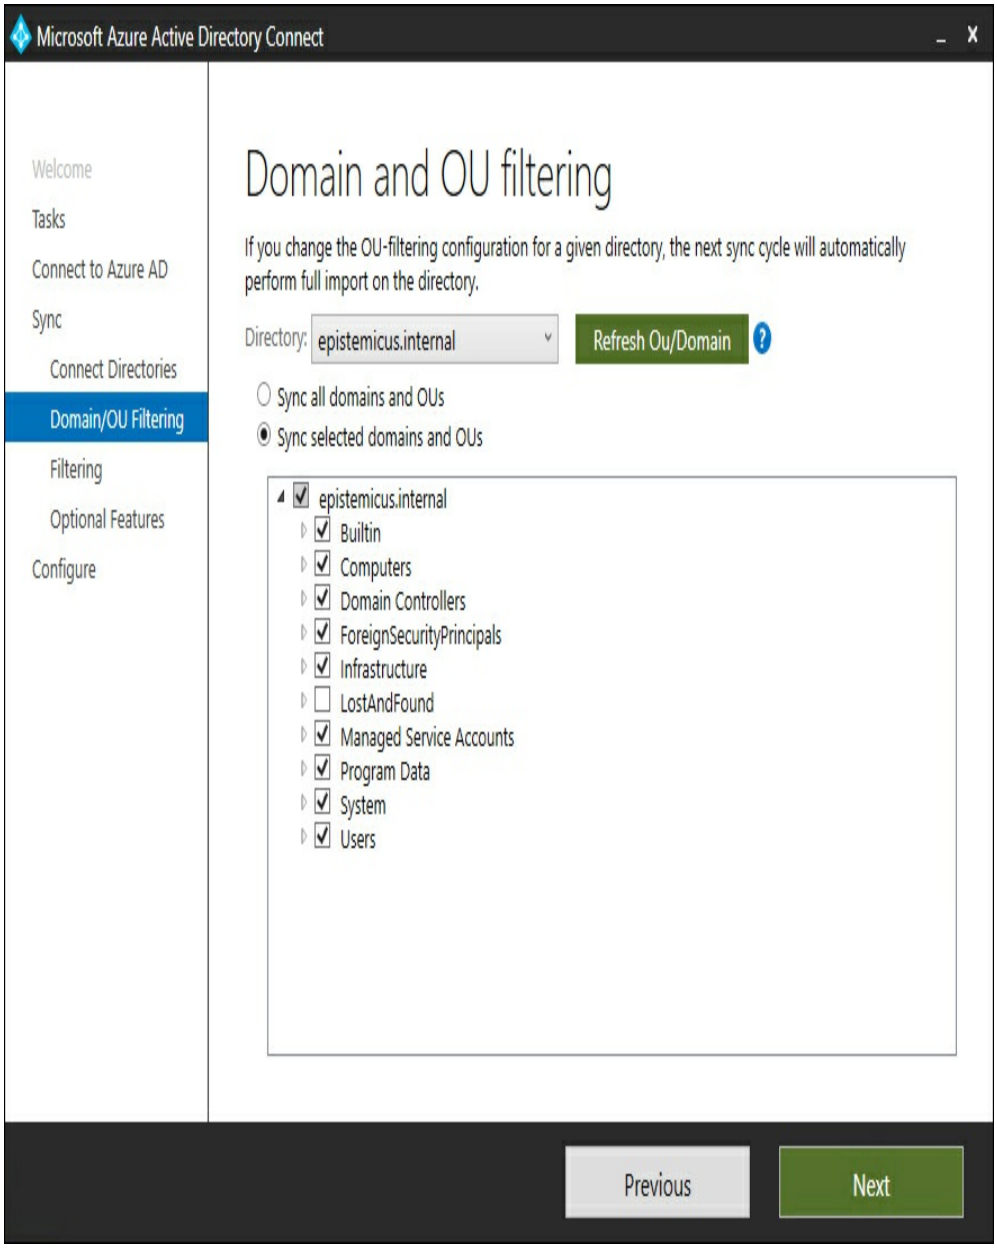


Figure 17-28 Domain And OU Filtering

You can also configure filtering on the basis of group membership. You can configure separate group-based filters for each forest, or you can configure domain-

synchronized filters using Azure AD Connect.

> **More Info**
>
> **Configure Filtering**
>
> You can learn more about Azure AD Sync filtering at https://docs.microsoft.com/en-us/azure/active-directory/connect/active-directory-aadconnectsync-configure-filtering.

While Azure AD Connect will address most organization's synchronization requirements, the most comprehensive tool that you can use to filter synchronization is the Synchronization Rules Editor, shown in Figure 17-29. You can use this tool to modify existing synchronization rules, but you also can create new rules. Rather than configuring synchronization on a per-domain or per-OU basis, you can tailor rules for individual objects and specific Active Directory attributes.

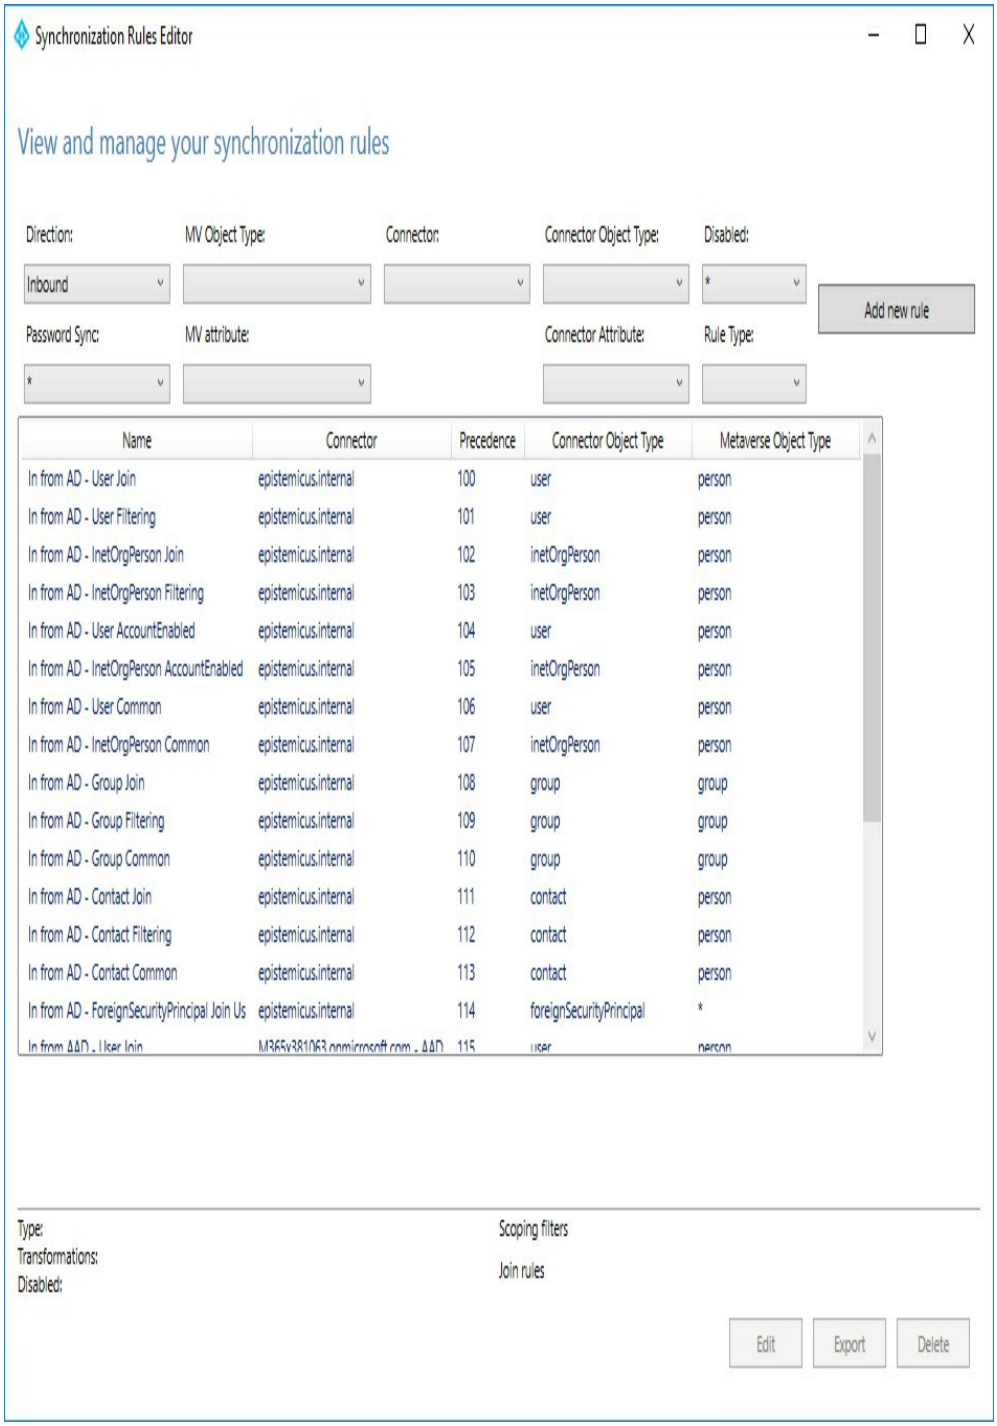


Figure 17-29 Synchronization Rules Editor

More Info

Synchronization Rules Editor

You can learn more about the Synchronization Rules Editor at https://docs.microsoft.com/en-us/azure/active-directory/connect/active-directory-aadconnectsync-change-the-configuration.

Implement and manage Azure AD self-service password reset

A self-service password reset is challenging to deploy in an on-premises environment, but it is relatively straightforward to deploy in an environment that uses Azure AD as a source of identity authority. A self-service password reset allows users to reset their own passwords when they forget them, rather than having to contact the service desk and have a member of the IT staff perform the password reset for them. Perform the following steps to enable self-service password reset:

1. Open the Azure Active Directory portal at *https://aad.portal.azure.com* with an account that has tenant administrator permissions.
2. In the Azure Active Directory Admin Center, click the Users node, which will open the Users blade.
3. In the Users blade of the Azure Active Directory admin center, click Password Reset.
4. On the Password Reset – Properties page, click All, as shown in Figure 17-30, to enable the self-service password reset for all users whose accounts are synchronized from your on-premises Active Directory instance to Azure AD.

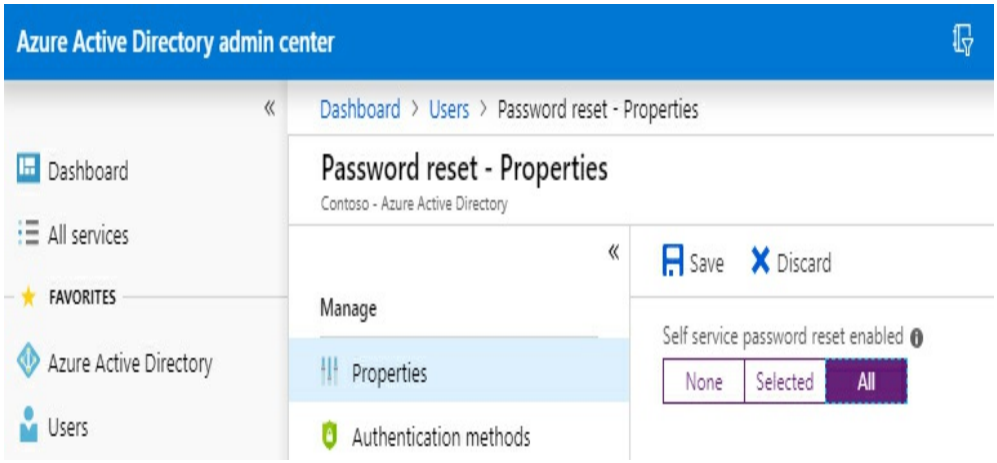


Figure 17-30 Enable self-service password reset

Once enabled, users will be prompted for additional information the next time they sign in; the password will be used to verify their identity if they use the Self Service Password Reset tool. Users can reset their passwords by navigating to the website *https://passwordreset.microsoftonline.com*, shown in Figure 17-31, and completing the form.

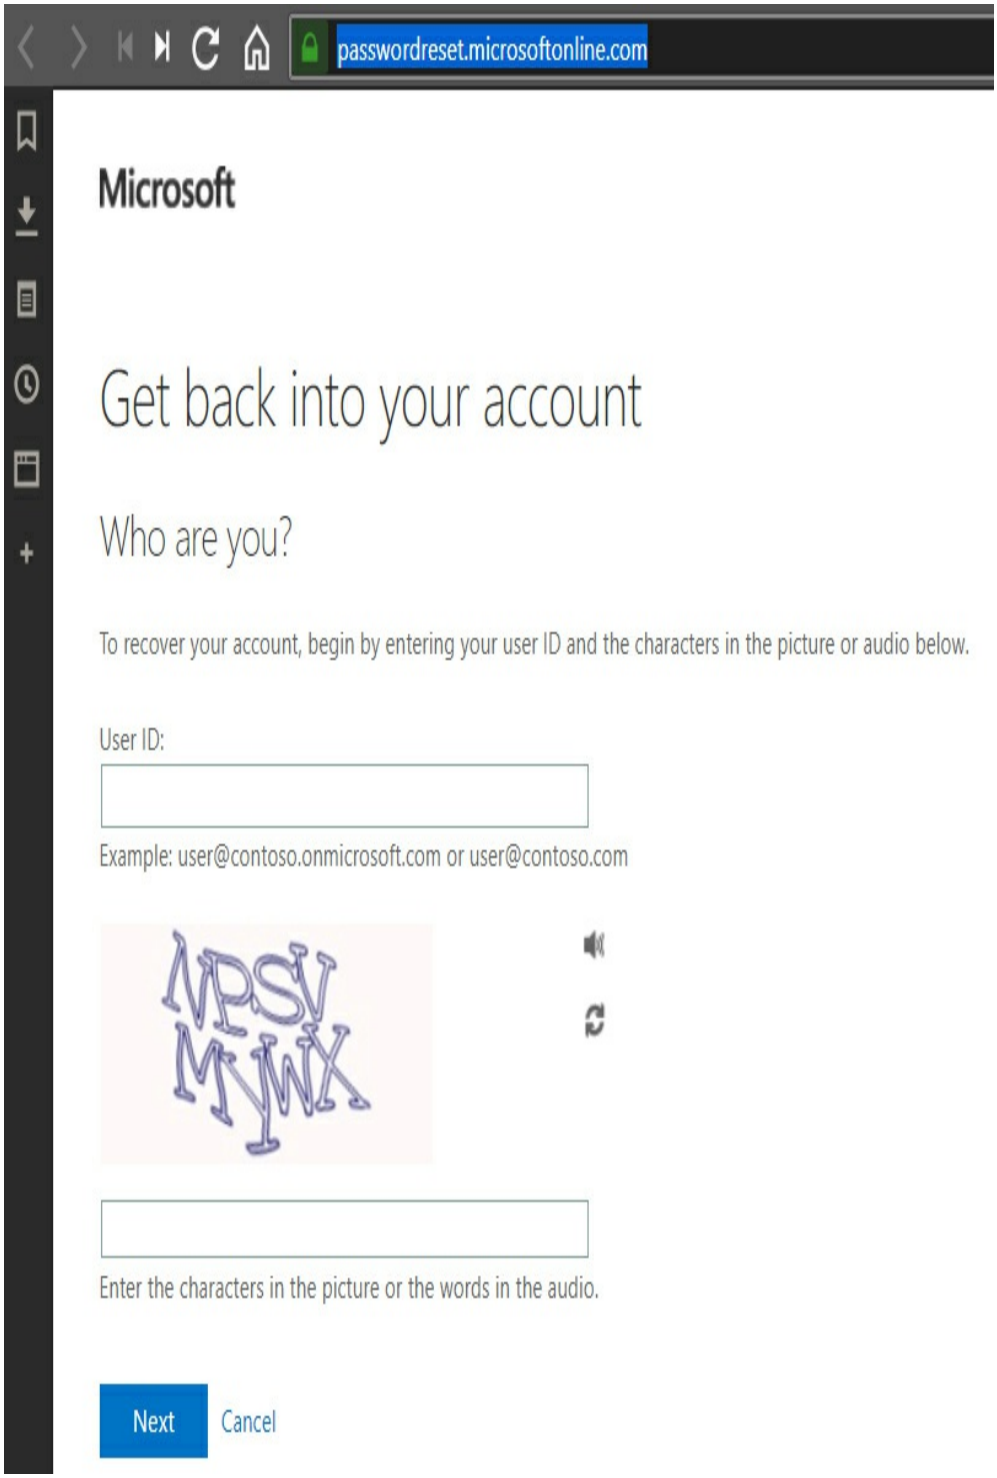


Figure 17-31 Self-service password reset website

More Info

Self-Service Password Reset

You can learn more about configuring Self-Service Password Reset at https://docs.microsoft.com/en-us/office365/admin/add-users/let-users-reset-passwords.

Azure AD Password Protection

Azure AD Password Protection allows you to block commonly used passwords as well as your own custom password list from being used in your on-premises Active Directory environment. Deploying Azure AD Password Protection involves deploying the Azure AD Password Protection Proxy Service on computers in your environment that are located on your organization's perimeter network and the Azure AD Password Protection DC agent on each domain controller in your environment. The proxy service allows you to ensure that your domain controllers are not directly communicating with the Internet.

Microsoft recommends that you initially deploy Azure AD Password Protection in audit mode. This will allow you to determine how many users are currently using insecure passwords and will also allow you to notify users about when you will be implementing the solution to block the use of commonly known passwords.

To deploy Azure AD Password Protection, you need the following:

- You need an account that has Windows Active Directory domain administrator privileges in the forest root domain of your Windows Active Directory environment. This will be used to register the forest with Azure AD.
- You need an account that has Global Administrator privileges for your organization's Azure tenant. This account will be used to register the Azure AD Password Protection Proxy service with Azure AD.
- You must enable the Key Distribution Service on all domain controllers.
- You must have network connectivity between at least one domain controller in each domain in the forest and a server that hosts the Azure AD Password Protection Proxy service.
- The computers that host the Azure AD Password Protection Proxy service must be able to access the following endpoints on the Internet:
    - *https://login.microsoftonline.com*
    - *https://enterpriseregistration.windows.net*
- You must deploy the Azure AD Password Protection Proxy service prior to installing the agents on domain controllers. Register the proxy service against Azure Active Directory using the *Register-AzureADPasswordProtectionProxy* cmdlet. Then register the on-premises Active Directory forest with the *Register-AzureADPasswordProtectionForest* PowerShell cmdlet.
- You must deploy the Windows Server Active Directory Domain Controller agent on each domain controller in your organization. The agent must be deployed on each domain controller because you cannot predict which domain controller in your organization will be used to process a password change.

Once the agents are installed, you switch on on-premises password protection by performing the

following steps:

1. In the Azure portal, open Azure Active Directory.
2. In the Manage section, click Security > Authentication Methods> Password Protection.
3. Select the Yes option under Enable Password Protection On Windows Server Active Directory. You'll only be able to do this once you've registered the proxy service with Azure AD.
4. You can then use this console to switch between Audit mode, which will report on a problematic password, and Enforced mode, which will block users from setting a problematic password.

> **More Info**
>
> **Azure AD Password Protection**
>
> You can learn more about Azure AD Password Protection at https://docs.microsoft.com/en-us/azure/active-directory/authentication/howto-password-ban-bad-on-premises-deploy.

Azure AD DS

A simple option for domain-joining an IaaS VM is to deploy an IaaS domain controller VM and to then join other VMs to the domain in the same way that you would in an on-premises environment. The challenge with this is that it means that you have to deploy, maintain, and pay for an IaaS VM configured as a domain controller. Or you can deploy, maintain, and pay for two IaaS VMs if you're following the best practice of having at least two domain controllers for any single domain.

Azure Active Directory Domain Services (AD DS)

provides a managed domain service for IaaS VMs running in Azure. This service provides domain join, Group Policy, lightweight directory access protocol (LDAP), and Kerberos and NTLM authentication, and it is compatible with Windows Server Active Directory. Because Azure AD DS is a managed service, Microsoft takes care of the management of the back-end domain controller infrastructure.

Azure AD DS pulls identity data from Azure AD. This includes identities synchronized from an on-premises Windows Active Directory instance through Azure AD Connect. When you deploy Azure AD DS, you select an Azure virtual network on which to make the service available. IaaS VMs placed on that virtual network can then domain join to Azure AD DS in the same manner as a computer running on a traditional network with a domain controller would.

More Info

Azure AD DS

You can learn more about Azure AD DS at https://docs.microsoft.com/en-us/azure/active-directory-domain-services/.

You can only perform a domain join to an Azure AD DS instance with an account that is part of the Azure AD tenant. You can't perform a domain join using an account that has been synchronized from an on-

premises Windows Active Directory instance to perform this task. As is the case when configuring Azure AD Connect, you should consider creating a special account using the default tenancy *onmicrosoft.com* suffix for performing domain-join operations rather than any custom domain name that you have assigned to the tenancy.

You perform the domain-join operation from within the Windows Server VM, either directly by connecting through an RDP or Azure Bastion session or through a remote PowerShell session.

More Info

IaaS VM domain join

You can learn more about domain joining a Windows Server IaaS VM at https://docs.microsoft.com/en-us/azure/active-directory-domain-services/join-windows-vm.

AZURE HYBRID CLOUD SERVICES

Microsoft is increasingly moving the products that you use for managing and securing your on-premises environment to Azure. While it's still possible to manage and monitor a Windows Server deployment using only on-premises products, Microsoft is moving more in the direction of monitoring, automation orchestration, and security being something that you deal with through cloud-based products. This means

functionality for which you once needed products such as System Center Operations Manager, System Center Data Protection Manager, and System Center Orchestrator will increasingly be available within Azure. This doesn't mean those products are going away any time soon; it just means that at some point in the future, you'll be able to achieve the same things, if not more, by leveraging a cloud service to perform those tasks rather than deploying on-premises workloads.

There are several advantages to having the management layer in the cloud rather than hosted in an on-premises datacenter. The cloud:

- Gives you almost limitless capacity to store monitoring telemetry. This gives you the option of ingesting almost all possible telemetry for later analysis. The more you can record about your on-premises environment, the more you can learn about what's going on in your on-premises environment.
- Allows you take advantage of machine learning and AI technologies to gain insights from the telemetry you have ingested.
- Allows you to use resources, such as Microsoft Graph, to take advantage of what Microsoft has discovered about certain patterns of telemetry. For example, Azure Sentinel can identify novel attack chains under exploitation by malicious actors in your environment once the characteristics of those attack changes have been recognized elsewhere.
- Is constantly updated with new functionality. Unlike on-premises monitoring solutions that infrequently add new functionality, cloud-based monitoring and management technologies add new functionality at a regular cadence. This includes adding an increasing library of workbooks

Connect Windows Admin Center

Before you can integrate Azure hybrid services with Windows Server 2019, you need to register the Windows Admin Center gateway with your Azure subscription. To connect Windows Admin Center to your Azure subscription, perform the following steps:

1. In Windows Admin Center, click Azure Hybrid Services and then click Register Your Windows Admin Center Gateway To Get Started.
2. You'll be presented with a device code, as shown in Figure 17-32. Copy the code. Click the Device Login link and then past the code in.



Figure 17-32 Register WAC with Azure

3. On the login page, paste the code. If you don't use this code within a certain amount of time, you'll need to get a new code.
4. You'll then be prompted to provide the credentials of an account that has permissions to grant an app access to the Azure AD tenancy associated with the subscription. Once you've provided those credentials and authenticated, you'll need to select a tenancy and click Register. You can also select an existing app that is registered with Azure AD if

this isn't the first instance of Windows Admin Center that you have registered with the subscription.

5. To register the Azure AD app, click the link to register the Azure AD app, click Settings > Required Permissions > Grant Permissions, as shown in Figure 17-33, and then click Yes.

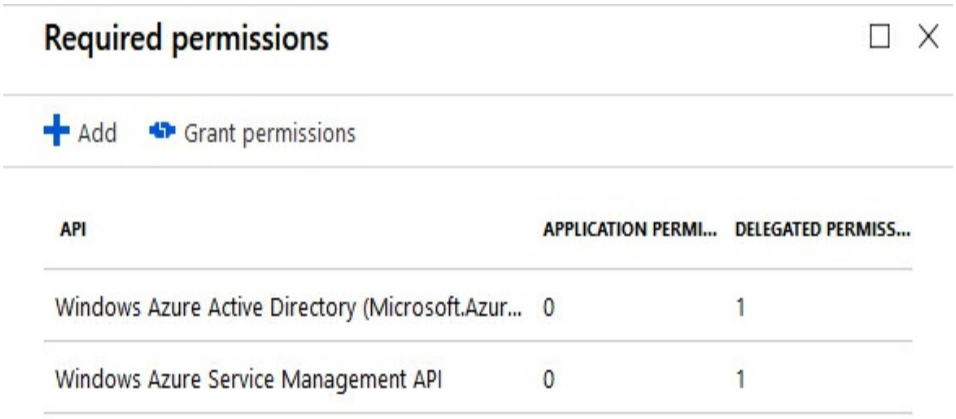


Figure 17-33 Grant WAC permissions

6. Once you have granted the permissions in Azure Admin Center, return to Windows Admin Center and click Continue. The notification section of Windows Admin Center will confirm that the application is successfully registered, as shown in Figure 17-34.

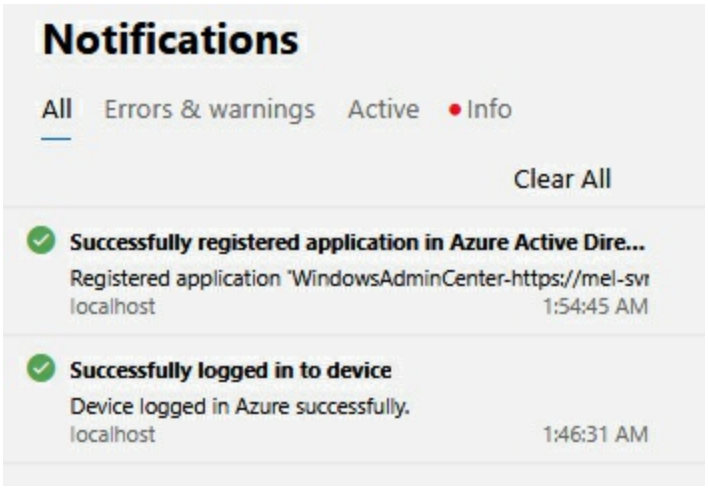


Figure 17-34 Successful registration

7. Once you receive this notification, you will be able to manage Azure

hybrid services using Windows Admin Center.

> **More Info**
>
> **Azure hybrid services**
>
> You can learn more about Azure hybrid services at https://docs.microsoft.com/en-us/windows-server/manage/windows-admin-center/azure/.

Creating Azure IaaS VMs from Windows Admin Center

Once you have configured hybrid services, you'll even be able to create VMs in Azure from your Windows Admin Center console. To do this, perform the following steps:

1. In Windows Admin Center, go to All Connections, click Add, and then scroll to select Azure VM, as shown in Figure 17-35. Click Create New.

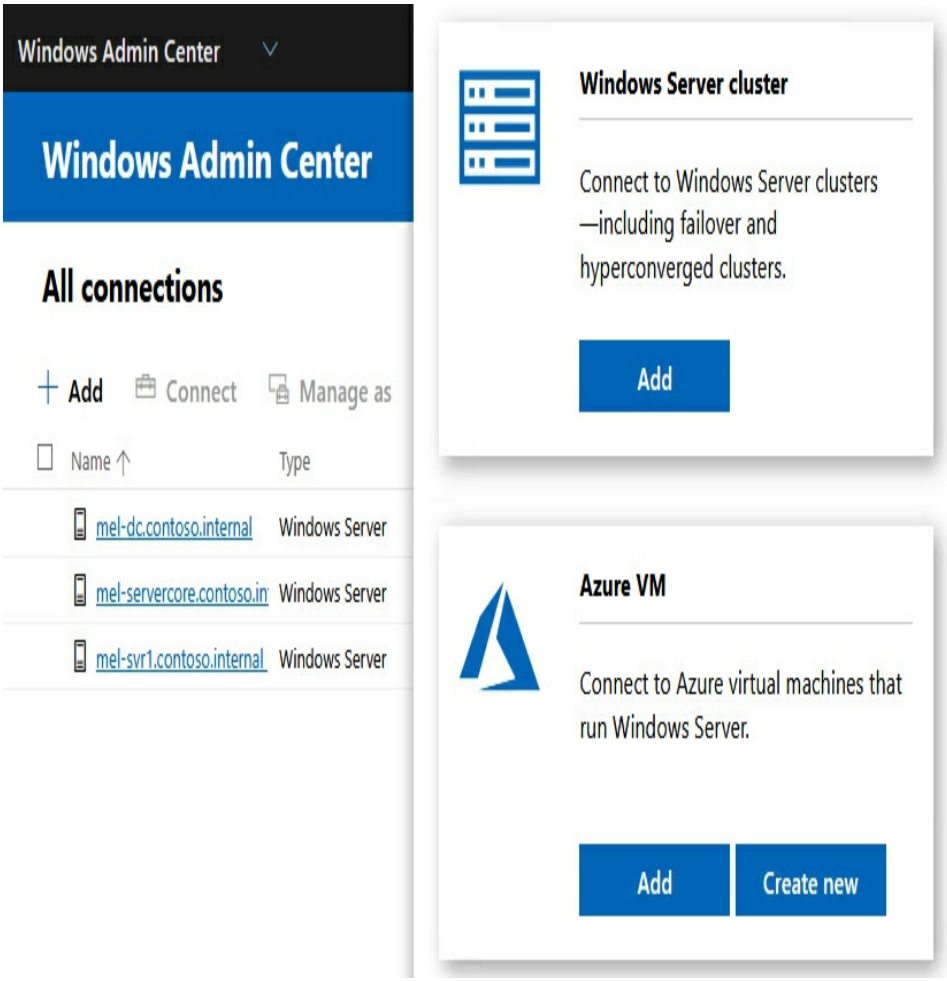


Figure 17-35 Use WAC to create Azure VM

2. On the Create An Azure VM page, shown in Figure 17-36, configure the following information and click Next:

    1. **Subscription.** Select which subscription the VM will be associated with.
    2. **Resource Group.** Specify which resource group the VM will be associated with.
    3. **VM Name.** Provide a name for the VM.
    4. **Region.** Specify which Azure region to which you wish to deploy the VM.

5. **Operating System Image.** Choose which operating system to deploy. Deployments through Windows Admin Center allow you to deploy all currently supported Windows Server images. At the time of writing, you couldn't use Windows Admin Center to deploy a Linux image.
6. **Username.** This is the username of the local administrator account. Remember that you can't use a name like Administrator or Root.

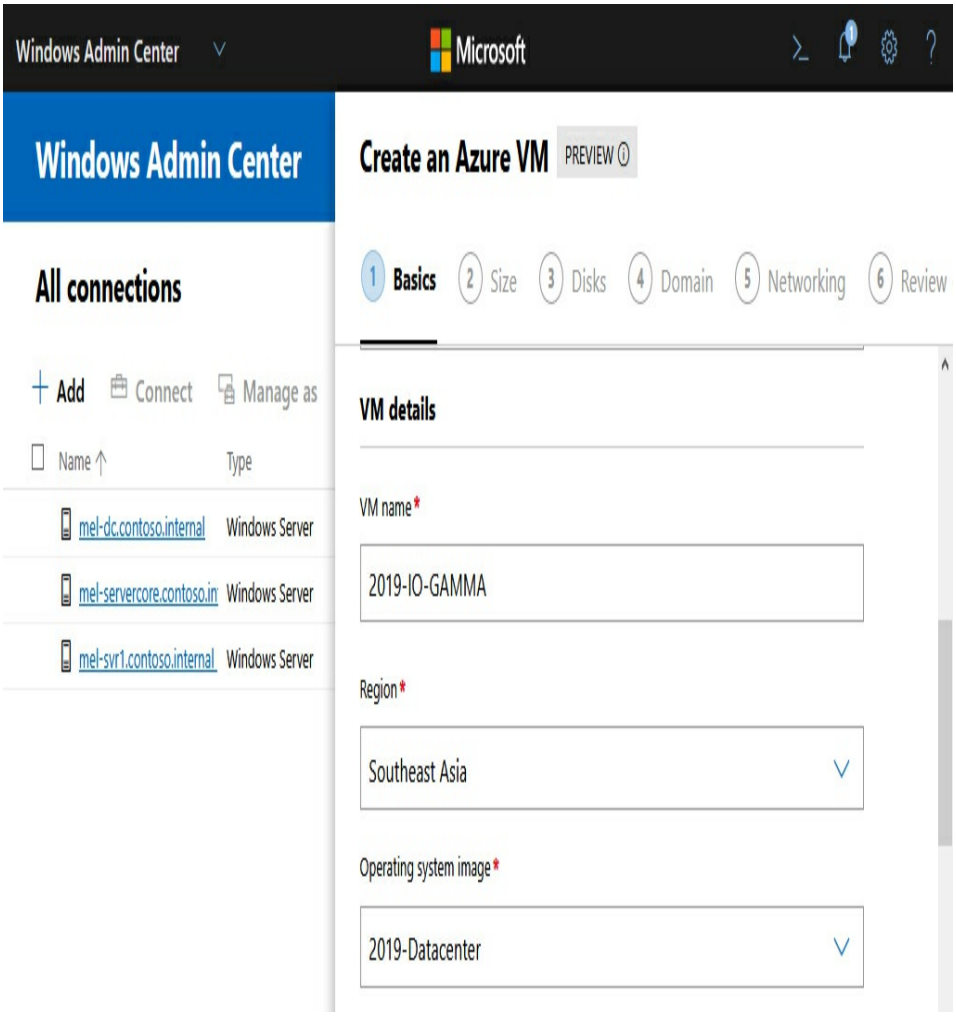


Figure 17-36 New IaaS VM through WAC

3. On the Size page, select which size the deployed VM should be. The page, shown in Figure 17-37, lists the VCPU, RAM, Data Disks, and

Technet24

Premium Disk Support characteristics of each VM size.

4. On the Disks page, select whether you want the IaaS VM to leverage Premium SSD and whether you want to connect any data disks to the VM.
5. On the Domain page, select whether you want the IaaS VM to be domain joined.
6. On the Networking page, specify which virtual network and subnet to which you wish to connect the VM.
7. On the final page, review the settings and create the VM.

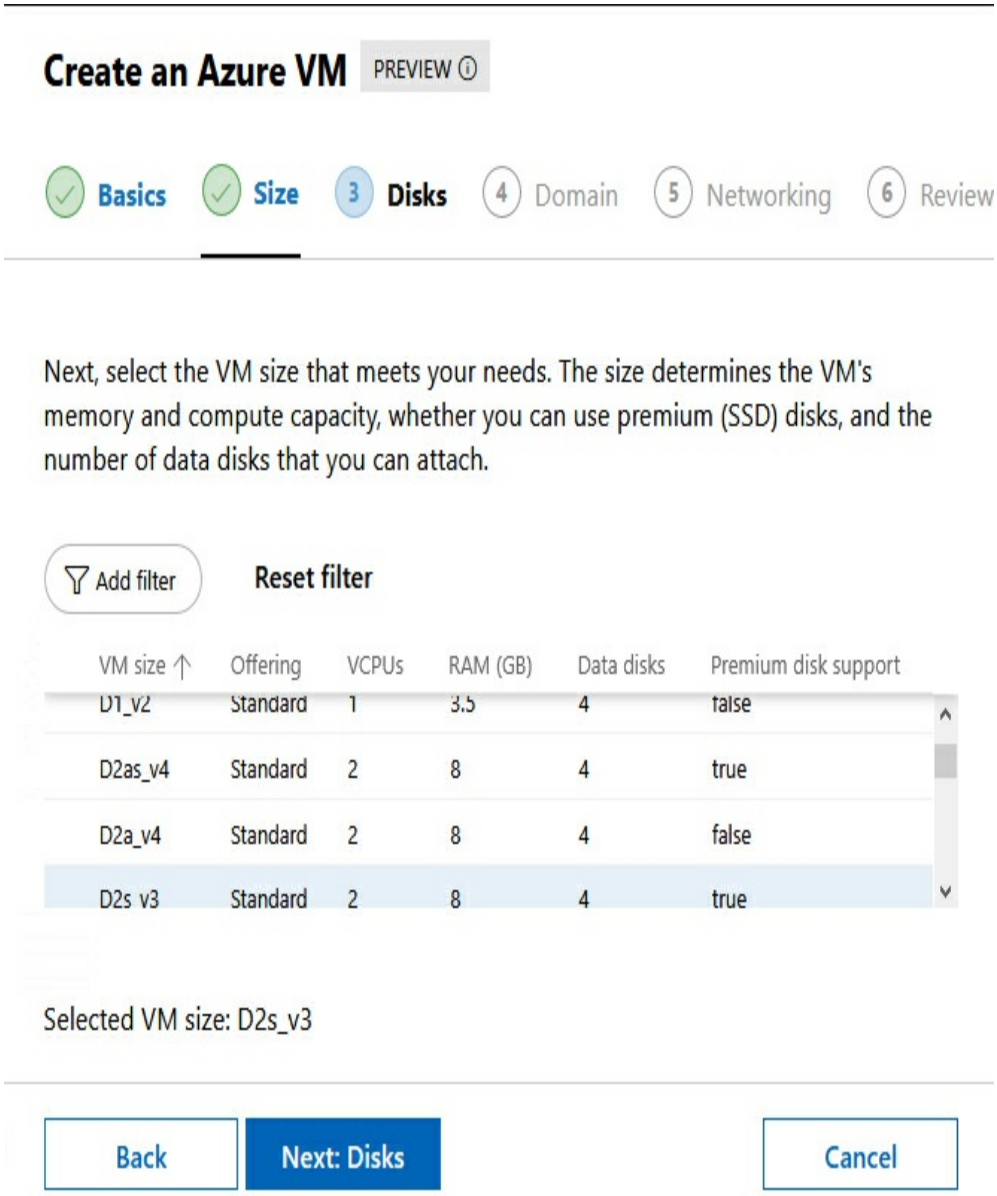


Figure 17-37 Set Azure IaaS VM size

Azure File Sync

Azure File Sync is one of the most useful hybrid technologies available for Windows Server. Anyone who has managed a file server knows the challenges they

pose, from having to remove disused files on a regular basis to ensure that there is enough space for new files, to challenges around ensuring that files are regularly backed up, and even to being able to restore a file that you might have removed at some point to save space because someone actually needs it now.

Important concepts related to Azure File Sync are:

- Azure File Sync Endpoint
- Cloud Tiering
- Azure File Share
- Azure File Sync Replication

Azure File Sync Endpoint

An Azure File Sync endpoint consists of a server and a path that are enrolled in an Azure File Sync Service. A Windows Server can host multiple endpoints, each of which has a different path. An Azure File Sync endpoint functions as the folder structure that underlies a normal file share. Administrators create a traditional shared folder and point it at the path that the Azure File Sync endpoint replicates to. You can also point a Distributed File System (DFS) namespace at this path, replacing DFS replication with Azure File Sync replication while still keeping the navigational advantages of the DFS way of identifying shared folders.

Cloud tiering

Azure File Sync uses a process called cloud tiering to ensure that there is capacity on the volume that hosts the share. Cloud tiering means that you don't need to worry about constantly freeing up space for new files. You can configure Azure File Sync on a per-file-share basis to tier files based on when the file was last accessed, how much free space there is on the volume that hosts the share, or both. For example, you might configure Azure File Sync so that any file that hasn't been accessed in 14 days on a particular share is automatically tiered to Azure. You could also configure Azure File Sync so that the least recently accessed files are automatically tiered to Azure in the event that the volume only has 30 percent free space remaining.

From the users' perspective, a tiered file still appears as though it's on the file server that they are accessing. If users try to open the file, it syncs down from the back-end Azure File Share to the Azure File Sync file share endpoint and then opens normally.

Azure File Share

The backend of Azure File Sync is an Azure File Share, also known as a cloud endpoint. This is a cloud file share of essentially limitless capacity that will store any file that is written to an Azure File Sync endpoint. The back-end Azure File Share stores the entire contents of what appears to be on the file share that is the front end

for Azure File Sync. You can configure Azure Backup to then back up this endpoint. The advantage of this is in the event of data corruption or deletion, you can just recover data to the Azure File Share in the cloud from the Azure console, and it will replicate down to all the Azure File Sync endpoints.

Azure File Sync replication

Azure File Sync replicates files back and forth between the on-premises endpoint and the Azure File Share backend. You can have multiple endpoints connected to the same Azure File Share backend. This collection of endpoints and Azure File Share backend is called a sync group. The Azure File Sync replication engine allows you to have the same file share content in multiple locations, and you can replace most DFS implementations with Azure File Sync. If two people have the same file open and click Save at nearly the same time, the most recently written change will keep the original file name with the older file having the source server name and a conflict number appended to the name.

Deploying Azure File Sync

To deploy Azure File Sync, perform the following steps:

1. In the Azure portal, create a storage sync service. You should plan for the storage sync service and the Azure File Share endpoint to be in the same region, and it is common to place these in the same resource

group.

2. Install the Azure File Sync agent on the server that currently hosts the files that you want to have replicated by Azure File Sync. Register this server with the storage sync service that you created earlier.
3. Create an Azure File Share that will function as the cloud endpoint.
4. Create a sync group. The sync group needs a name—the name of the storage account it will use and the name of the Azure File Share that it will use.
5. Create a server endpoint by adding the server that you registered with the storage sync service and specifying the local path to the files that you want to replicate using Azure File Sync. When creating the server endpoint, you also specify the cloud tiering settings in terms of how much free space should always be available on the local volume that hosts the files and how many days after a file was last accessed should pass before the file is tiered. After the endpoint is created, any files in the path specified will be replicated up to the Azure File Share that functions as the cloud endpoint.
6. Register additional servers with the storage sync service, add them to the sync group, and then configure them as server endpoints. The files that are stored in the cloud endpoint will then start replicating down to the newly added server endpoints.

You'll need to create file shares that point to the endpoint paths on each new file server that you add to the sync group. File shares that serve as frontends for Azure File Sync endpoints should have the same share permissions. You should also consider setting up Azure Backup to back up the cloud endpoint on a daily basis, so that you can recover any data that might accidentally become deleted or corrupted.

More Info

Azure File Sync

You can learn more about Azure File Sync at https://docs.microsoft.com/en-us/azure/storage/files/storage-sync-files-deployment-guide?tabs=azure-portal.

Azure Arc

At present, Azure Arc allows you to perform a limited number of management tasks on enrolled computers in your organization's in-premises or cloud environment. Over time, the number of tasks that you can perform with Azure Arc will increase, but as of early 2020, a server enrolled in Azure Arc can have its event viewer logs viewed through the Azure console, and it is possible to assess Azure policies against enrolled endpoints. For example, you could use an Azure policy to assess a server's security settings against a compliance standard such as ISO 27001. Eventually, you'll be able to use Azure Arc to remediate security configuration settings that don't meet compliance benchmark requirements.

More Info

Azure Arc for Servers

You can learn more about Azure Arc for Servers at https://docs.microsoft.com/en-us/azure/azure-arc/servers/overview.

Azure Site Recovery

Azure Site Recovery allows you to create replica VMs of on-premises workloads in Azure. Azure Site Recovery allows you to replace a disaster recovery (DR) site that

would be used in the event that a primary site failed. Instead of failing over to the DR site, you can fail your virtual machines over to Azure.

Performing failover with Azure Site Recovery doesn't mean that you are stuck with your workloads in Azure. You can fail workloads back to your on-premises infrastructure once that infrastructure has been returned to an orderly state. Workloads are replicated and stored in Azure storage, with VMs only being provisioned within Azure when you perform failover. You can perform test failovers to check that your disaster recovery plan is functioning properly. When you do this, temporary workloads are created in Azure that allow you to assess workload and disaster recovery plan integrity without performing an actual failover to Azure.

Azure Site Recovery supports application-consistent snapshots, and replication frequency can be as high as every 30 seconds for workloads hosted on an on-premises Hyper-V server. You can create specialized recovery plans that orchestrate the failover of multiple VMs in multi-tier applications.

Azure Site Recovery works with virtual and physical workloads running Windows and Linux operating systems. You can protect workloads running on Hyper-V, VMware, and Azure Stack.

More Info

Azure Site Recovery

You can learn more about Azure Site Recovery at https://docs.microsoft.com/en-us/azure/site-recovery/site-recovery-overview.

Azure Network Adapter

Azure Network Adaptor allows you to create a point-to-site VPN connection to an Azure virtual network directly from Windows Admin Center on a computer running Windows Server 2019. Because it is a point-to-site VPN, the computer that you're configuring with Azure Network Adapter only needs a connection to the Internet; it does not need to have a public IP address.

While it's possible to create point-to-site VPN connections from the Azure console, as is the case with creating an Azure VM from Windows Admin Center, this allows you to perform all the necessary tasks from the primary tool that you use to manage Windows Server. Windows Admin Center will also assist in the creation of Azure VPN gateways if one isn't already present on the virtual network to which you are creating a connection.

Once you have configured Azure Network Adapter, you'll be able to connect to and manage IaaS VMs running Windows Server operating systems connected to the virtual network the adapter connects to. You can configure multiple Azure virtual network adapters, each

of which connects to a different virtual network.

More Info

Azure Network Adapter

You can learn more about Azure Network Adapter at https://aka.ms/WACNetworkAdapter.

Chapter 18

Windows Subsystem for Linux

Linux on Windows Server

Installing WSL

WSL 2.0

Windows Subsystem for Linux (WSL) allows you to run a variety of Linux distributions on Windows computers. Windows Subsystem for Linux has been available for Windows 10 for several years, and distributions are available from the Microsoft Store. You can even use X-Windows server on Windows 10 to run Linux GUI applications.

Though you can't use the Microsoft Store on Windows Server 2019, WSL is available and relatively easy to install. In this chapter, you'll learn how to set up Windows Subsystem for Linux on Windows Server and

how to install, configure, and manage Linux distributions.

LINUX ON WINDOWS SERVER

Linux has been available on Windows Server through a variety of methods over the years, including Windows Services for UNIX that were available with Windows NT and Windows 2000 and third-party products, such as Cygwin; also, you can run Linux or BSD virtual machine on Hyper-V.

When it comes to running Linux utilities or hosting Linux workloads, Windows Server provides multiple options.

- If you need access to Linux command-line utilities on a Windows Server computer, WSL provides you access to almost all utilities available in the most popular Linux distributions. You install and manage them in the same way you'd install or manage them from the command line on a Linux host.
- If you need to host Linux container workloads, you can use Linux Containers On Windows (LCOW). You learned about LCOW in Chapter 10, "Containers." In small- to medium-sized environments, it may be simpler to host Linux container workloads on a Windows Server instance than to manage both Windows and Linux servers separately. At present, you can't host Windows container workloads on Linux computers.
- If you need to host more complicated Linux workloads, many of the most popular Linux distributions are fully supported on Hyper-V. Also, you can integrate a Hyper-V IaaS deployment with an Azure IaaS deployment, shifting virtual machines running Linux to and from the cloud as required.

Technet24

Although it is possible to run workloads such as the NGINX web server through WSL on a Windows Server host, you'd be far better off running such a workload using a container or virtual machine. WSL works best as a way of accessing additional command-line tools. If you want to host Linux-based workloads on Windows Server, Linux virtual machines and containers are far more robust solutions.

INSTALLING WSL

The first step involved in installing WSL on Windows Server 2019 is to enable the Windows Subsystem for Linux optional feature. You can do this by running the following command and then restarting the server:

Click here to view code image

```
Enable-WindowsOptionalFeature -Online -FeatureName
Microsoft-Windows-Subsystem-Linux
```

Unlike Windows 10, where you can install a WSL distribution directly from the store, WSL on Windows Server requires that you download a prepared WSL distro. WSL on Windows Server 2019 supports the following Linux distributions:

- Ubuntu 18.04 and Ubuntu 16.04
- Debian GNU/Linux
- Kali Linux
- OpenSUSE Leap 42

- SUSE Linux Enterprise Server 12
- Fedora Remix for WSL

You can download these distributions from the locations listed in Table 18-1.

Table 18-1 Obtaining WSL distributions

Distribution	Location
Ubuntu 18.04	*https://aka.ms/wsl-ubuntu-1804*
Ubuntu 16.04	*https://aka.ms/wsl-ubuntu-1604*
Debian GNU/Linux	*https://aka.ms/wsl-debian-gnulinux*
Kali Linux	*https://aka.ms/wsl-kali-linux-new*
OpenSUSE Leap 42	*https://aka.ms/wsl-opensuse-42*
SUSE Linux Enterprise Server 12	*https://aka.ms/wsl-sles-12*
Fedora Remix for WSL	*https://github.com/WhitewaterFoundry/WSLFedoraRemix/releases/*

You can download distributions from the addresses listed in Table 18-1 using *curl.exe*, which is included with Windows Server 2019. For example, to download a distribution to Ubuntu 18.04, use the following command:

Technet24

Click here to view code image

```
Curl.exe -L -o Ubuntu.zip https://aka.ms/wsl-ubuntu-1804
```

This will download a compressed version of the distribution to the *Ubuntu.zip* file. Once the distribution archive has been downloaded, extract it to the directory that you wish to have host the distribution files using the *Expand-Archive* PowerShell cmdlet. For example, to extract *Unbuntu.zip*, run this PowerShell command:

Click here to view code image

```
Expand-Archive .\Ubuntu.zip .\Ubuntu
```

It's a good idea to extract the distribution files in an obvious place, such as a folder named *C:\Distros*, as shown in Figure 18-1.

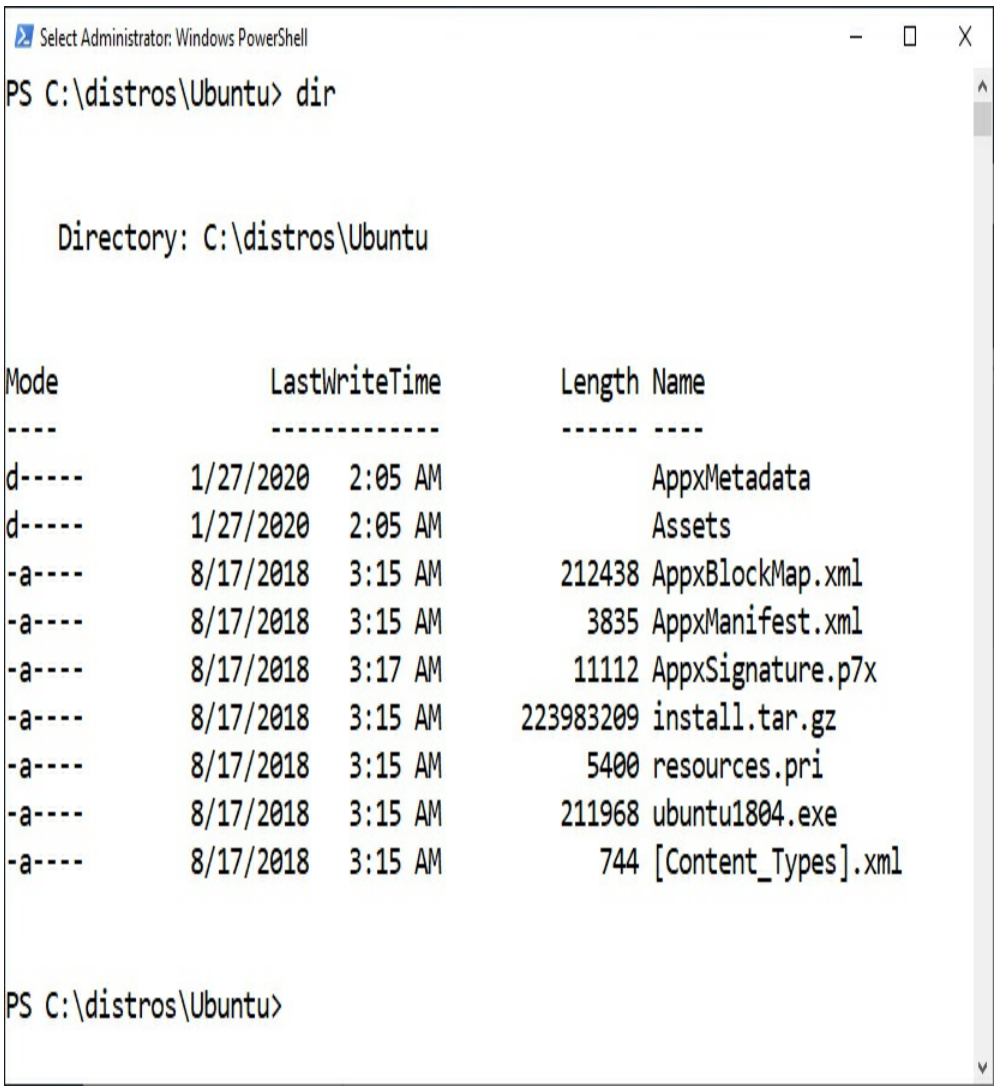


Figure 18-1 Extracted distribution

Once the archive has been unzipped, change into the distribution directory and run the distribution launcher application. This will be named <*distro*>.*exe*. Running this executable will set up the distribution. You'll need to set a new UNIX username and password as a part of this process. Once you've signed in, you'll be able to

Technet24

perform typical Linux administration activities. Figure 18-2 shows the output of the *lsb_release -a* command in Windows Subsystem for Linux.

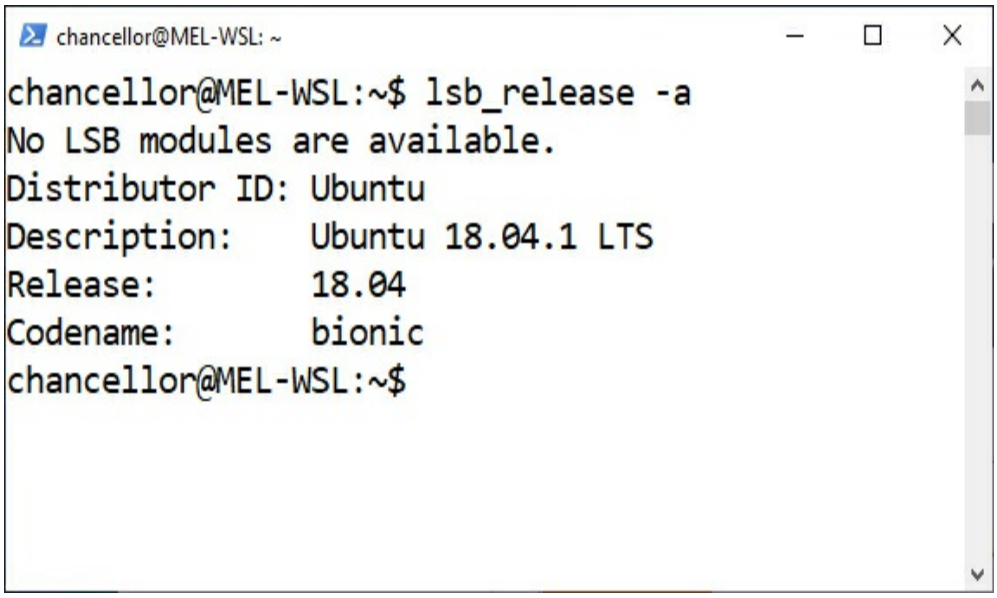


Figure 18-2 Running Ubuntu on Windows Server

You can install multiple Linux distributions on the same Windows Server computer as long as the distributions are in separate directories. You can also run different distros concurrently.

More Info

Updated WSL distributions

You can find the most up-to-date list of WSL distributions at https://docs.microsoft.com/en-us/windows/wsl/install-manual.

Inside OUT

Linux containers on Windows

Chapter 10, "Containers," provided information on how to run Linux Containers on Windows (LCOW). LCOW works with both Windows 10 and Windows Server 2019, but you must use Docker in Experimental mode to do this.

When you run a WSL distribution, the Windows file system will be located in the */mnt/<drive letter>* folder, as shown in Figure 18-3.

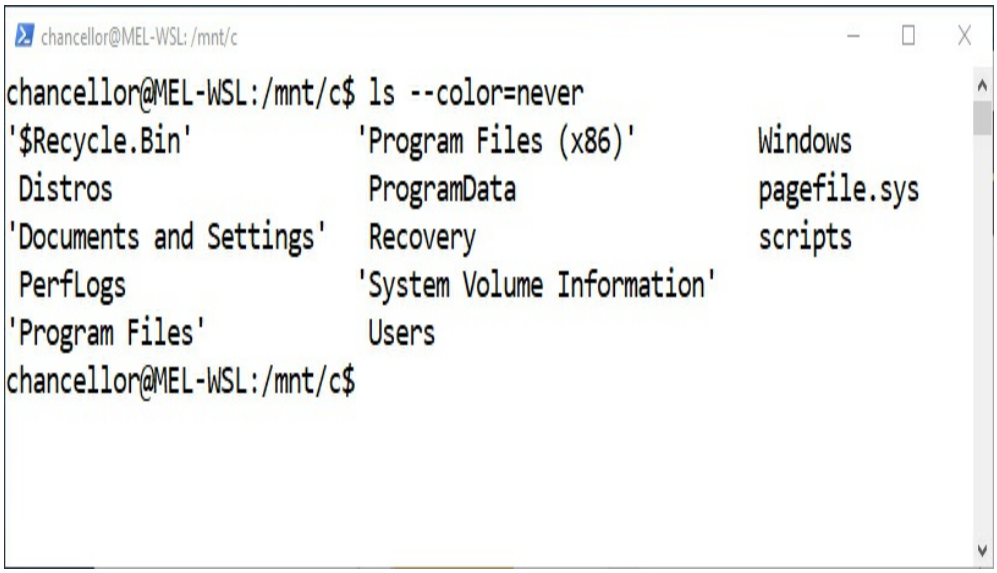


Figure 18-3 Windows file system directory listing output

Once WSL is installed, you can run Linux utilities directly from the Windows Server command line by using the *wsl.exe <command>* syntax. For example, Figure 18-4 shows the output of the *ps -aux* command.

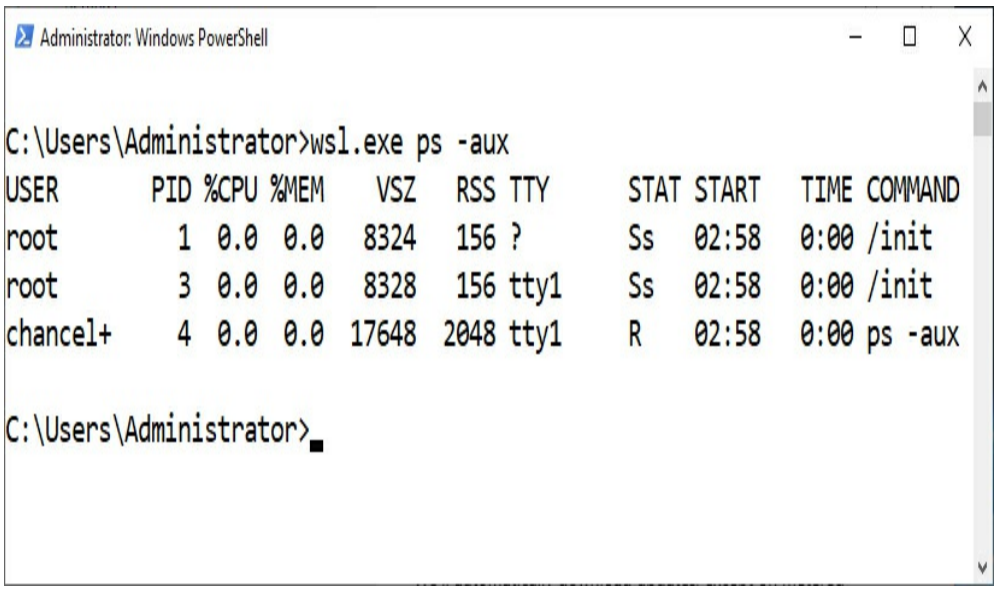


Figure 18-4 List of WSL processes

Inside OUT

Remember to update

WSL will not be automatically updated by running Windows Update, and you need to manually update the distribution using the appropriate distribution update tools, such as apt-get for Ubuntu or Debian. Once you've installed a WSL distribution, the first thing you should do is run a software update!

WSL 2.0

WSL 2.0 is a newer version of Windows Subsystem for Linux that uses virtualization to enable WSL to interact with a Linux kernel, which leads to better performance. Unlike WSL 1, which uses a translation layer to allow Linux workloads to run on Windows, WSL 2 leverages virtualization to host a lightweight kernel to achieve this goal.

WSL 2 requires that the host supports virtualization. If you want to install WSL 2 on a Windows Server 2019 virtual machine, the virtual machine will need to be configured to support nested virtualization. At the time of writing, you can only use WSL 2 on a preview release of Windows Server available through the Windows Server Insider's program.

In addition to the steps outlined earlier, you'll need to run the following commands and then restart the computer to enable WSL 2.0 on Windows Server.

Click here to view code image

```
Enable-WindowsOptionalFeature -Online -FeatureName Microsoft-
Windows-Subsystem-Linux
Enable-WindowsOptionalFeature -Online -FeatureName
VirtualMachinePlatform
```

To ensure that all future WSL distributions use WSL 2.0 instead of WSL, run the following command:

Click here to view code image

```
Wsl --set-default-version 2
```

You can verify which Linux version your WSL distributions are configured for by using the following command, as shown in Figure 18-5.

Click here to view code image

```
Wsl --list –verbose
```

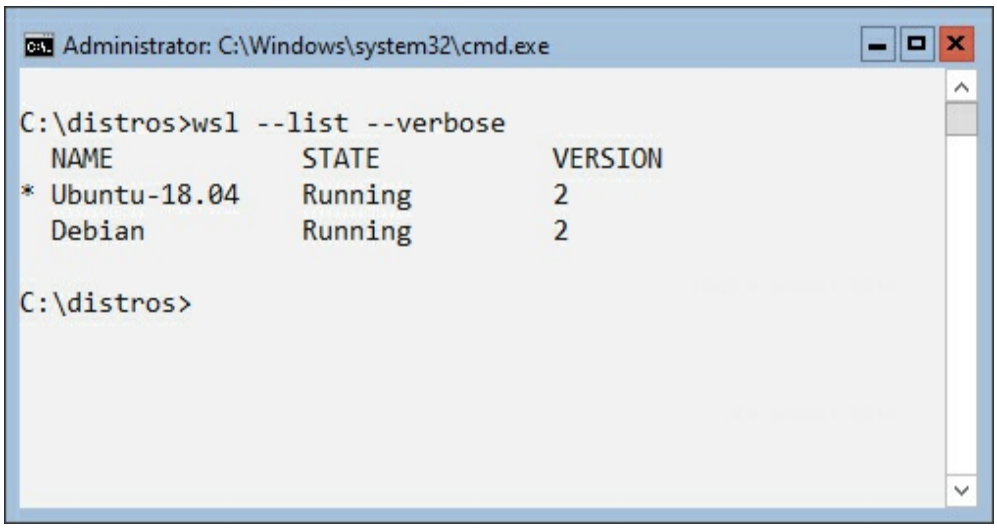


Figure 18-5 List of WSL distribution versions

If the *wsl --list --verbose* command does not function on a Windows Server computer, that version of the operating system does not support WSL 2.0.

More Info

WSL 2

You can find out more about WSL at: https://docs.microsoft.com/en-us/windows/wsl/wsl2-index.

Chapter 19

Hardening Windows Server and Active Directory

Hardening Active Directory

Hardening Windows Server

Shielded VMs

Windows Server 2019 is Microsoft's most secure version of the Windows Server operating system, but you can configure the operating system to be far more secure than it is in a default deployment.

Hardening is the process of configuring security controls to improve security. The drawback of hardening a service such as Active Directory or Windows Server 2019 is that it reduces compatibility. Applying security controls is always a matter of balancing security with the requirements of the environment. For example,

Technet24

while you should disable older authentication protocols such as NTLM, doing so might block an older business-critical system in your environment from being able to authenticate against Active Directory.

Always perform hardening activities that are commensurate with your organization's risk profile. If you're managing the Active Directory environment for a federal government department, you have a substantially different risk profile than if you are managing the Active Directory environment for the local school district.

Inside OUT

Tighten and test slowly

Be sure to implement security controls incrementally. Most security controls aren't enabled by default because turning them on often breaks compatibility with something. If you turn on all possible security controls at once, it's likely that a raft of compatibility problems will arise, and you will have a substantial amount of trouble figuring out which control caused them. If you implement one security control at a time and a compatibility problem arises, you'll easily be able to guess which control was responsible for the problem.

HARDENING ACTIVE DIRECTORY

Active Directory controls the authentication, authorization, and the basic security configuration infrastructure of a Windows Server environment. You should make it as hard as possible for attackers to compromise your Active Directory infrastructure by enabling as many security controls as possible.

Hardening domain controllers

Domain controllers are the vaults that host the “crown jewels” of any Windows Server network. Attackers who can compromise a domain controller can compromise the Active Directory instance hosted on that domain controller. This gives attackers the ability to grant themselves any Active Directory domain privilege they desire. You want to ensure that your organization’s domain controllers are as secure as possible. Some obvious steps you can take to secure your organization’s domain controllers include:

- **Only run the Server Core installation option.** The Server Core installation option of Windows Server has a smaller attack surface than the Server with Desktop Experience version of Windows Server.
- **Only host AD DS and DNS on domain controllers.** Generally, you should also avoid hosting workloads other than DNS on a computer configured as a domain controller as a further way of minimizing the attack surface. The more workloads a domain controller hosts, the more components there are for an attacker to attempt to exploit to gain control of that computer.
- **Run the most recent version of Windows Server.** Domain controllers should always run the most recent version of Windows Server. When a new version of Windows Server is released, the first thing you should do is deploy domain controllers running the new operating system and retire domain controllers running previous versions of Windows Server. Newer versions of Windows Server are inherently more secure than previous versions of Windows Server.
- **Stay up to date with monthly software updates.** You should ensure that you keep domain controllers patched with the most

recent software updates. You should aggressively update domain controllers. As you should not run workloads beyond Active Directory and DNS on a domain controller, you are unlikely to encounter compatibility problems with new software updates (though you should still test on one domain controller before rolling out to all of the others).

- **Enable Device Guard.** Only authorized applications and processes should be able to run on an Active Directory domain controller. Application whitelisting is the process by which only authorized applications and processes can be run; any application or process that is not explicitly authorized will be unable to run. The most effective method of applying application whitelisting is enabling Device Guard. You'll learn more about Device Guard in Chapter 20, "Security systems and services."
- **Block domain controllers from directly communicating with hosts on the Internet.** You should configure perimeter network firewalls so that domain controllers on internal networks can't directly communicate with hosts on the Internet. Software updates from domain controllers should be obtained from sources on the internal network, such as a WSUS server.
- **Only accept administrative connections from authorized hosts.** Domain Controllers should only accept RDP and remote PowerShell sessions from specific authorized hosts. These hosts should be privileged access workstations or jump servers, which you learned about in Chapter 1, "Administration tools."
- **Keep domain controllers in domain controllers OU.** Avoid moving domain controllers out of the default domain controllers OU because this can lead to unexpected configurations being assigned through group policy.
- **Ensure that Directory Services Restore Mode (DSRM) passwords are changed regularly.** This password allows out-of-band access to a domain controller. Passwords should be different for each domain controller; passwords should be updated every couple of months; and passwords should be stored in a secure

location, such as a safe.

Least privilege

Least privilege is the practice of assigning only the minimum required privileges to an account for the account to perform its assigned tasks. Put another way, accounts should not be assigned rights and privileges in excess of those required to perform the specific tasks for which the account is used.

The benefit of using least privilege is that in the event that an attacker manages to compromise an account, that attacker is limited to acting with the privileges assigned to that account. Least privilege also reduces the chance that damage might inadvertently be done to a system. For example, an account used for backup and recovery can't be used to add or remove software from the server if that account doesn't have permission to do anything beyond backup and recovery.

When it comes to assigning privileges, a very bad habit of many IT professionals is to simply add an account to the local Administrators group on the computer where the privilege is required. This is especially true with service accounts where proper procedure is to edit the local security group policies and assign very specific rights; however, many administrators take a quick shortcut and assign all necessary rights by adding the service account to the local Administrators group. This

is one of the reasons that attackers go after service accounts, which often use a single, unchanging, and organization-wide password and that also are members of the local Administrator groups on which they are installed.

Inside OUT

Day-to-day accounts

IT operations staff and developers in your organization should have separate accounts for performing day-to-day activities, such as reading email, browsing the Internet, and filling in TPS reports. The day-to-day accounts of IT operations staff and developers should not have local administrative privileges. Should a change in the configuration of IT staff workstations need to occur, a specific privileged account that exists to perform that task should be used. In highly secure and well-resourced organizations, privileged accounts can only be used on special privileged access workstations (PAWs). These workstations are locked down using technologies such as Device Guard and Credential Guard, which makes them especially resistant to malware, credential theft, and reuse attacks.

More Info

Least privilege

You can learn more about least-privilege administrative models at https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/plan/security-best-practices/implementing-least-privilege-administrative-models.

Role-Based Access Control

Role-Based Access Control (RBAC) operationalizes the least-privilege principle. Instead of having all-powerful administrator accounts that can perform any task on any system, you instead parcel out administrative privileges that limit what an account can do and where

an account can do it.

For example, rather than giving help desk support staff the ability to reset the passwords of any user in the domain, you instead delegate them the right to only be able to reset the passwords of user accounts in a specific organizational unit. This way, you allow the support staff to reset the passwords of a specific class of user account; for example, you could grant support staff the authority to reset passwords that are used by people in one department without giving them the ability to reset the passwords of domain administrator accounts.

You delegate privileges for organizational units in Active Directory using the Delegation Of Control Wizard, as shown in Figure 19-1.

Technet24

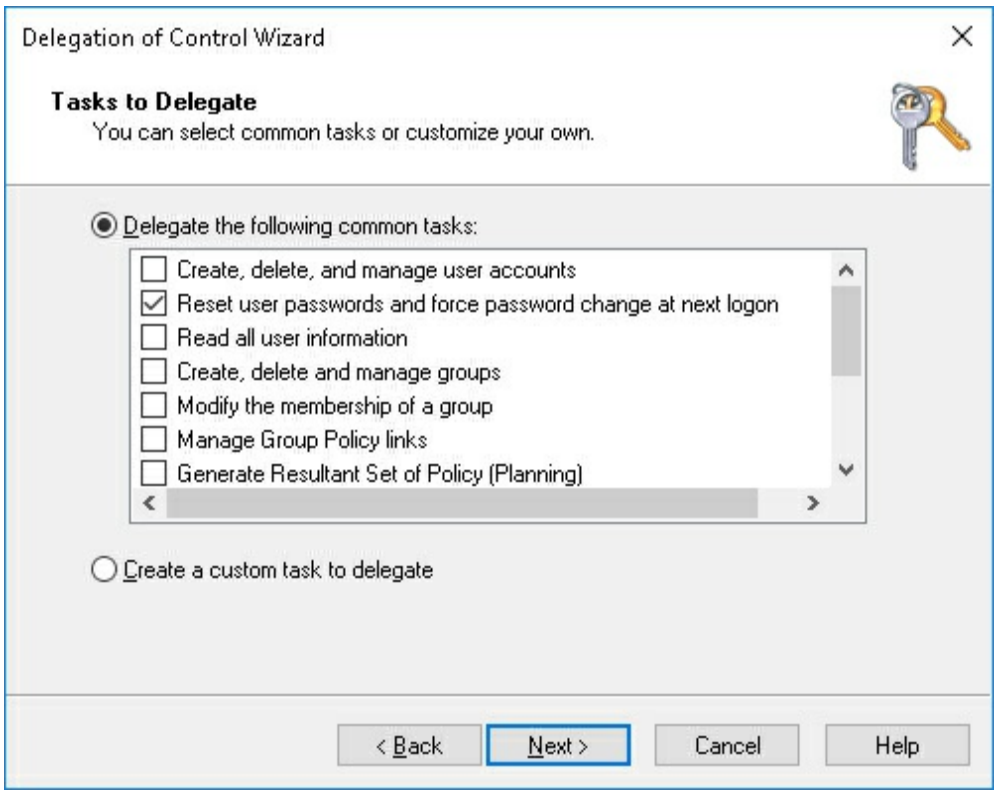


Figure 19-1 Delegating the reset password right

Password policies

Account policies include both password policies and account lockout policies. You configure the domain password and lockout policies by configuring the default domain GPO. You can override the domain policies for specific security groups or user accounts using fine-grained password policies. You configure fine-grained password policies using the Active Directory Administrative Center console.

Default password policy settings include:

- **Number Of Remembered Passwords.** Users cannot change

their passwords to a password remembered by the system. By default, Active Directory remembers the previous 24 passwords that have been used by an account.

- **Maximum Password Age.** This is the maximum amount of time in days that a password can remain the same before it needs to be changed. The default value is 42 days.
- **Minimum Password Age.** This is the minimum amount of time in days that a newly changed password must be kept. This policy prevents users from cycling through password changes to exceed the password history setting, so they can return to using the same password.
- **Minimum Password Length.** This is the minimum number of characters required for a password. The default is 7. Services such as Azure now require a minimum of 12 characters for administrator account passwords on IaaS virtual machines.
- **Password Complexity Requirement.** This setting determines whether passwords must include three of the four following characteristics: uppercase letters, lowercase letters, numbers, and symbols. This setting is enabled by default.

Lockout policy determines:

- The period of time an account is locked out when a user enters a specific number of incorrect passwords in a specific amount of time. The default is that an account is not locked out when incorrect passwords are entered. This means that an attacker or someone who is just forgetful can brute-force guess passwords without any lockout occurring.
- The number of incorrect passwords that trigger a lockout over a specific period of time. When configuring this policy, you should allow users to make a reasonable number of guesses. If you don't, you'll end up flooding the service desk with password reset calls.
- The amount of time that must elapse before the account lockout counter is reset. A correctly entered password also resets the lockout

counter.

Account security options

There are additional security options that you should consider configuring for highly privileged accounts. These include configuring the following settings:

- **Logon Hours.** Use this setting to configure when users can use an account. AD DS does not authenticate someone attempting to sign in outside of these hours. By default, accounts are configured so that users can sign in at all times. Consider configuring privileged accounts so that users can use them only at specific times. This stops a user from signing in after hours, either specifically or with a compromised account, to get up to no good when they assume that no one is watching. Organizations that do need to use privileged accounts to perform after-hours maintenance tasks can create special accounts for those tasks, or they can temporarily modify an existing account.
- **Logon Workstations.** Use this setting to limit the computers to which a particular account can sign in. By default, users can use an account to sign in to any computer in the domain. Use this setting to ensure that users can use privileged accounts only on specific, specially configured administrative workstations or certain servers. Limiting which computers' users can use an account reduces the locations where that account can be used.
- **Password Never Expires.** You should use this setting reluctantly when configuring privileged accounts because it absolves the account from the domain password policy. You can use products such as Microsoft Identity Manager to assist with password management for privileged accounts.
- **Smart Card Is Required For Interactive Logon.** Use this setting to ensure that a smart card must be present for the account sign in to occur. In high-security environments, you should deploy smart cards and enable this option to ensure that only an

authorized person who has both the smart card and the account credentials can use the privileged account.

- **Account Is Sensitive And Cannot Be Delegated.** Use this setting to ensure that trusted applications cannot forward the account credentials to other services or computers on the network. Enable this setting for highly privileged accounts.
- **This Account Supports Kerberos AES 256 Bit Encryption.** Use this setting to allow Kerberos AES 256-bit encryption. Where possible, you should configure this option for privileged accounts and have them use this form of Kerberos encryption over the AES 128-bit encryption option.
- **Account Expires.** Use this setting to configure an expiration date for an account. Configuring privileged accounts with expiration dates ensures that privileged accounts do not remain available to use beyond when they are required.

When you join a computer to a domain for the first time, it creates a computer account in the Computers container. If you want to create an account in a specific container, you'll need to pre-stage the computer account and place it in the appropriate OU. A pre-staged computer account must have a name that matches that of the computer that is joining the domain.

Computer accounts can be made members of the domain security group. This group membership influences how the Local System and Network Service accounts on the computer function. This is because, by default, these services use the computer's credentials when interacting with resources on the network. This gives the computer account the rights and privileges

assigned to any security group to which the computer is a member.

Computer accounts have automatically assigned passwords that are updated every 30 days. If a computer doesn't connect to a domain within 30 days, a new password is assigned the next time a connection is established. If you disable a computer account, the computer cannot connect to the domain, and domain users are unable to sign in to the computer until the account is re-enabled.

Resetting a computer account removes the relationship between the computer and the domain. To fix this, you'll need to either join the computer back to the domain, or reestablish the broken trust relationship using the following PowerShell command, specifying the credentials of a member of the Domain Administrator group:

Click here to view code image

```
Test-ComputerSecureChannel -credential <domain>\<admin> -Repair
```

If you delete a computer account, it's necessary to rejoin the computer to the domain manually. When you delete a computer account, all information associated with the account is removed from Active Directory. You should only delete computer accounts once a computer has been decommissioned.

Protected accounts

The Protected Users group, shown in Figure 19-2, provides you with a method of protecting highly privileged user accounts from being compromised. It does this by blocking the use of less-secure security and authentication options.

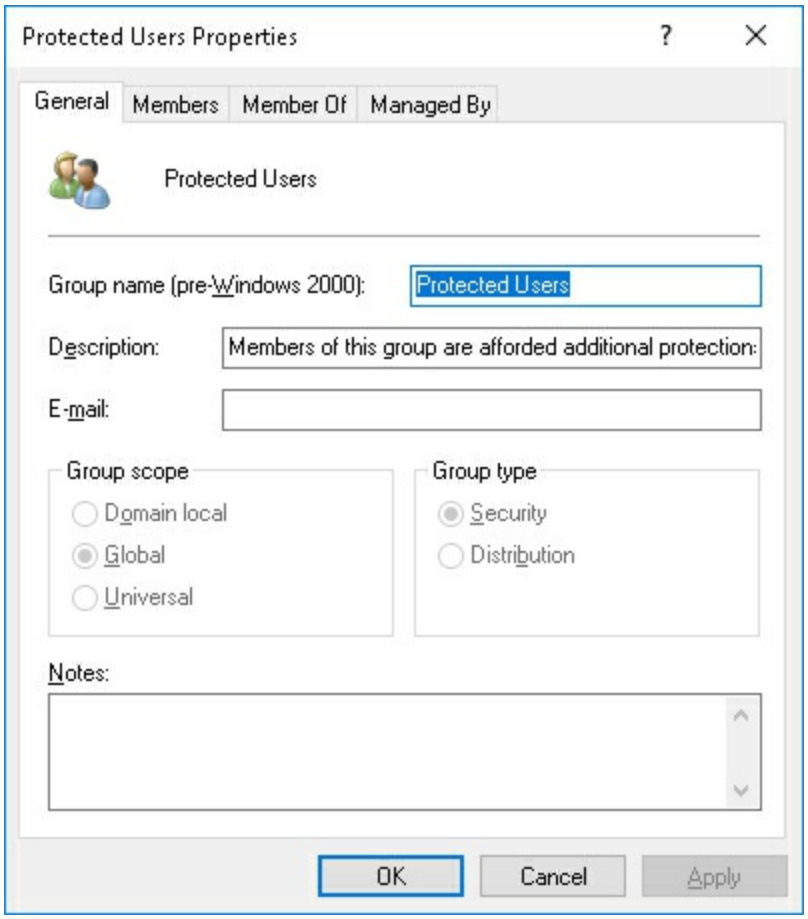


Figure 19-2 Protected Users Properties

For example, credentials are not cached if an account that is a member of the Protected Users group signs in

to a computer running Windows 8.1, Windows 10, Windows Server 2012 R2, Windows Server 2016, or Windows Server 2019. Additionally, accounts that are members of this group cannot use the following security options:

- Default credential delegation (CredSSP)
- Windows Digest
- NTLM—NTOWF
- Kerberos long-term keys
- Sign in offline

If the domain functional level is Windows Server 2012 R2 or higher, user accounts that are members of the Protected Users group cannot:

- Use NT LAN Manager (NTLM) for authentication
- Use DES for Kerberos pre-authentication
- Use RC4 cipher suites for Kerberos pre-authentication
- Be delegated using constrained delegation
- Be delegated using unconstrained delegation
- Renew user tickets (TGTs) past the initial 240-minute lifetime

Only user accounts should be added to the Protected Users group. You should not add computer accounts or service accounts to this group. In secure environments, all privileged accounts should be members of this group, with exceptions occurring only when specific incompatibilities arise that cannot be resolved in

another way.

More Info

Protected Users security group

You can learn more about the Protected Users security group at https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/dn466518(v=ws.11).

Authentication policies silos

Authentication policy silos allow you to define relationships between the user, computer, and managed service accounts. A user, computer, and managed service account can only belong to a single authentication policy silo. This provides a more robust security method that goes beyond configuring log in restrictions and restricting which accounts can access specific servers in your environment.

Accounts in an authentication policy silo are associated with a silo claim. You can use this silo claim to control access to claims-aware resources. For example, you can configure accounts so that only accounts that are associated with a specific silo claim can access particularly sensitive servers. This means only accounts that are associated with a Certificate Services silo claim can sign in to a computer that has the Active Directory Certificate Services role installed.

Authentication policies allow you to configure settings,

such as the TGT lifetime and access-control conditions, which specify conditions that must be met before a user can sign in to a computer. For example, you might configure an authentication policy that specifies a TGT lifetime of 120 minutes and limit a user account so that users can only use it with specific devices, such as privileged access workstations. You can only use authentication policies if you have configured the domain functional level at the Windows Server 2012 R2 or higher.

You configure authentication policies and authentication policy silos using Active Directory Administrative Center, as shown in Figure 19-3.

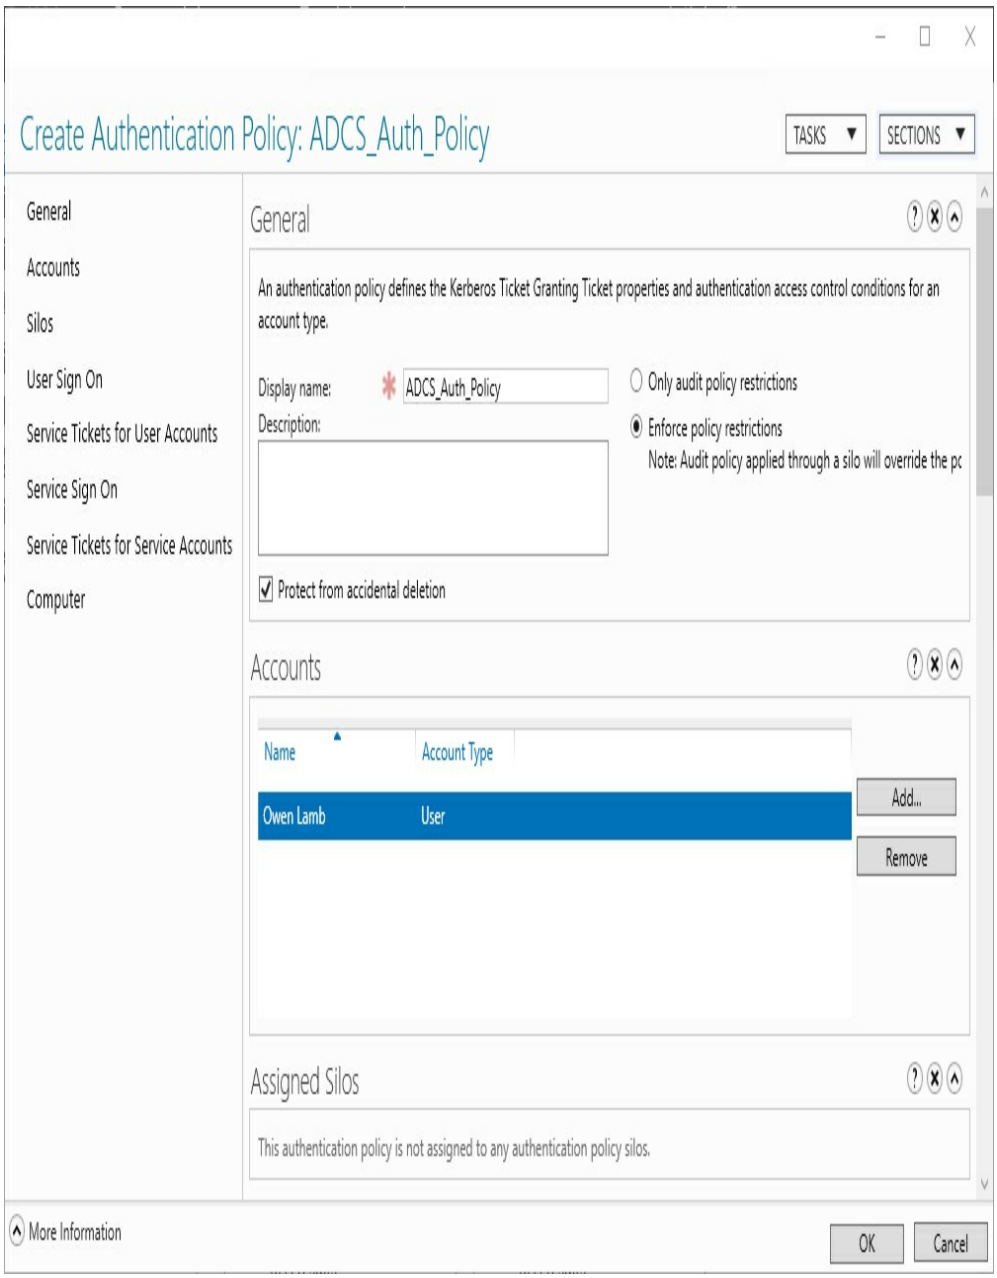


Figure 19-3 Authentication policies

More Info

Authentication policies and silos

You can learn more about authentication policies and silos at https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-

server-2012-R2-and-2012/dn486813(v=ws.11).

Disable NTLM

NTLM (NT Lan Manager) is an older authentication protocol that is not as secure as Kerberos. You can disable NTLM by configuring the Network Security: Restrict NTLM: NTLM Authentication In This Domain policy in the default domain GPO and setting the policy to Disable, as shown in Figure 19-4. This policy can be found in the *Policies\Windows Settings\Security Settings\Local Policies\Security Options* node of the default domain GPO. Prior to disabling NTLM authentication, you should configure the Network Security: Restrict NTLM: Audit Incoming NTLM Traffic policy to determine if and where the NTLM protocol is still being used in your organization's environment.

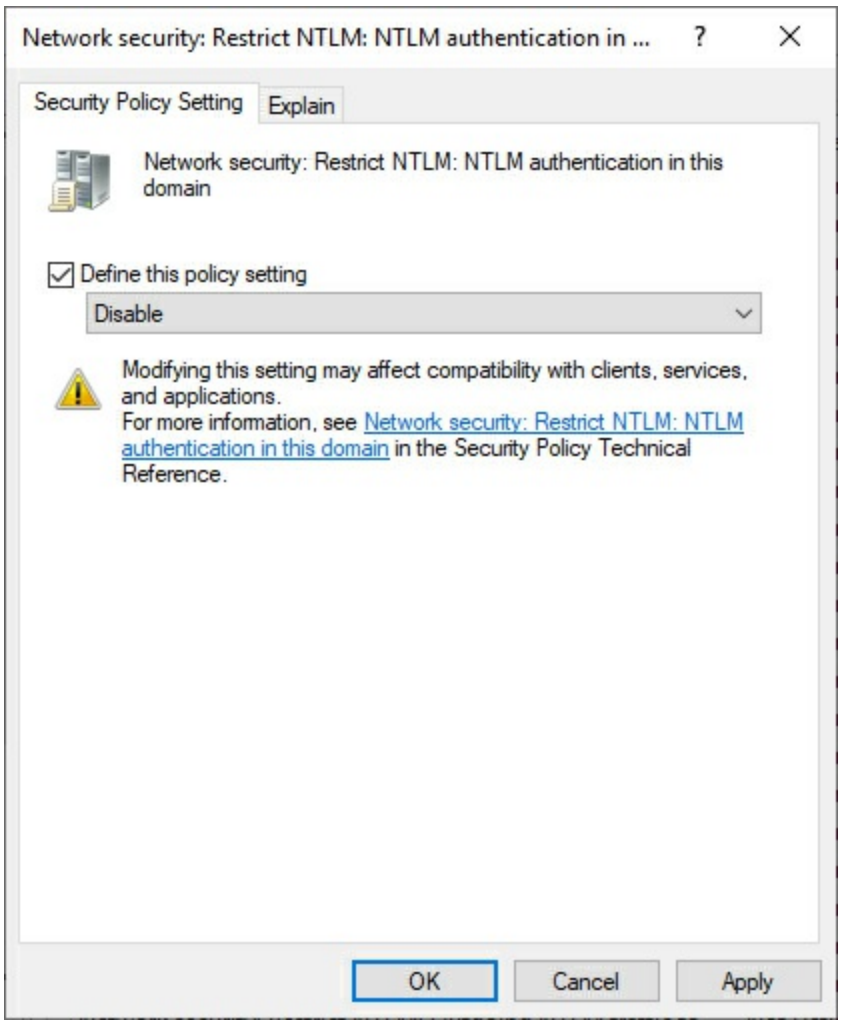


Figure 19-4 Disable NTLM

More Info

Disable NTLM

You can learn more about disabling NTLM at https://docs.microsoft.com/en-us/windows/security/threat-protection/security-policy-settings/network-security-restrict-ntlm-ntlm-authentication-in-this-domain.

Block server operators from scheduling tasks

Configuring the Domain Controller: Allow Server Operator To Schedule Tasks policy as disabled, as shown in Figure 19-5, will block someone who is signed in to a server using the *AT* command to configure scheduled jobs. By default, *AT* jobs run under the SYSTEM account, and attackers often exploit the *AT* command to elevate privileges. If you configure this policy, it will be necessary to specify account credentials when configuring a scheduled task.

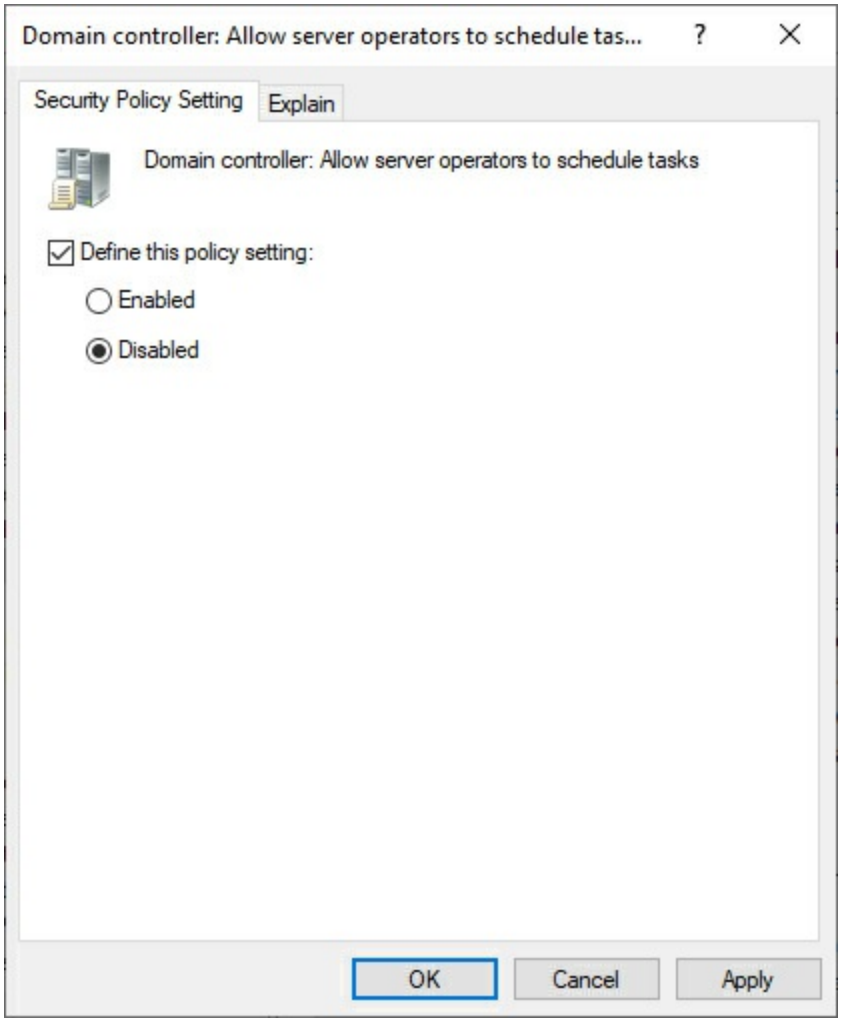


Figure 19-5 Block task scheduling using *AT*

Enable Local Security Authority protection

There are multiple techniques that you can employ to reduce the chance of an attacker successfully executing a “pass the hash” attack. Enabling LSA protection makes it more difficult for the Local Security Authority to be compromised. To configure LSA protection on a domain controller:

1. Open the Registry Editor (*RegEdit.exe*) and navigate to the registry key that is located at: *HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa*.
2. Create (or set the value of) this registry key: *"RunAsPPL"=dword:00000001*.
3. Restart the computer.

> **More Info**
>
> **Local Security Authority protection**
>
> You can learn more about configuring Local Security Authority protection at https://docs.microsoft.com/en-us/windows-server/security/credentials-protection-and-management/configuring-additional-lsa-protection.

KRBTGT account password

While you should definitely change the KRBTGT account password if your domain is compromised by an attacker or when someone who has had Domain Admin privileges leaves the organization, it is also good security hygiene to change the KRBTGT account password every

six months. Changing the KRBTGT password involves an initial password change, waiting for replication to occur, and then changing the password again. Many people wait 24 hours between the first and second changes. You must perform the password operation twice because of the way the password history of the KRBTGT account functions. Updating the KRBTGT password twice means that replication cannot occur using an older password. If you only do it once, and the domain has been compromised, an attacker might leverage the second stored password.

You can reset the KRBTGT account password by performing the following steps:

1. Open Active Directory Users And Computers.
2. In the View menu, click Advanced Features.
3. In the Users container of the domain, click the KRBTGT user account, as shown in Figure 19-6.

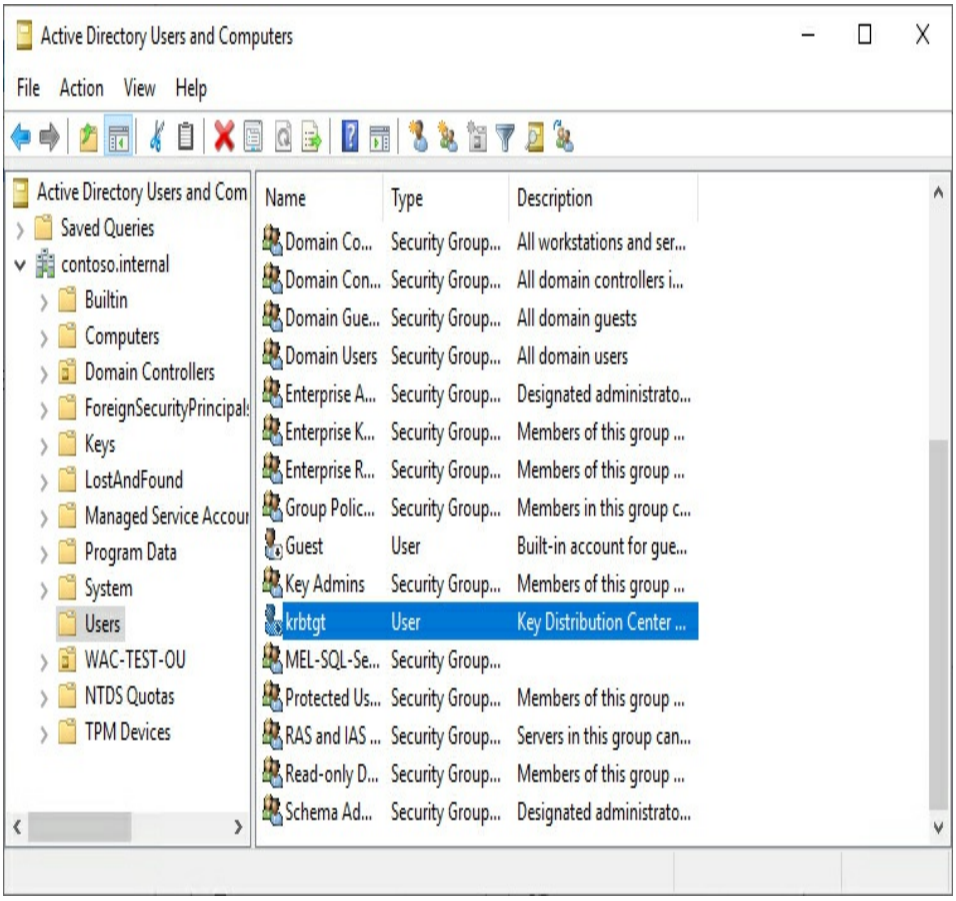


Figure 19-6 KRBTGT account

4. From the Action menu, select Reset Password. You should then type in and confirm a new password. The password you enter in this dialog box isn't important because Active Directory will trigger the creation of a separate strong password that is independent of the password you configure.

Enhanced Security Administrative Environment forest

Enhanced Security Administrative Environment (ESAE) forests are an approach to Active Directory privileged access design in which you deploy a dedicated administrative Active Directory forest to host privileged

Technet24

accounts, privileged groups, and privileged access workstations for a production forest. In this design, the ESAE forest is configured with a one-way trust relationship with a production forest. A one-way trust relationship means that accounts from the ESAE forest can be used in the production forest; however, accounts in the production forest cannot be used in the ESAE forest.

The production forest is the Active Directory forest in which an organization's day-to-day activities occur. You configure the production forest so that administrative tasks in the production forest can only be performed by using accounts that the ESAE forest hosts.

The ESAE forest design provides the following benefits:

- **Locked-down accounts.** A standard non-privileged user account in the ESAE forest can be configured as a highly privileged user account in the production forest. For example, you make a standard user account in the ESAE forest a member of the Domain Admins group in a domain in the production forest. You can lock down the standard user account hosted in the ESAE forest so that it cannot sign in to hosts in the ESAE forest and can only be used to sign in to hosts in the production forest. Should the account be compromised, the attacker cannot use that account to perform administrative tasks in the ESAE forest.
- **Selective authentication.** ESAE forest design allows you to leverage a trust relationship's selective authentication feature. For example, you can configure sign-ins from the ESAE forest so that they can only occur on specific hosts in the production forest. One strategy is to configure selective authentication so that privileged accounts can only be used on privileged access workstations or

jump servers in the production forest.

- **A simple way to improve security.** ESAE forest design provides substantive improvement in the security of existing production forests without requiring a complete rebuilding of the production environment. The ESAE forest approach has a small hardware/software footprint and only affects IT users. In this design, standard user accounts remain hosted in the production forest. Only privileged administrative accounts are hosted in the ESAE forest. As part of the transition to the ESAE forest model, you eventually remove almost all privileged accounts that are native to the production forest. Because they are hosted in a separate forest, privileged administrative accounts can be subject to more stringent security requirements than standard user accounts in the production forest.

Inside OUT

Failsafe admin account

With an ESAE forest design, you should remove almost all privileged accounts from the production forest. You should leave a failsafe admin account in the production forest that allows you to fix things should everything go pear shaped, but this admin account should be given an insanely long and complex password that's stored securely in a safe that someone who is not a member of the IT operations staff has access to and should only be used as a last resort.

An ESAE forest design should have the following properties:

- Limit the function of the ESAE forest to hosting accounts of administrative users for the production forest.
- Do not deploy applications or additional resources in the ESAE forest. An exception to this is when you deploy technologies such as Microsoft Identity Manager, which is used to increase security when implementing Privileged Access Management (PAM).
- The ESAE forest should be a single-domain Active Directory forest.

The ESAE forest only hosts a small number of accounts that need to have strict security policies applied.

- Only use one-way trusts. You should only configure a one-way trust where the production forest trusts the ESAE forest. This means that accounts from the ESAE forest can be used in the production forest but that accounts from the production forest cannot be used in the ESAE forest. Accounts used for administrative tasks in the production forest should be standard user accounts in the ESAE forest. If an account is compromised in the production forest, it cannot be used to elevate privileges in the ESAE forest.

Configure ESAE forest servers in the following ways:

- Validate installation media. Use the available checksums to validate that the media hasn't been modified.
- All servers should run the most recent version of the Windows Server operating system.
- Servers should be updated automatically with security updates.
- Security configuration toolkit baselines should be used as the starting point for server configuration.
- Servers should be configured with secure boot, BitLocker volume encryption, Credential Guard, and Device Guard.
- Servers should be configured to block access to USB storage.
- Servers should be on isolated networks. Inbound and outbound Internet connections should be blocked.

More Info

Securing privileged access

You can learn more about securing privileged access at https://docs.microsoft.com/en-us/windows-server/identity/securing-privileged-access/securing-privileged-access-reference-material.

HARDENING WINDOWS SERVER

Just as you should harden Active Directory, you should secure Windows Server hosts as much as possible without triggering compatibility issues. Beyond the advice offered in this chapter, multiple third-party guides exist that list recommended security configurations for Windows Server. These guides are extremely lengthy and provide an excellent way of understanding exactly which security controls are available to harden Windows Server. The two most important organizations providing guides are:

- **Center for Internet Security.** Benchmarks exist for Microsoft Exchange, Microsoft SQL Server, and Windows Server 2019. You can find out more at *https://www.cisecurity.org/cis-benchmarks*.
- **Defense Information Systems Agency (DISA).** DISA Security Technical Implementation Guide (STIG) is a respected security benchmark. You can access benchmarks for workloads including Active Directory, .NET Framework, Microsoft Exchange, Microsoft SharePoint, Microsoft SQL Server, and Windows Server 2019. You can find out more at *https://public.cyber.mil/stigs/downloads/*.

In Chapter 20, "Security systems and services," you'll also learn about the Security Configuration Toolkit, which allows you to assess the security of your Windows Server workloads against some existing security benchmarks.

User rights

Rather than adding a user or service account to the local

Technet24

Administrators group, you should use Group Policy to assign the specific rights that a user account needs to perform the tasks it needs to perform. Assign rights as close to the object as possible. For example, don't give a user account the ability to log in through Remote Desktop Services to all computers in an OU if the user account only needs remote desktop access to one or two computers. You can use Group Policy to configure the rights outlined in Table 19-1.

Table 19-1 User rights assignment policy

User rights assignment policy	Function
Access Credential Manager As A Trusted Caller	Used by Credential Manager during backup and restore. Do not assign this privilege to user accounts.
Access This Computer From The Networ	Specifies which accounts and groups may connect to the computer from the network. Does not affect Remote Desktop Services.

k	
Act As Part Of The Operating System	Allow a process to impersonate an account without requiring authentication. Processes that require this privilege are often assigned the Local System account.
Add Workstations To Domain	Grants the ability to join workstations to the domain.
Adjust Memory Quotas For A Process	Configures which security principals can adjust the maximum amount of memory assigned to a process.
Allow Log On Locally	Specify which accounts can sign in locally to a computer. Alter this policy on privileged access workstations to remove members of the Users group. This limits which accounts can sign in to a computer. By default, any authenticated user can sign in to any workstation or server except for a domain controller.
Allow Log On Through	Determines which accounts and groups can sign in remotely to the computer using Remote Desktop. Configure this policy to allow users to access enhanced session mode for Hyper-V VMs.

Technet24

Remote Desktop Services	
Back Up Files And Directories	Grant permission to back up files, directories, registry, and other objects that the user account wouldn't normally have permission to. Assigning this right gives indirect access to all data on a computer. A person who holds the right can back up that data and then recover it in an environment over which they have complete control.
Bypass Traverse Checking	Gives users with this right the ability to traverse directories on which they don't have permission. Does not allow the user to list the contents of that directory.
Change The System Time	Accounts with this right can alter the system time. System time is separate from the computer's time zone.
Change The Time Zone	Accounts with this right can alter the time zone but not the system time.
Create A Pagefile	Allows the account assigned this right the ability to create and modify the page file.

Create A Token Object	Specifies which accounts can be used by processes to create tokens that allow accesses to local resources. You should not assign this right to any user who you do not want to give complete control of the system. This privilege can be used to elevate to local Administrator privileges.
Create Global Objects	Determines which accounts can create global objects that are available to all sessions. You should not assign this right to any user who you do not want to give complete control of the system. This privilege it can be used to elevate to local Administrator privileges.
Create Permanent Shared Objects	Specify which accounts can create directory objects by using the object manager.
Create Symbolic Links	Specify which accounts can create symbolic links from the computer they are signed in to. Assign this right only to trusted users because symbolic links can expose security vulnerabilities in apps that aren't configured to support them.
Debug Programs	Specify which accounts can attach a debugger to processes within the operating system kernel. Only required by developers writing new system components. Rarely, if ever, necessary for developers writing applications. Removing the default Administrator from this account will substantially limit the likelihood that utilities, such as Mimikatz, can be used to scrape credentials.
Deny Access	Blocks specified accounts and groups from accessing the computer from the network. This setting overrides the policy

To This Computer From The Network	that allows access from the network.
Deny Log On As A Batch Job	Blocks specified accounts and groups from signing in as a batch job. Overrides the Log On As A Batch job policy.
Deny Log On As A Service	When assigned, blocks service accounts from registering a process as a service. Overrides the Log On As A Service policy. Does not apply to Local System, Local Service, or Network Service accounts.
Deny Log On Locally	When assigned, blocks accounts from signing on locally. This policy overrides the Allow Log On Locally policy.
Deny Log On Through Remote Desktop Services	When assigned, blocks accounts from signing in by using Remote Desktop Services. This policy overrides the Allow Sign In Through Remote Desktop Services policy.

Enable Computer And User Accounts To Be Trusted For Delegation	Determines whether you can configure the Trusted For Delegation setting on a user or a computer object.
Force Shutdown From A Remote System	Accounts assigned this right are allowed to shut down computers from remote locations on the network.
Generate Security Audits	Determines which accounts can use processes to add items to the security log. Because this right allows interaction with the security log, it presents a security risk when you assign this to a user account.
Impersonate A Client After Authentication	Allows apps running on behalf of a user to impersonate a client. This right can be a security risk, and you should assign it only to trusted users.

Increase A Process Working Set	Accounts assigned this right can increase or decrease the number of memory pages visible to the process in RAM.
Increase Scheduling Priority	Accounts assigned this right can change the scheduling priority of a process.
Load And Unload Device Drivers	Accounts assigned this right can dynamically load and unload device drivers into kernel mode. This right is separate from the right to load and unload Plug and Play drivers. Assigning this right is a security risk because it grants access to the kernel mode.
Lock Pages In Memory	Accounts assigned this right can use a process to keep data stored in physical memory, blocking that data from being paged to virtual memory.
Log On As A Batch Job	Accounts assigned this right can be signed in to the computer by means of a batch-queue facility. This right is only relevant to older versions of the Windows operating system, and you should not use it with newer versions, such as Windows 10 and Windows Server 2016.
Log On As A Service	Allows a security principal to sign in as a service. You need to assign this right to any service that is configured to use a user account, rather than one of the built-in service accounts.

Manage Auditing And Security Log	Users assigned this right can configure object access auditing options for resources such as files and AD DS objects. Users assigned this right can also view events in the security log and clear the security log. Because attackers are likely to clear the security log as a way of hiding their tracks, you should not assign this right to user accounts that would not normally be assigned local Administrator permissions on a computer.
Modify An Object Label	Users assigned this right can modify the integrity level of objects, including files, registry keys, or processes owned by other users.
Modify Firmware Environment Values	Determines which users can modify firmware environment variables. This policy is used primarily for modifying the boot configuration settings of Itanium-based computers.
Perform Volume Maintenance Tasks	Determines which accounts can perform maintenance tasks on a volume. Assigning this right is a security risk because it might allow access to data stored on the volume.
Profile Single Process	Determines which accounts can leverage performance monitoring tools to monitor nonsystem processes.

Technet24

Profile System Performance	Determines which accounts can leverage performance-monitoring tools to monitor system processes.
Remove Computer From Docking Station	When assigned, the user account can undock a portable computer from a docking station without signing in.
Replace A Process Level Token	When assigned, the account can call the *CreateProcessAsUser* API so that one service can trigger another.
Restore Files And Directories	Allows users assigned this right the ability to bypass permissions on files, directories, and the registry, so that they can overwrite these objects with restored data. This right represents a security risk because a user account assigned this right can overwrite registry settings and replace existing permissions.
Shut Down The System	Assigns the ability for a locally signed-in user to shut down the operating system.
Synchronize	Assigns the ability to synchronize AD DS data.

Directory Service Data	
Take Ownership Of Files Or Other Objects	When assigned, the account can take ownership of any securable object, including AD DS objects, files, folders, registry keys, processes, and threads. Represents a security risk because it allows an assigned user the ability to take control of any securable object.

Service accounts

Service accounts allow services running on a computer to interact with the operating system as well as resources on the network. Windows Server 2019 uses three types of built-in service accounts, each of which is suitable for a specific set of circumstances. These accounts are as follows:

- **The Local System** (*NT AUTHORITY\SYSTEM*) account has privileges that are equivalent to those assigned to a user account that is a member of the local Administrators group on the computer. A service that uses this account can act by using the computer account's credentials when interacting with other resources on the network.
- **The Local Service** (*NT AUTHORITY\LocalService*) account has privileges that are equivalent to those assigned to a user account that is a member of the local Users group on the computer. A service that uses this account can access resources on the network without credentials. You use this account when a service does not

need to interact with resources on the network and does not need local Administrator privileges on the computer on which it is running.

- **The Network Service** (*NT AUTHORITY\NetworkService*) account has privileges that are equivalent to those assigned to a user account that is a member of the local Users group on the computer. A service that uses this account accesses resources on the network by using the computer account's credentials.

In the past, when you've needed to create a custom service account, you've probably created a user account and then assigned it an appropriate set of permissions. The challenge with this type of account is password management. Many organizations configure custom service accounts with passwords that never expire.

A Group Managed Service Account (gMSA) is a special type of account that has AD DS manage its password. The password of a gMSA is updated every 30 days. You don't need to know the password of a gMSA, even when configuring a service to use that password, because gMSA accounts provide you with a domain-based service account identity without all the hassle of service account password management.

You can only use a gMSA if the forest is running at the Windows Server 2012 functional level or higher. You also must have deployed the master root key for AD DS. You can deploy the master root key by running the following command using an account that is a member of the Enterprise Admins group:

Click here to view code image

```
Add-KdsRootKey -EffectiveTime ((get-date).addhours(-10))
```

You create the gMSA using the *New-ADServiceAccount* cmdlet. When creating the account, you specify a hostname as well as service principle names. These should be associated with the purpose for which you are creating the gMSA. For example, run this command to create a new gMSA named *SYD-SVC1* that is associated with the hostname *SYD-SVC1.adatum.com*:

Click here to view code image

```
New-ADServiceAccount SYD-SVC1 -DNSHOSTNAME SYD-SVC1.adatum.com
```

gMSAs are stored in the Managed Service Accounts container, which you can view using Active Directory Administrative Center, as shown in Figure 19-7.

Technet24

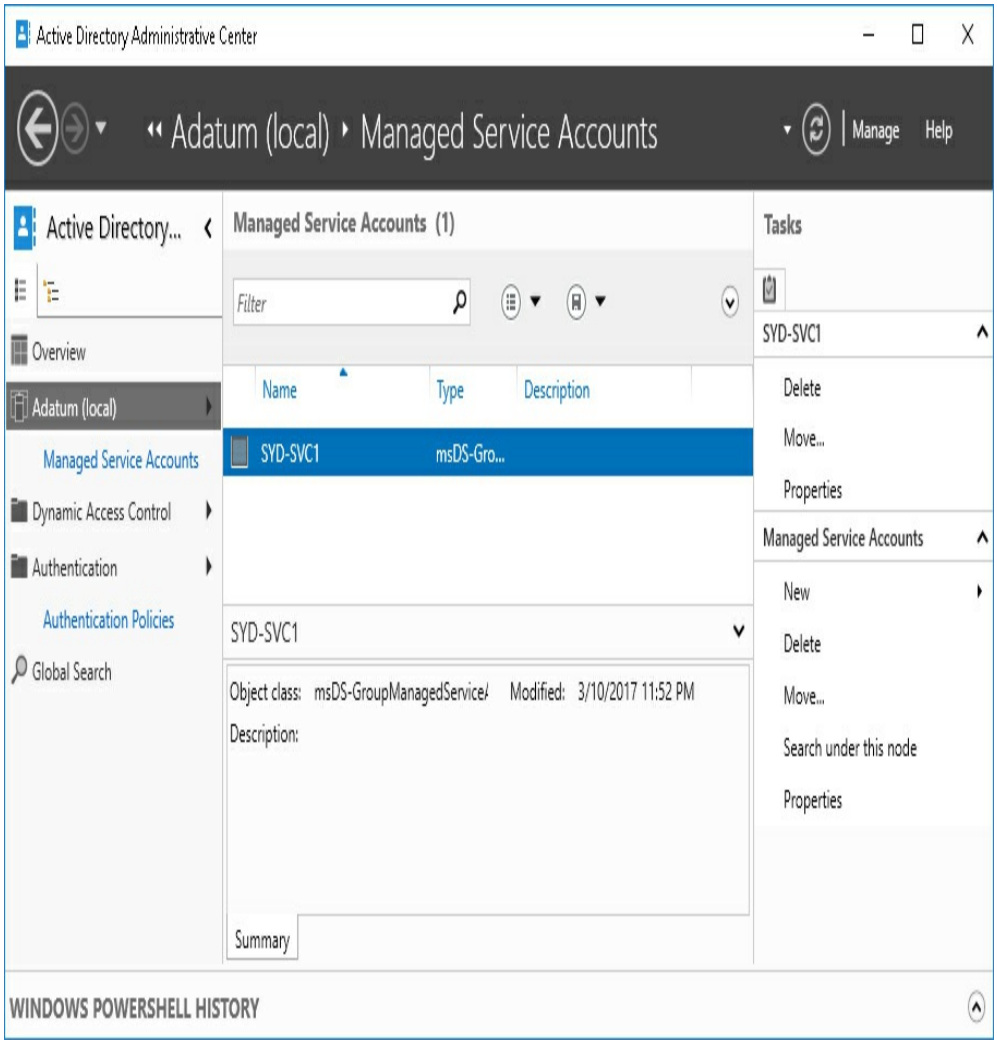


Figure 19-7 Managed Service Accounts container

Once the gMSA is created, you need to configure permissions for specific computers to be able to install and use the account. The simplest way to do this is to create a security group, and then use the *Set-ADSeviceAccount* cmdlet to assign the group permissions to the account. You then add computers to the group. For example, to allow computers in the *SYD-*

SVRS group to use the *SYD-SVC1* gMSA, run this command:

Click here to view code image

```
Set-ADServiceAccount -Identity SYD-SVC1 -
PrincipalsAllowedToRetrieveManagedPassword
SYD-SVRS
```

Once you've added a computer to the group that you've given permission to retrieve the account, you can install the account on the computer using the *Install-ADServiceAccount* cmdlet. Once the account is installed, you can assign it to a service by setting the Log On As settings for the service, as shown in Figure 19-8. You only need to specify the account name and can browse to the account. You don't need to specify the password setting because the password is managed by Active Directory.

Technet24

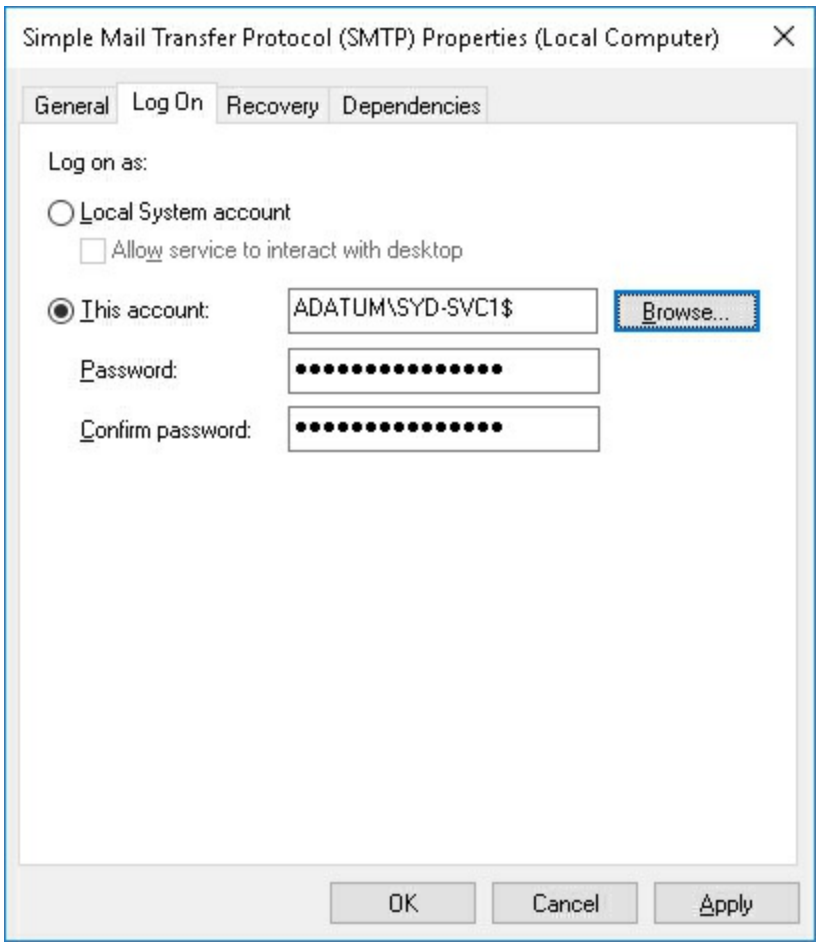


Figure 19-8 Group managed service account

A virtual account is the computer local equivalent of a group Managed Service Account. You can create a virtual account by editing the properties of a service and setting the account name to *NT SERVICE\<ServiceName>*. As with a gMSA, you don't need to specify a password because this is managed automatically by the operating system.

More Info

Group Managed Service Accounts

You can learn more about group Managed Service Accounts at https://docs.microsoft.com/en-us/windows-server/security/group-managed-service-accounts/group-managed-service-accounts-overview.

Just Enough Administration

Just Enough Administration (JEA) allows you to implement Role-Based Access Control (RBAC) functionality through Windows PowerShell remoting. JEA allows you to specify which PowerShell cmdlets and functions can be used when connected to a specific endpoint. You can go further and specify which parameters within those cmdlets and functions are authorized and even specify which values can be used with those parameters.

For example, you could create a JEA endpoint where a user is able to run the *Restart-Service* command, but only where the *Name* parameter is set to DHCPServer. This would allow the user to restart the DHCPServer on the computer they connected to, but it would not restart any other service on the computer.

You can also configure a JEA endpoint to allow other command-line commands to be run, such as *whoami.exe*, though the drawback of this is that you don't have the same level of control when restricting how that command can be run.

JEA endpoints can leverage virtual accounts. This

means that activities performed on the computer through the endpoint use a special temporary virtual account rather than the user's account. This temporary virtual account has local Administrator privileges but is constrained to only using the cmdlets, functions, parameters, and values defined by JEA. The benefits of this include:

- The user's credentials are not stored on the remote system. If the remote system is compromised, the user's credentials are not subject to credential theft and cannot be used to traverse the network to gain access to other hosts.
- The user account used to connect to the endpoint does not need to be privileged. The endpoint simply needs to be configured to allow connections from specified user accounts.
- The virtual account is limited to the system on which it is hosted. The virtual account cannot be used to connect to remote systems. Attackers cannot use a compromised virtual account to access other protected servers.
- The virtual account has local administrator privileges but is limited to performing only the activities defined by JEA. You have the option of configuring the virtual account with the membership of a group other than the local Administrators group to further reduce privileges.

JEA works on Windows Server 2019, Windows Server 2016, and Windows 10. It also functions on previous versions of Windows client and server as long as Windows Management 5.0 is installed. Virtual accounts are not available when using JEA on Windows 7 or Windows 2008 R2, and all activities are performed

using the privileges assigned to the account connecting to the JEA endpoint.

> **Inside OUT**
>
> **JEA drawbacks**
>
> The biggest drawback to JEA is the amount of time that it takes to configure. You'll need to customize each JEA endpoint, which requires that you understand exactly which cmdlets, functions, parameters, and values are required to perform specific tasks. The other drawback of JEA is that it relies upon administrative tasks being performed using the command line. While most Windows Server administrators are comfortable performing tasks using PowerShell, some may be reluctant to be placed in a position where they can only use PowerShell for specific tasks rather than falling back on the RSAT consoles.

Role-capability files

A role-capability file is a special file that allows you to specify what tasks can be performed when connected to a JEA endpoint. Only tasks that are explicitly allowed in the role-capability file can be performed.

You can create a new blank role-capability file by using the *New-PSRoleCapabilityFile* cmdlet. Role-capability files use the *.psrc* extension. For example, run this command to create a new role capability file for a role that allows someone to manage a DNS server:

Click here to view code image

```
New-PSRoleCapabilityFile -Path .\DNSOps.psrc
```

> **Inside OUT**
>
> **Authoring JEA**

When creating JEAconfigurations, it usually makes sense to put them in their own folders. For example, create a DNSOps folder to store all the files related to configuring the DNSOps JEAconfiguration.

Once the *.psrc* file is created, you edit the role-capability file and add the cmdlets, functions, and external commands that are available when a user is connected to the endpoint. You can allow entire Windows PowerShell cmdlets or functions or list which parameters and parameter values can be used.

You can edit a role-capability file in PowerShell ISE, as shown in Figure 19-9. Editing the file involves commenting out the appropriate sections and filling them in with the configuration items that you want to set.

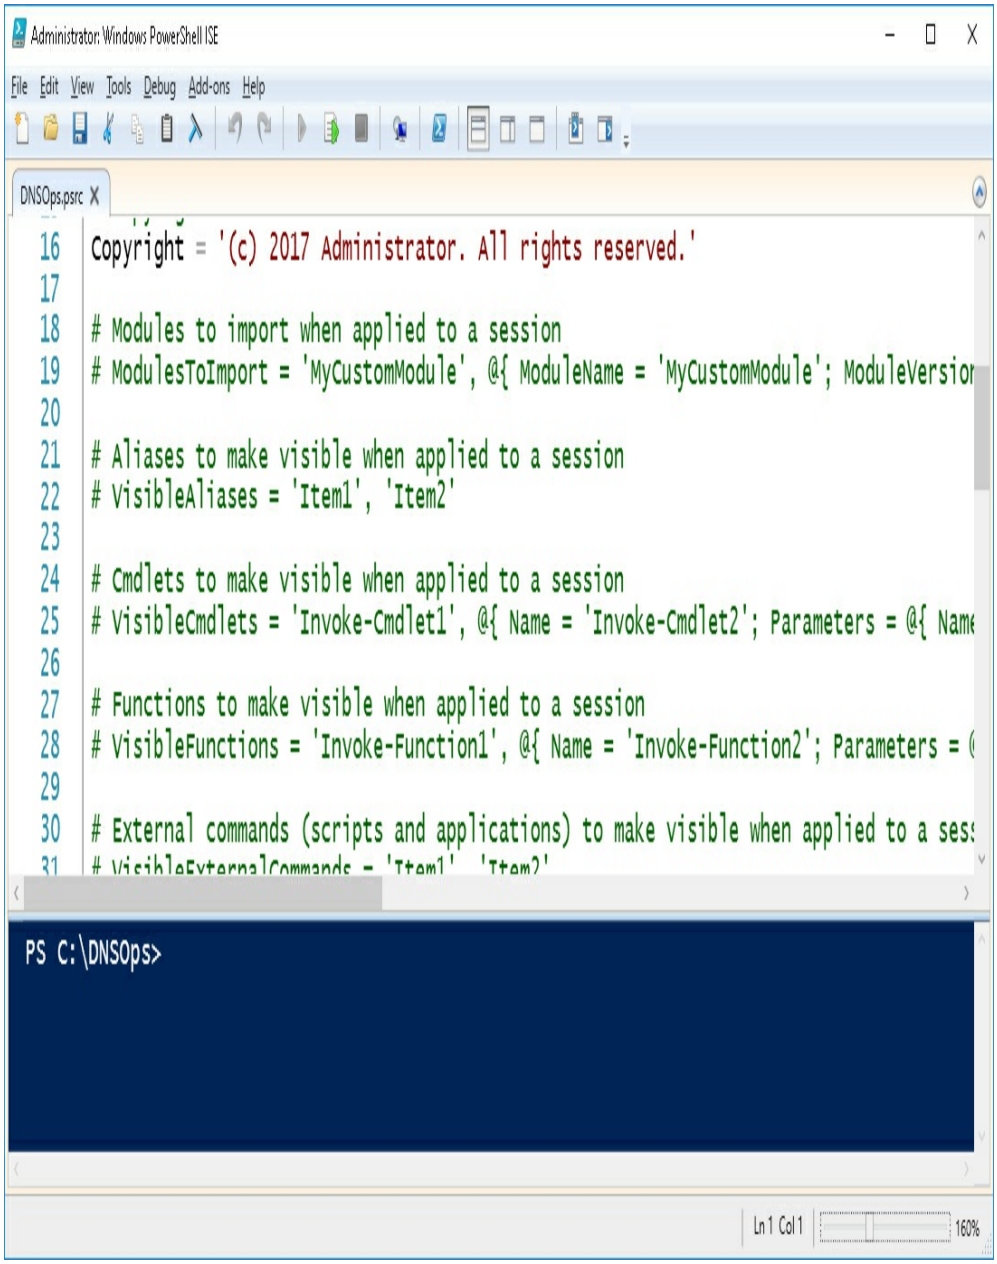


Figure 19-9 JEA role-capability file

Inside OUT

Cmdlets, functions, and aliases

Authoring role-capability files is one of those few times when you need to know whether something in PowerShell is a cmdlet or function. Mostly, people refer to

commands in PowerShell as cmdlets, but some are actually functions and others are aliases. You need to know the appropriate type when configuring a role-capability file because if you put a function in as an allowed cmdlet, you won't get the expected result. You can figure out which designation is appropriate using the Get-Command cmdlet.

Table 19-2 describes the different options that you can configure in a role-capability file.

Table 19-2 Role-capability files

Capability	Description
ModulesToImport	JEA auto-loads standard modules, so you probably don't need to use this unless you need to import custom modules.
VisibleAliases	Specifies which aliases to make available in the JEA session. Even if an aliased cmdlet is available, the alias won't be available unless it's here.

s	
Visible Cmdlets	Lists which Windows PowerShell cmdlets are available in the session. You can extend this by allowing all parameters and parameter values to be used, or you can limit cmdlets to particular parameters and parameter values. For example, use the following syntax if you wanted to allow the *Restart-Service* cmdlet to only be used to restart the DNS service: VisibleCmdlets = @{ Name = 'Restart-Service'; Parameters = @{ Name='Name'; ValidateSet = 'DNS'}}
VisibleFunctions	This field lists which Windows PowerShell functions are available in the session. You can choose to list functions, allowing all parameters and parameter values to be used, or you can limit functions to particular parameters and parameter values. For example, if you wanted to allow the *Add-DNSServerResourceRecord*, *Get-DNSServerResourceRecord*, and *Remove-DNSServerResource* functions to be used, you would use the following syntax: VisibleFunctions = 'Add-DNSServerResourceRecord', 'Get-DNSServerResourceRecord','Remove-DNSServerResourceRecord'
VisibleExternalCommands	This field allows users who are connected to the session to run external commands. For example, you can use this field to allow access to *c:\windows\system32\whoami.exe* so that users connected to the JEA session can identify their security context by using the following syntax: VisibleExternalCommands = 'C:\Windows\System32\whoami.exe'

Technet24

mmands	
VisibleProviders	This field lists Windows PowerShell providers that are visible to the session.
ScriptsToProcess	This field allows you to configure Windows PowerShell scripts to run automatically when the session is started.
AliasesD	This field allows you to define Windows PowerShell aliases for the JEA session.

ef in iti o n s	
F u n ct io n D ef in iti o n s	This field allows you to define Windows PowerShell functions for the JEA session.
V a ri a bl e D ef in iti o n s	This field allows you to define Windows PowerShell variables for the JEA session.

EnvironmentVariables	This field allows you to specify environment variables for the JEA session.
TypesToProcess	This field allows you to configure Windows PowerShell type files to load for the JEA session.
FormatsT	This field allows you to configure Windows PowerShell formats to load for the JEA session.

oProcess	
AssembliesToLoad	This field allows you to specify which assemblies to load for the JEA session.

Session-configuration files

Session-configuration files determine which role capabilities are mapped to specific security groups. For example, if you wanted to allow only members of the *CONTOS\DNSOps* security group to connect to the JEA endpoint that is defined by the DNSOps role-capability file, you would configure this in the session configuration file.

You use the *New-PSSessionConfigurationFile* cmdlet to create a session configuration file. These files use the

Technet24

.pssc extension. For example, to create a new session configuration file for the DNSOps role, run the following command:

Click here to view code image

```
New-PSSessionConfigurationFile -Path .\DNSOps.pssc -Full
```

Session configuration files include the elements described in Table 19-3.

Table 19-3 Session configuration files

Field	Explanation
SessionType	This field allows you to configure the session's default settings. If you set this to R*estrictedRemoteServer*, you can use the *Get-Command*, *Get-FormatData*, *Select-Object*, *Get-Help*, *Measure-Object*, *Exit-PSSession*, *Clear-Host*, and *Out-Default* cmdlets. The session execution policy is set to *RemoteSigned*. Example: SessionType = 'RestrictedRemoteServer'
RoleDefinitions	You use the *RoleDefinitions* entry to assign role capabilities to specific security groups. These groups do not need to have any privileges and can be standard security groups. Example: RoleDefinitions =@{'CONTOSO\DNSOps' = @{RoleCapabilities='DNSOps'}}

i *o* *n* *s*	
R *u* *n* *A* *s* *V* *ir* *t* *u* *a* *l* *A* *c* *c* *o* *u* *n* *t*	When enabled, this field allows JEA to use a privileged virtual account created just for the JEA session. This virtual account has local Administrators privileges on member servers and is a member of the Domain Admins group on a domain controller. Use this option to ensure that credentials are not cached on the server that hosts the endpoint. Remember that you can configure the virtual account to be a member of groups other than the local Administrators group.
T *r* *a* *n* *s* *c* *ri* *p* *t* *D* *ir*	This field allows you to specify the location where JEA activity transcripts are stored.

e *c* *t* *o* *r* *y*	
R *u* *n* *A* *s* *V* *ir* *t* *u* *a* *l* *A* *c* *c* *o* *u* *n* *t* *G* *r* *o* *u* *p* *s*	If you do not want the virtual account to be a member of the local Administrators group (or Domain Admins on a domain controller) you can instead use this field to specify the groups in which the virtual account is a member.

JEA endpoints

A JEA endpoint is a Windows PowerShell endpoint that

you configure so that only specific authenticated users can connect to it. When those users do connect, they only have access to the Windows PowerShell cmdlets, parameters, and values defined by the appropriate session-configuration file that links security groups and role capabilities. When you use endpoints with virtual accounts, the actual activity that a user performs on the server that hosts the endpoint occurs using the virtual account. This means that no domain-based administrative credentials are stored on the server that hosts the endpoint.

A server can have multiple JEA endpoints, and each JEA endpoint can be used for a different administrative task. For example, you could have a DNSOps endpoint to perform DNS administrative tasks and an IISOps endpoint to perform Internet Information Server related administrative tasks. Users are not required to have privileged accounts that are members of groups, such as the local Administrators group, to connect to an endpoint. Once connected, users have the privileges assigned to the virtual account configured in the session-configuration file.

Inside OUT

Non-privileged groups

JEAlets you provide IT Ops staff with administrative privileges on specific computers without adding them to traditional privileged groups, such as the Domain Admins or Enterprise Admins group.

You create JEA endpoints by using the *Register-PSSessionConfiguration* cmdlet. When using this cmdlet, you specify an endpoint name and a session-configuration file hosted on the local machine.

For example, to create the endpoint DNSOps using the *DNSOps.pssc* session-configuration file, enact the following command and then restart the WinRM service:

Click here to view code image

```
Register-PSSessionConfiguration -Name DNSOps -Path .\DNSOps.pssc
```

You can use the *Get-PSSessionConfigurationFile* to determine which endpoints are present on a computer. A user wanting to connect to a JEA session endpoint uses the *Enter-PSSession* cmdlet with the *ConfigurationName* parameter. For example, to connect to the DNSOps JEA endpoint on server MEL-DNS1, you would use the command:

Click here to view code image

```
Enter-PSSession -ComputerName MEL-DNS1 -ConfigurationName DNSOps
```

Once you've verified that JEA works, you'll need to lock down the default PowerShell endpoint. By default, only members of the local Administrators group can connect to this default endpoint, and if you've implemented JEA properly, this group shouldn't need to have very many members anyway.

More Info

Just Enough Administration

You can learn more about JEA at https://aka.ms/JEA.

Privileged Access Management

Privileged Access Management (PAM), also known as Just In Time (JIT) Administration, works on the concept that users can be granted permission for a finite amount of time rather than permanently. It uses temporary membership of a security group that has been delegated privileges, rather than permanent membership of a security group that has been delegated privileges, to accomplish this goal.

For example, a user named Rooslan requires that his password be reset and submits a ticket to Oksana, who works at the service desk, to have Oksana perform this task. In a traditional environment, Oksana would use an account that is a member of a security group that has been delegated the reset-password privilege. In a PAM environment, Oksana requests the privilege to change Rooslan's user account password. The PAM process adds Oksana's user account temporarily to a group that has been delegated the reset password privilege on Rooslan's user account. Oksana then resets Rooslan's user account password. After a certain period of time has elapsed, the PAM process removes Oksana's account

from this group, removing the reset-password privilege for Rooslan's user account from Oksana's user account.

Privileges are assigned temporarily rather than permanently. Rather than assigning the IT Operations team privileged accounts, the IT Operations team members have standard accounts that are assigned privileges only when those privileges are required. PAM can be configured so that privileges are assigned when requested, assigned only after approval has been sought and granted, or a mixture of these methods depending on which privilege is being requested. For example, requesting permission to reset the password of a non-privileged user account might occur automatically, whereas permission to add a domain controller to the domain might require approval.

PAM benefits

When properly implemented, PAM can provide the following security improvements:

- You can configure your environment so that all accounts that the IT Operations team uses are standard user accounts. Privileges are only granted to these accounts after being requested and approved, and they are temporary. If a user account used by a member of the IT Operations team becomes compromised, the attacker gains no additional privileges beyond those assigned to a standard user account because any privileges that the account might hold are only ever assigned on a temporary basis.
- All requests for privileges are logged. In high-security environments, privilege use can be further tracked, with the

activities of a user account granted temporary privileges written to a log file.

Once privileges are granted, a user needs to establish a new session by opening a new Windows PowerShell session or by signing out and signing in again to leverage the new group memberships configured for their account.

PAM components

A PAM deployment consists of the following elements:

- **Administrative forest.** This is an ESAE forest in which you have deployed Microsoft Identity Manager (MIM) 2016. Accounts from this forest are used to perform administrative tasks. Also known as the Bastion forest.
- **A production forest.** This is the forest that hosts the resources used on a daily basis by the organization.
- **Shadow principals.** A shadow principal is a copy of an account or group that exists in the production or source forest. Shadow security groups that are used in PAM use a process called SID History mirroring. SID History mirroring allows the shadow security group in the administrative forest to have the security identifier of the security group in the production or source forest. When a user is added to a PAM role, MIM adds the user's shadow account in the administrative forest to the shadow group in the administrative forest. When a user logs in by using this account, the user's Kerberos token then includes a security identifier that matches the security identifier of the original group from the production forest. This functionally makes the user a member of the group that the shadow group is mirroring.
- **PAM Client.** This client includes a collection of Windows PowerShell cmdlets or a custom solution that interacts with the

PAM REST API to allow a PAM user to query and interact with the MIM PAM functionality. The PAM client is only required on computers where it is necessary to request access to a PAM role.

- **MIM Service.** The MIM server manages the privileged account management process, including adding and removing shadow principals from shadow groups when a PAM role is granted and when it expires.
- **MIM Portal.** The MIM Portal is a SharePoint site. It provides management and configuration functionality. For example, IT Operations users can sign in to this site to request privileges.
- **MIM Service Database.** The MIM service database can be hosted on SQL Server 2012 or later. This database holds configuration and identity information that the MIM service uses.
- **PAM Component Service.** This manages the role-expiration process, removing the shadow account from the shadow group in the administrative forest.
- **PAM Monitoring Service.** This monitors the production forest and duplicates changes to the administrative forest or the MIM service.
- **PAM REST API.** This can be used to enable a custom client to interact with PAM.

PAM users and groups

You use the *New-PAMUser* cmdlet, which is available when you install Microsoft Identity Manager, to create a shadow user. When doing this, an existing PAM-created trust relationship between the ESAE forest and the production must exist. You must also specify the source domain and source account name. For example, to create a new shadow user based on the Oksana account

in the *Adatum.com* domain, you would use the following command:

Click here to view code image

```
New-PAMUser –SourceDomain Adatum.com –SourceAccountName Oksana
```

You use the *New-PAMGroup* cmdlet, available when you install Microsoft Identity Manager, to create a shadow group. When creating a shadow group, you need to specify the source group name, the source domain, and a source domain controller in the source domain. For example, to create a new shadow group for the *CorpAdmins* group in the *Adatum.com* domain by using the *mel-dc1.adatum.com* domain controller, issue the following command:

Click here to view code image

```
New-PAMGroup –SourceGroupName "CorpAdmins" –SourceDomain
adatum.com
–SourceDC mel-dc1.adatum.com
```

PAM roles

A PAM role represents the privileged role to be requested. For example, you might configure a specific PAM role for performing password resets or another for configuring Active Directory Certificate Services. You can configure PAM role settings by viewing the properties of the PAM role in the MIM web interface. A PAM role consists of the following options:

- **Display Name.** This is the name of the PAM role.
- **PAM Privileges.** This is a list of security groups to which a user who is granted access to the role is temporarily added. This does not list the privileges actually assigned to this group. For example, if you delegate the Reset User Password privilege to the *CorpAdmins* group, this dialog box does not indicate that this privilege has been delegated. This dialog box also lists the name of the shadow principal rather than the group in the original domain.
- **PAM Role TTL(sec).** This is the maximum amount of time that the member can be granted this role. The default is 3,600 seconds (1 hour). You should configure the maximum lifetime for a Kerberos user ticket in line with the Time To Live (TTL) value to ensure that a user's ticket is updated to reflect the change in group membership after the PAM Role TTL expires.
- **MFA Enabled.** MIM's PAM functionality can be integrated with Azure Multifactor Authentication. When you enable this option, requests for a PAM role require two forms of authentication. This second form of authentication can include a text message or a telephone call to a registered number. When a text message is sent, the requesting user must enter the code sent to him or her. When a telephone call is made, the user must enter a PIN that he or she has previously configured.
- **Approval Required.** PAM roles can be configured so that PAM role membership is only granted if a PAM administrator approves the request. You configure PAM administrators by adding a user to the PAM Admins set. This can be done through the Sets section under Management Policy Rules in the MIM web console.
- **Availability Window Enabled.** When you configure an availability window for a role, the PAM role can only be used during certain hours. This provides an additional layer of security for a PAM role. For example, you might configure a role that allows modification to sensitive groups to only be used during business hours.

- **Description.** Use this option to provide a description of the PAM role.

> **Inside OUT**
>
> **PAM and JEA**
>
> With JEA, you limit who can access a JEA endpoint on the basis of security group membership. With PAM, you add user accounts to security groups on a temporary rather than an ongoing basis. You can combine PAM with JEA to allow users to only connect to JEA endpoints for a limited amount of time and to only be able to do so after they have explicitly requested access to a specific JEA endpoint.

> **More Info**
>
> **Privileged Access Management**
>
> You can learn more about PAM at https://docs.microsoft.com/en-us/microsoft-identity-manager/pam/privileged-identity-management-for-active-directory-domain-services.

Local Administrator Password Solution

Local Administrator Password Solution (LAPS) is a tool that you can use to manage the passwords of local administrator accounts of each domain-joined computer in your environment. LAPS allows you to avoid the problem where a common default local Administrator account password is set for every computer in the domain and that password is set to never expire.

- Local Administrator passwords are unique on each computer managed by LAPS.
- Local Administrator passwords are randomized and changed on a regular basis.
- LAPS stores local Administrator passwords and secrets securely

within AD DS.

- Access to passwords and secrets is controlled by configurable permissions.
- Passwords retrieved by LAPS are transmitted to the client in a secure encrypted manner.

You install the LAPS agent on each computer that you wish to manage. The LAPS agent is then managed through a Group Policy Client Side Extension. Installing LAPS requires an update to the Active Directory schema. You perform this update by running the *Update-AdmPwdADSchema* cmdlet, which is included in a PowerShell module made available when you install LAPS on a computer. To run this update, the account running the cmdlet must be a member of the Schema Admins group, and this cmdlet should be run on a computer that is in the same Active Directory site as the computer that holds the Schema Master role for the forest.

When a group policy refresh occurs, LAPS does the following:

- LAPS checks whether the password of the local Administrator account has expired.
- If the password hasn't expired, LAPS does nothing.
- If the password has expired, LAPS performs the following steps:
    - Changes the local Administrator password to a new random value based on the configured parameters for local Administrator passwords.

- Transmits the new password to AD DS, where it is stored in a special confidential attribute associated with the computer account of the computer that has had its local Administrator Account password updated.
- Transmits the new password expiration date to AD DS, where it is stored in a special confidential attribute associated with the computer account of the computer that has had its local Administrator account password updated.

By default, accounts in the Domain and Enterprise Administrators group can retrieve local Administrator accounts. Domain and Enterprise Administrators also can trigger a local Administrator password change on a specific computer.

To configure computers so that their passwords are managed by LAPS, perform the following steps:

- Move computer accounts of computers that you want to use LAPS to manage passwords to an OU.
- Install LAPS on the computers for which you want to manage the local Administrator password. You'll also need to install LAPS on an administrative computer on which you can use to install the necessary template.
- Install the GPO templates into AD DS. You can do this by running the installer and selecting the PowerShell module and the GPO Editor templates installation option.
- Use the *Set-AdmPwdComputerSelfPermission* cmdlet to assign computers in an OU the ability to update the password of their local Administrator account when it expires. For example, for computers in the *Adelaid* OU, run the command:

Click here to view code image

```
Set-AdmPwdComputerSelfPermission -Identity “Adelaide”
```

Once installed, you’ll be able to configure the following policies:

- **Enable Local Admin Password Management.** This policy enables LAPS and allows you to manage the local Administrator account password centrally.
- **Password Settings.** This policy allows you to configure the complexity, length, and maximum age of the local Administrator password. The default is that upper and lowercase letters, numbers, and special characters are used. The default password length is 14 characters, and the default password maximum age is 30 days.
- **Do Not Allow Password Expiration Time Longer Than Required.** When enabled, the password is updated according to domain password expiration policy.
- **Name Of Administrator Account To Manage.** Use this policy to identify custom local Administrator accounts.

To view local Administrator passwords, you can use the client that comes with the LAPS installer, the *Get-AdmPwdPassword* cmdlet, or view a computer's *ms-Mcs-AdmPwd* attribute.

More Info

Local Administrator Password Solution

You can learn more about LAPS at https://docs.microsoft.com/en-us/previous-versions/mt227395(v=msdn.10).

Advanced auditing

There are two sets of audit policies in a Group Policy

Object (GPO): traditional audit policies and advanced audit policies. The traditional audit policies are located in the *Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Audit Policies* node. These are the audit policies that have been available with the Windows Server operating system since Windows 2000. The drawback of these policies is that they are general, and you can't be specific in the way you configure auditing. When you use these policies, you not only audit the events that you're interested in, but you also end up auditing many events that you don't need to know about.

The advanced audit policies enable you to be more specific in the types of activity you audit. The advanced audit policies are located under the *Computer Configuration\Policies\Windows Settings\Security Settings\Advanced Audit Policy Configuration* node.

There are 10 groups of audit policy settings and 61 individual audit policies available through Advanced Audit Policy Configuration. The audit policy groups contain the following settings:

- **Account Logon.** You can audit credential validation and Kerberos-specific operations.
- **Account Management.** You can audit account management operations, such as changes to computer accounts, user accounts, and group accounts.
- **Detailed Tracking.** You can audit encryption events, process

creation, process termination, and RPC events.

- **DS Access.** You can audit Active Directory access and functionality.
- **Logon/Logoff.** You can audit logon, logoff, and other account activity events, including IPsec and Network Policy Server (NPS) events.
- **Object Access.** You can audit access to objects including files, folders, applications, and the registry.
- **Policy Change.** You can audit changes to audit policy.
- **Privilege Use.** You can audit the use of privileges.
- **System.** You can audit changes to the security subsystem.
- **Global Object Access Auditing.** You can configure expression-based audit policies for files and the registry.

Expression-based audit policies

Traditional object audit policies involve specifying a group and configuring the type of activities that trigger an event to be written to the security log. For example, you can specify that an audit event is written each time a member of the Managers group accesses a file in a specific folder.

Expression-based audit policies enable you to go further. These policies enable you to set conditions for when auditing might occur. For example, you might want to configure auditing so that tracking only occurs for members of the *Managers* group who have access to sensitive files when they access those files from computers that aren't part of the *Managers_Computers*

group. This way, you don’t bother tracking access when members of this group access sensitive files from within the office, but you do track all access to those sensitive files when members of this group are accessing them from an unusual location.

You can integrate expression-based audit policies with Dynamic Access Control (DAC) to create targeted audit policies based on user, computer, and resource claims. Instead of just adding claims based on user or device group membership, the claim can be based on document metadata, such as confidentiality settings and site location. You can configure expression-based audit policies at the file or folder level or apply them through Group Policy by using policies in the Global Object Access Auditing node of Advanced Audit Policy Configuration.

File and folder auditing

After you configure object access auditing either through traditional or advanced audit policies, you can configure auditing at the file and folder level. The simplest way to configure auditing is at the folder level because you can configure all folders and subfolders to inherit those auditing settings. If you change the auditing settings at the folder level, you can use the Replace All Child Object Auditing Entries option to apply the new auditing settings to the folder’s child files

and folders.

You can configure auditing for a specific file and folder through the Advanced button on the Security tab of the object's properties. You can configure basic success and failure auditing. You can also configure expression-based auditing so that a specific security group member's activity is audited only if other conditions, such as membership of other security groups, are also met.

The advantage of using Global Object Access Auditing is that when you configure it, you can use file classification to apply metadata to files and then automatically have auditing enabled for those files. For example, by using file classification and DAC, you can configure a Windows Server 2019 file server so that all files containing the phrase "code secret" are marked as Sensitive. You can then configure Global Object Access Auditing so that all attempts to access files marked as Sensitive are automatically audited. Instead of having an administrator track down all the files that are sensitive and configuring auditing on those files, the process is automatic. All that needs to happen to trigger the audit is that the phrase "code secret" is included in the file.

Using auditpol.exe with auditing

Auditpol.exe is a command-line utility that you can use to configure and manage audit policy settings from an

elevated command prompt. You can use *auditpol.exe* to perform the following tasks:

- View the current audit policy settings with the */Get* subcommand
- Set audit policy settings with the */Set* subcommand
- Display selectable policy elements with the */List* subcommand
- Back up and restore audit policies using the */Backup* and */Restore* subcommands
- Delete all per-user audit policy settings and reset the system policy settings using the */Clear* subcommand
- Remove all per-user audit policy settings and disable all system policy settings using the */Remove* subcommand

For example, to enable success and failure auditing for the File System subcategory of Object Access, execute the following command:

Click here to view code image

```
Auditpol.exe /set /subcategory:"File System" /success:Enable
/failure:Enable
```

To view the current audit policy settings for all audit policies, issue the following command:

Click here to view code image

```
Auditpol.exe /get /category:*
```

To view the current audit policy settings for a specific category, such as Object Access, issue the following command:

Click here to view code image

```
Auditpol.exe /get /category:"Object Access"
```

Windows Firewall with Advanced Security

Windows Firewall with Advanced Security (WFAS) is a bidirectional, host-level firewall. This means that not only can you set it to protect a host from incoming traffic, but you can also have it block all or specific outgoing traffic. Generally, you can control outbound traffic only on very sensitive hosts. Controlling outbound traffic minimizes the chance of malware infecting the server that is phoning home. On sensitive hosts, you should allow outbound traffic on a per-application rule. For example, you should create a rule to allow a specific browser to communicate with the Internet rather than just allowing outbound traffic on port 80.

WFAS allows for the creation of rules based on program, port, or predefined rules. There are 45 predefined rules, which include categories such as Remote Desktop, Windows Firewall Remote Management, File and Printer Sharing, and Core Networking. One of the useful things about WFAS is that it generally enables the appropriate firewall rules when you enable a specific role, role service, or feature. These default rules are appropriate most of the time, but you may wish to edit the properties of these rules to make them more secure.

Firewall profiles

Firewall profiles allow you to treat traffic differently depending on the network to which the computer is connected. For example, the same server might be connected to the following:

- A remote site through a VPN connection that has the private profile applied
- A perimeter network by a network adapter, where the connection has the public profile applied
- The organization's internal network through a second adapter that has the domain profile applied

The advantage of separate profiles is that each one allows you to have separate sets of firewall rules applied. You might allow Distributed File System (DFS) traffic on one profile, web traffic on another, and SMTP traffic on a third.

Profiles are independent of each other; the settings that you configure for one profile do not apply to other profiles. You can configure the following properties for each profile:

- **Firewall State.** You can set this to On (default) or Off.
- **Inbound Connections.** You can set this to Block (default), Block All, or Allow. The Block All setting means firewall rules that allow connections are ignored.
- **Outbound Connections.** You can set this to Block or Allow (default).
- **Settings.** This allows you to configure whether notifications are

displayed, whether unicast responses are transmitted to multicast or broadcast traffic, and whether to merge rules when rules apply both locally and through Group Policy.

- **Logging.** This allows you to specify logging settings.

Firewall logs are written to the *%systemroot%\system32\LogFiles\Firewall\pfirewall.log* file. The default settings do not log dropped packets or successful connections, which means that, unless you change the defaults, nothing is logged. Firewall logging is most useful when you've determined that a particular service is unavailable to the network because of the firewall.

You can turn on firewall logging on a per profile basis by clicking the Customize button next to Logging on each profile's properties page. This displays the Customize Logging Settings dialog box shown in Figure 19-10. You can then enable logging for dropped packets and successful connections. If you are having trouble with a newly installed service on a server, enable the logging of dropped packets to determine the properties of the packets dropped, so that you can create a firewall rule that allows your service to be accessed from the network.

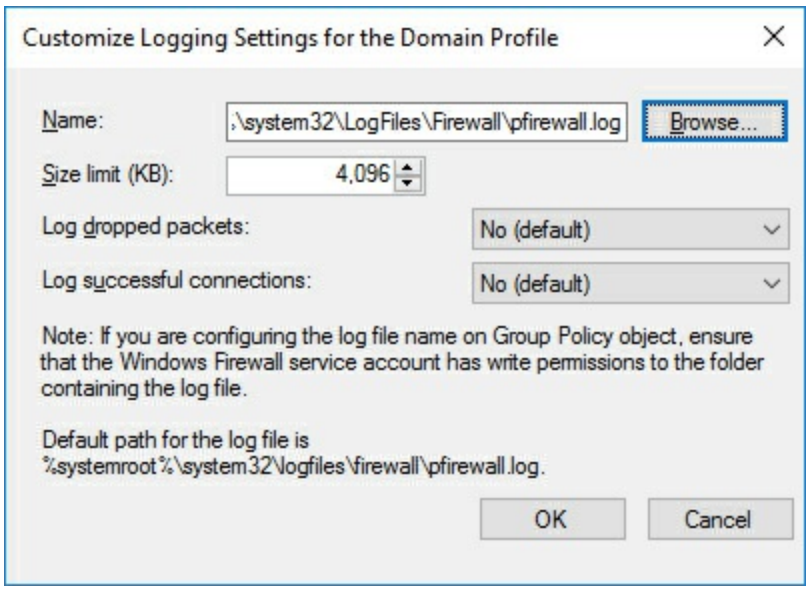


Figure 19-10 Custom logging settings

Inbound rules

Inbound rules are based on program, port, or one of 45 predefined categories, such as BranchCache Content Retrieval or Network Policy Server. You can also create custom rules that include a mixture of these, where you specify a program and a port as well as a rule scope. For example, you can use a custom rule to block all incoming traffic on port 80 to a particular application but not block port 80 traffic to other applications on the server. The basic aspects of creating a firewall rule involve:

- **Specifying a program or port.** When you specify a port, you must choose whether the rule applies to TCP or UDP traffic.
- **Specifying what action should be taken.** You can allow the connection, in which case all traffic that matches the rule is allowed

Technet24

by the firewall. You can allow the connection if it is authenticated, in which case traffic that meets the IPsec authentication requirements and the connection security rules is allowed, but traffic that is not properly authenticated according to these conditions is dropped.

- **Specifying the network profiles in which the rule applies.** In general, rules should apply in all profiles, but there might be circumstances for which you want to allow traffic from an interface connected to a domain network but block the same traffic if it comes from an interface connected to a public network.

After you have created the rule, you can then edit the rule's properties. Editing the rule's properties allows you to configure more advanced options than are present in the Rule Creation Wizard. By editing a rule's properties, you can:

- Configure a rule to apply to a service rather than just a program.
- Limit the computers that can make authenticated connections. By default, if you configure the Allow Traffic If The Connection Is Authenticated option, any computer that can authenticate is able to successfully transmit traffic. By editing a firewall's rules, you can limit traffic to specific computers rather than all authenticated computers. You can do the same with user accounts, limiting successful connections to specific users when authentication has successfully occurred.
- Edit the rule's scope, which is the local and remote IP address ranges to which the rule applies. You can also do this when you create a custom rule.
- Configure whether the rule applies to specific network interfaces, instead of just network profiles. For example, if your computer has two network adapters, you can configure a rule to apply so that it allows traffic on one adapter but not the other when both adapters have the same profile set.

- Configure whether to allow packets that have passed across a Network Address Translation (NAT) device.

Creating outbound rules

Default settings for Windows Firewall with Advanced Security do not block outbound traffic. In high-security environments, you should consider using outbound rules to block all but authorized outbound traffic.

The process for creating an outbound rule is almost identical to the process for creating an inbound rule. You specify the rule type on the New Outbound Rule Wizard, you specify whether the connection is blocked, allowed, or allowed only if authenticated, and you specify the network profiles in which the rule is active. By editing the properties of the rule after the rule is created, you can configure all the advanced options configurable for inbound rules, such as rule scope, specific network interfaces, and to which computers or users the rule allows outbound connections.

Configuring IPsec

The IPsec Settings tab of Firewall Properties allows you to configure how IPsec is used when applied in connection security rules. On the tab itself, shown in Figure 19-11, you can configure whether you want to customize the IPsec defaults, exempt Internet Control Message Protocol (ICMP) traffic from IPsec, and whether you want to configure IPsec tunnel

Technet24

authorization.

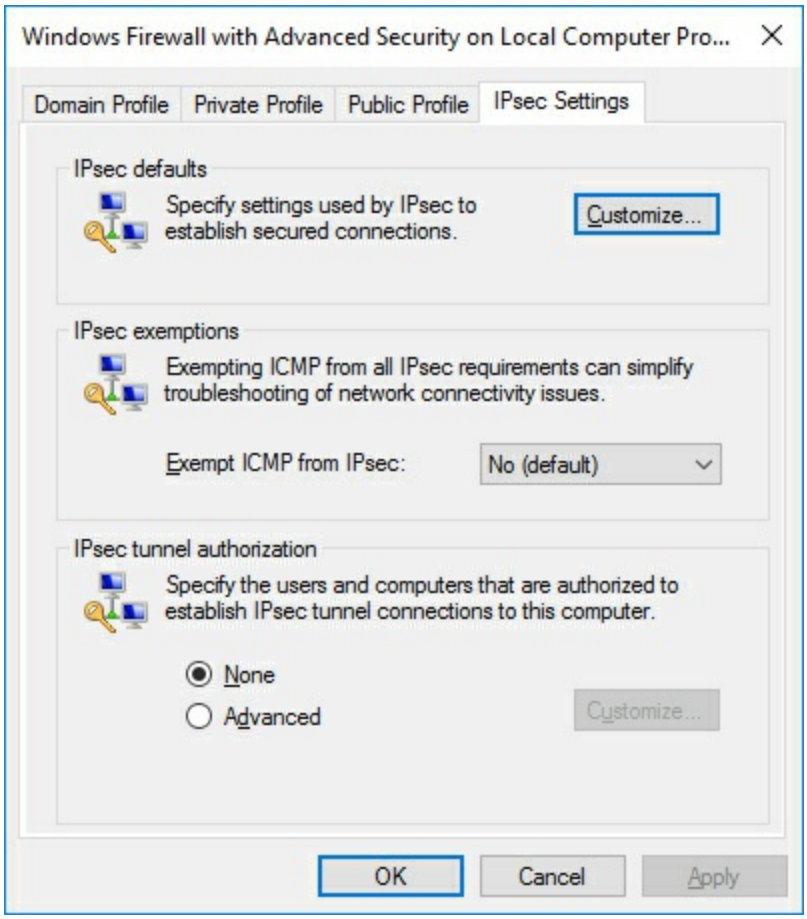


Figure 19-11 IPsec settings

Exempting ICMP traffic can be useful for diagnostic purposes because many administrators use the ping utility to diagnose whether a host has network connectivity. If connection security rules are enabled, the default IPsec settings mean that a successful IPsec negotiation must occur prior to an ICMP response being sent. If the IPsec negotiation is the problem, enabling ICMP response allows you to verify that there is

network connectivity and that the problem just lies a bit further up the network stack.

If you click the Customize button next to the IPsec defaults, you can change the Key Exchange, Data Protection, and Authentication Method (see Figure 19-12). The default key exchange uses Diffie-Hellman Group 2 key exchange algorithm with a key lifetime of 480 minutes. You can modify these settings so that they use more secure methods, but this may reduce the ability to interact with some third-party devices.

Technet24

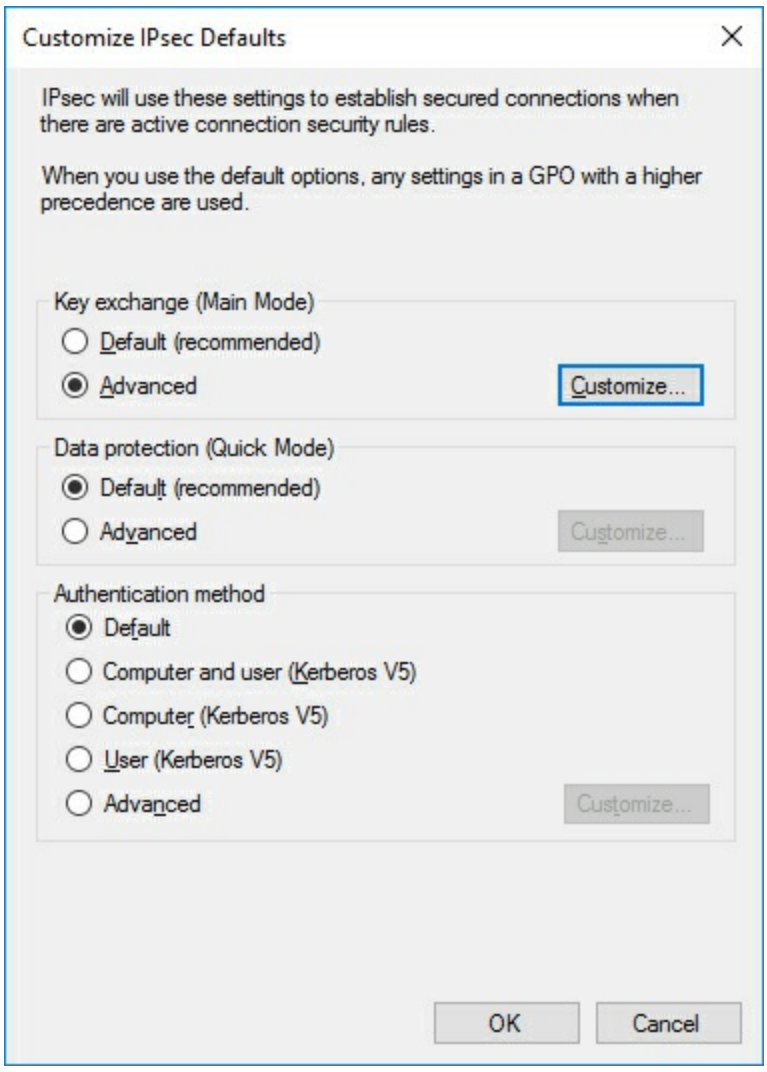


Figure 19-12 Customize IPsec defaults

The data protection settings allow you to configure the algorithms used for data integrity and encryption. Normally, there isn't much reason to change this; however, if you feel you need the strongest encryption possible to protect your organization's network traffic, you can use a different algorithm, such as AES-CBC with a 256-bit key length. In general, the stronger the

encryption, the greater the resources needed to support that encryption. Although you can make the protection a lot stronger, the benefit of doing so might be only marginal given the value of the data being protected.

The default authentication method used for IPsec connections is Computer (Kerberos V5). This means that a domain controller must be present to verify the identity of each computer before an IPsec session can be established. You can also use Computer (NTLMv2), a computer certificate from this certification authority (CA), or a pre-shared key to authenticate IPsec connections. A pre-shared key is not a recommended method of authentication, but it might be necessary if there is not a certificate services infrastructure and computers aren't members of a domain.

Connection security rules

Connection security rules are a more intelligent type of firewall rule than a simple port-based rule. An interesting adaption to firewalls has been an increase of traffic through ports traditionally used for common applications. Connection security rules are about trusting incoming connections based on the identity of the remote host rather than the remote host's network address or the communication port it uses.

Authentication exemptions Authentication exemptions allow you to create very specific holes in

connection security rules. These are useful in the event that you want to remotely manage a server, but all the servers that would normally authenticate your connection, such as certificate servers or domain controllers, are unavailable. Authentication exemptions override existing connection security rules.

When you create an authentication exemption, you define the exemption on the basis of the source computer's IP address. As Figure 19-13 shows, you do this using a single IPv4 or IPv6 address, an IP address range, or a predefined set of computers. For security reasons, you should limit the number of IP addresses for which you create authentication exemptions to one or two specific privileged access workstations.

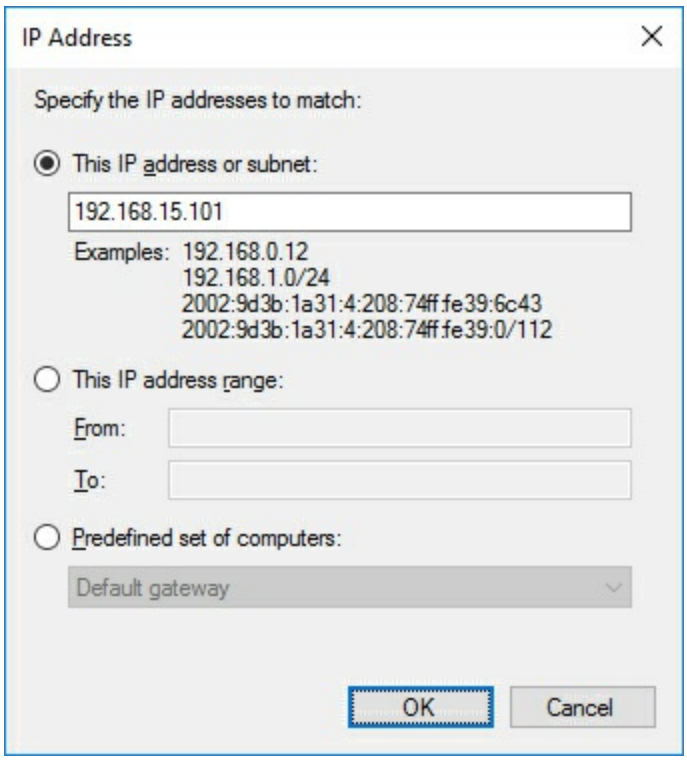


Figure 19-13 Authentication exemption

To create an authentication exemption, perform the following steps:

1. In Windows Firewall With Advanced Security, right-click the Connection Security Rules node and then click New Rule. This opens the New Connection Security Rule Wizard. Select Authentication Exemption.
2. In the Exempt Computers dialog box, click Add. In the IP Address dialog box, enter the IP address, subnet, or IP address range of the computers that you wish to exempt from connection security rules. You can also choose from a list of predefined addresses, including the Default Gateway, DHCP servers, the local subnet, or DNS servers.
3. Select the profile to which you want the rule to apply and specify a rule name.

Isolation rules Isolation rules allow you to limit communication so that a server or computer communicates only with computers that have authenticated with it. Sometimes, these are called domain isolation policies because they limit hosts to only communicating with other servers that are members of the same Active Directory forest. When you configure an isolation rule, it applies to all traffic from all hosts; this is unlike tunnel or server-to-server rules, which apply to traffic to and from specific hosts.

To create an isolation rule, perform the following steps:

1. In Windows Firewall With Advanced Security, right-click the Connection Security Rules node and then click New Rule. This opens the New Connection Security Rule Wizard. Select Isolation and then click Next.
2. On the Requirements page, as shown in Figure 19-14, choose one of the following:
    1. Request Authentication For Inbound And Outbound Connections
    2. Require Authentication For Inbound Connections And Request Authentication For Outbound Connections
    3. Require Authentication For Inbound And Outbound Connections
3. Requiring authentication on both inbound and outbound connections is the strongest form of protection because it means that communication can be performed only with authenticated hosts.
4. Choose the authentication method. The default is to use the authentication method specified on the IPsec Settings tab of Firewall Properties. You can also choose between computer and user leveraging Kerberos version 5, or you can configure advanced authentication options, which include certificate-based authentication. It is also

possible, though not recommended, to use a pre-shared key for authentication, which is useful when you need to communicate with computers running third-party operating systems.

5. The final step in setting up a connection security rule is to configure the profiles in which the rule applies. Unless there is a good reason to do otherwise, you should apply connection security rules in all profiles.

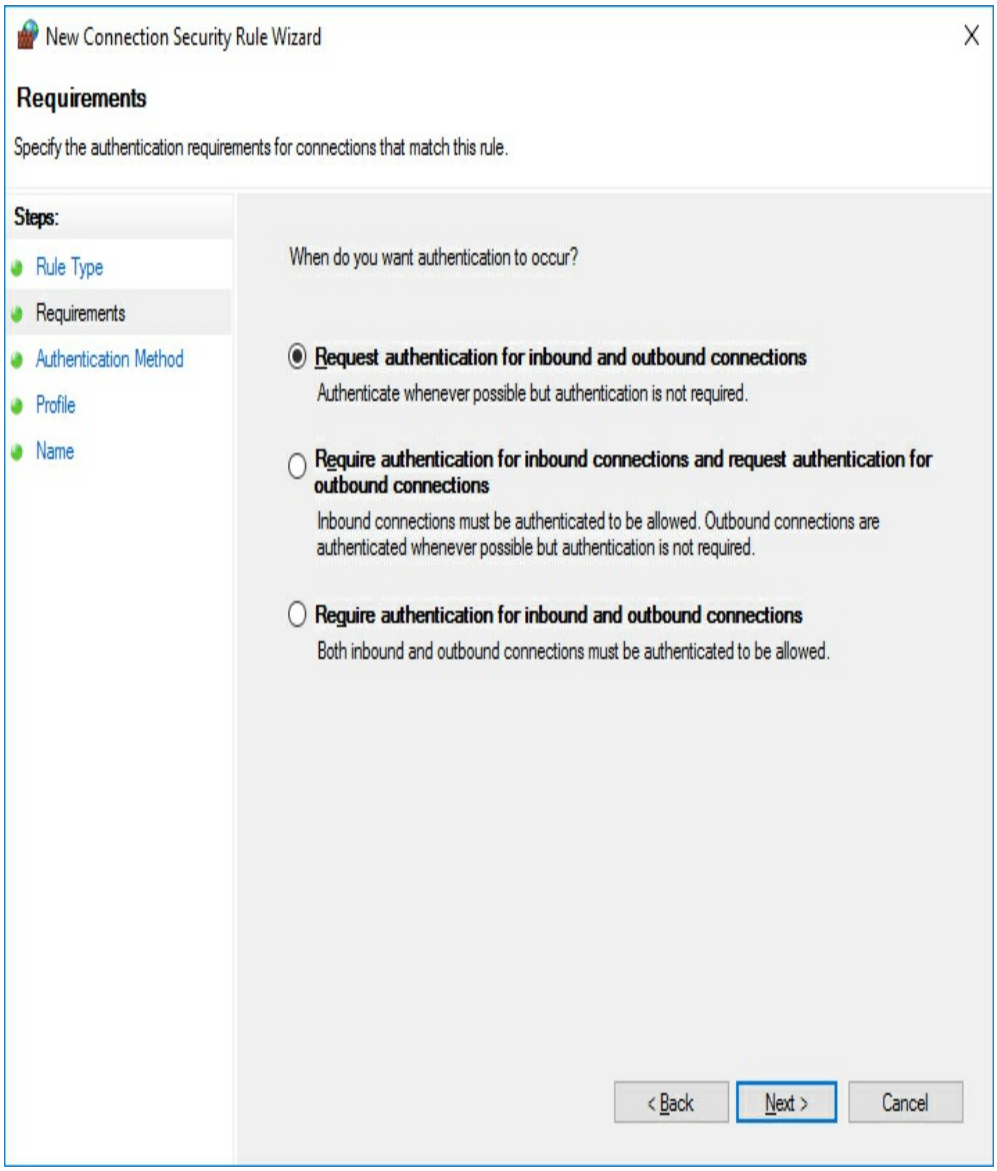


Figure 19-14 Connection security rule

Server-to-server rules Server-to-server rules are used to authenticate communication between two groups of computers. This can be a rule authenticating and encrypting a single computer-to-computer connection, such as between a web server and a database server or between computers on two separate subnets. Server-to-server rules differ from isolation rules; isolation rules apply to communication from all hosts, whereas server-to-server rules apply to specific hosts.

To create a server-to-server rule, perform the following steps:

1. In Windows Firewall With Advanced Security, right-click the Connection Security Rules node and then click New Rule. This opens the New Connection Security Rule Wizard. Select the Server-To-Server rule.
2. In the Endpoints dialog page, enter the IP addresses of the computers that are at one end of the connection and the IP addresses of the computers that are at the other end of the connection. Figure 19-15 shows a rule defining one endpoint as the subnet 192.168.15.0/24, and the other endpoint as subnet 192.168.16.0/24.

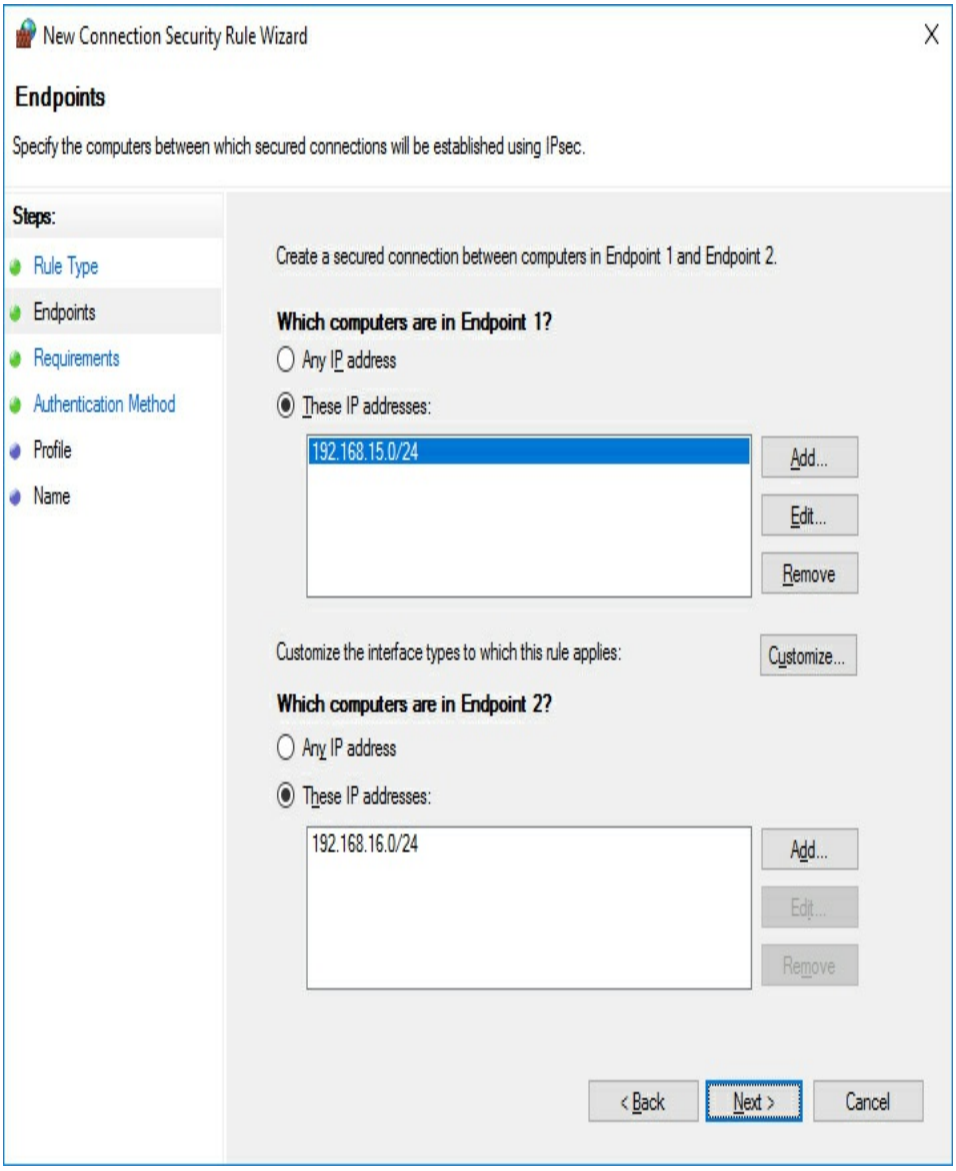


Figure 19-15 Connection security rule endpoints

3. On the Requirements page, specify how you want authentication to occur. If you want only this computer to communicate on the profile using encrypted and authenticated connections, select Require Authentication For Inbound And Outbound Connections.

4. Specify the authentication method. The default is to use a computer certificate issued by a designated CA. You can also specify computer-based Kerberos, NTLMv2, or a pre-shared key.

5. Specify the firewall profiles in which this connection security rule applies and add a name for the rule.

Tunnel rules Tunnel rules allow client computers to communicate with computers on a secure network behind a remote gateway. For example, if you have a single server located at a branch office that you want to connect to an internal network at another office, using a tunnel rule, you specify the location of a host that functions as a gateway to that secure network. This allows you to create an IPsec tunnel through which secure communication can occur.

To create a tunnel rule, perform the following steps:

1. In Windows Firewall With Advanced Security, right-click the Connection Security Rules node and then click New Rule. This opens the New Connection Security Rule Wizard. Select Tunnel Rule.
2. On the Tunnel Type page, as shown in Figure 19-16, determine the type of tunnel that you want to create. To create the type of rule outlined above, select Client-To-Gateway.

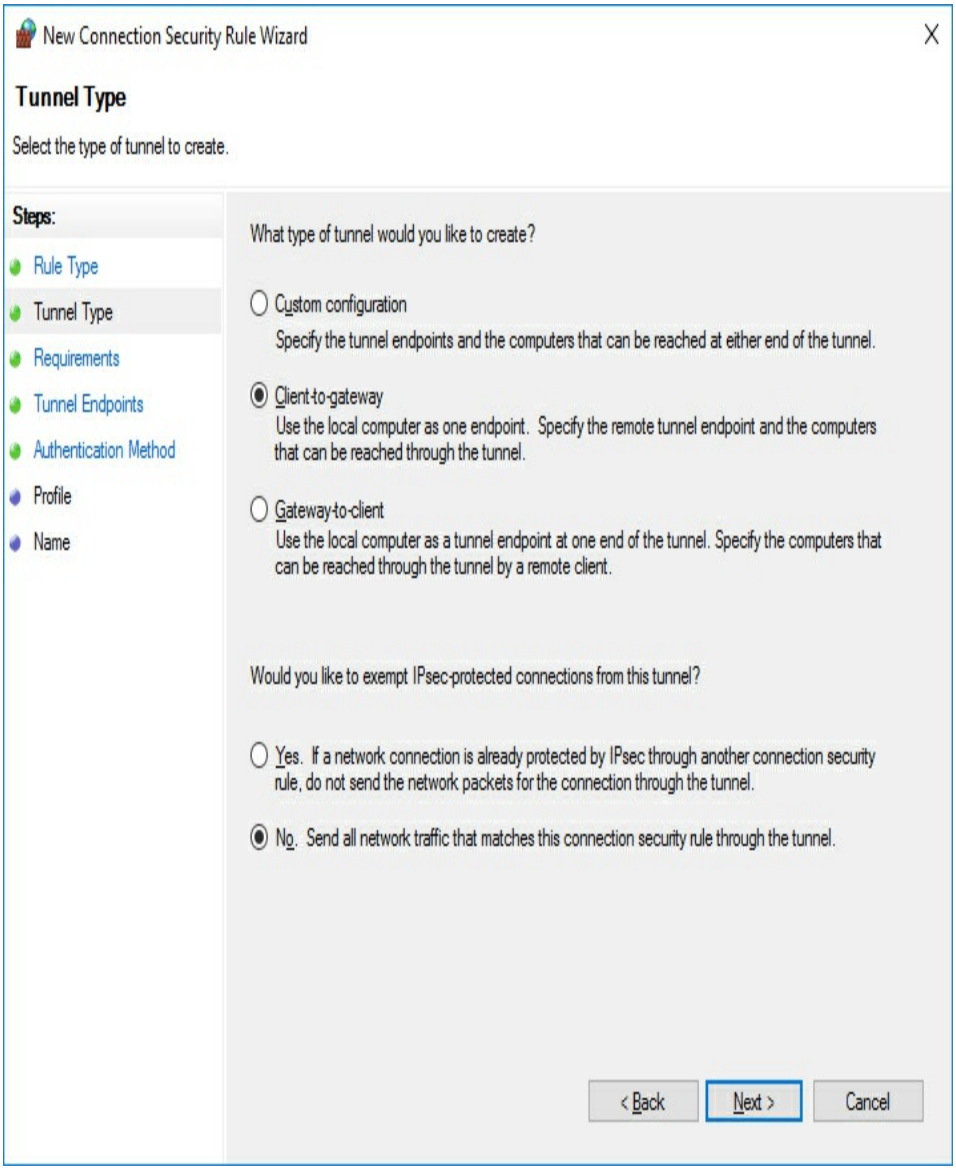


Figure 19-16 Tunnel Type

3. Choose whether you wish to require authentication for inbound and outbound connections.
4. Specify the address of the computer that functions as the gateway to the secure network.
5. Specify the authentication method. The default is to use a computer certificate issued by a designated CA. You can also specify computer-

based Kerberos, NTLMv2, or a pre-shared key.

6. Specify the firewall profiles in which this connection security rule applies and add a name for the rule.

> **More Info**
>
> **Windows Firewall with Advanced Security**
>
> You can learn more about deploying and managing Windows Firewall with Advanced Security, even though it deals with a previous version of Windows Server, at https://docs.microsoft.com/en-us/windows/security/threat-protection/windows-firewall/windows-firewall-with-advanced-security.

SHIELDED VMS

Shielded VMs are a special type of virtual machine configuration that are protected against attacks originating from the virtualization fabric. A virtualization fabric attack is one that comes from the infrastructure that hosts the virtual machine itself. Attacks that come from the virtualization infrastructure include:

- **Virtual machine export attacks.** Rather than compromise a running virtual machine, an authorized administrator can simply use Hyper-V's built in export functionality to export a virtual machine. Once exported, administrators can then examine the contents of the exported virtual machine at their leisure without any evidence existing of their access. Anything that can be done by an authorized administrator can also be done by a nefarious third party who has gained access to the administrator's account or has local Administrator access on Hyper-V host. For example, someone with the appropriate privileges can export a virtualized domain controller, extract the *ntds.dit* Active Directory file, and then run an offline brute force analysis of that file to extract username and

password data. They could then use that username and password data to gain administrative privileges and access that domain and its resources in a manner that seemed completely legitimate.

- **Virtual machine movement attacks.** Rather than exporting a virtual machine, an attacker might attempt to move a shielded VM to an untrusted host. The intention in doing so would be to remove the protection associated with the VM. However, with shielded VMs, if the new host is not recognized as trusted based on an identifier related to the host's TPM, the shielded VM does not function and the contents of the VM remain protected.

Inside OUT

Exfiltrate all things

In the past, you could ensure the physical security of a server by putting it in a locked cage in a highly secure datacenter that was also protected by security guards and might even be surrounded by a piranha filled moat. Someone trying to steal the server had to be able to get past the piranhas, the security guard, and the locked cage. Today, the vast majority of servers are virtualized and are functionally just a bunch of very large files. Rather than getting past the piranhas, stealing a server is mostly a case of uploading the files that comprise the virtual machine to a location you control somewhere on the Internet. Shielded VMs provide a way of protecting against this type of attack.

When you are thinking about shielded VMs, you separate the owner and operator of the VM—called the tenant—from the owner and operator of the virtualization fabric. (In most cases, this is you as the Hyper-V virtualization administrator.) Shielded VMs allow VM operators to run their workloads without the concern that anyone with local Administrator privileges on the Hyper-V host server can access the data and applications stored on that virtual machine.

Shielded VMs include the following security features and requirements:

- **Secure Boot must be present and enabled.** This protects the boot environment of the VM from attacks that might modify the environment to introduce malware. This malware could be leveraged to gain access to the virtual machine.
- **Virtual TPM must be present and enabled.** Virtual TPMs function in a manner similar to hardware-based TPMs in that they provide a secure location for the storage of cryptographic keys and secrets.
- **VM must have BitLocker enabled.** BitLocker protects the integrity of the boot environment and ensures that the VM's virtual hard disks are encrypted. If the virtual hard disk file was copied, it would be inaccessible unless the person attempting to access it had the necessary BitLocker recovery keys.
- **Only certain integration components are allowed.** Tools such as PowerShell Direct and Data Exchange are blocked. Blocking these integration services components blocks a Hyper-V administrator from indirectly accessing data stored in the virtual machine.
- **Virtual Machine Connection (console) cannot be enabled for the VM.** Access must be through remote access methods enabled within the VM operating system, such as PowerShell remote sessions or Remote Desktop. The only way to access the content of the VM is where the VM operating system has granted access to that content.
- **Restricted keyboard and mouse access.** Human Interface Devices, such as keyboards and mice, are disabled for the VM and can't be used to interact with it. For example, this means that an administrator couldn't interrupt the VM's boot process or use another attack to inject code into the VM in the way they might do so when attacking a physically deployed computer.

- **COM/serial ports.** Access through other virtual ports to the VM is disabled. This blocks a nefarious administrator using out-of-band management tools to access the VM.
- **Debugger access.** A debugger cannot be attached to the VM process on a shielded VM.

The following operating systems support being run as shielded VMs:

- Windows Server 2019, Windows Server 2016, Windows Server 2012 R2, and Windows Server 2012
- Windows 10, Windows 8, and Windows 8.1

Shielded VMs can only be hosted on specially configured Windows Server 2016 and Windows Server 2019 Hyper-V servers that are part of a shielded fabric. Shielded VMs must use virtual hard disks in VHDX format that have the following properties:

- Configured as GPT disks.
- Must be basic rather than dynamic disks.
- Disk that hosts operating systems must be configured with at least two partitions. One partition stores the Windows operating system; the other stores the active partition, which includes the bootloader.
- Volumes on the disk must be formatted as NTFS.

While it's possible to convert an existing VM to a shielded VM, Microsoft recommends that shielded VMs be deployed from scratch from trusted installation media or digitally signed VM templates.

More Info

Shielded VMs

You can learn more about shielded virtual machines at https://docs.microsoft.com/en-us/windows-server/security/guarded-fabric-shielded-vm/guarded-fabric-configuration-scenarios-for-shielded-vms-overview.

Guarded fabric

Before you can deploy a shielded VM, you must configure a guarded fabric on which to deploy the VM. A guarded fabric is a special cryptographically secure virtualization host environment. A guarded fabric is made up of one Host Guardian Service (HGS), usually a cluster of three nodes, a number of guarded virtualization hosts, and the shielded VMs that run on those hosts.

Host Guardian Service

The Host Guardian Service (HGS) is a Windows Server 2019 role that manages attestation and the Key Protection Service (KPS). The KPS stores and is responsible for securely releasing keys to shielded VMs. One of the keys that the KPS stores is the transport key. The transport key is used to encrypt a shielded VM's virtual TPM. Microsoft recommends that you deploy the HGS role on a physically secure physical server running the Server Core operating system.

Validating a Hyper-V host involves verifying both its identity and configuration. Once the attestation service verifies both its identity and configuration, the KPS

provides the transport key to that host that is required to unlock and run shielded VMs on that host.

Configuring a HGS server involves the deployment of a dedicated single-domain AD DS forest. An HGS cluster comprises at least three nodes, and each HGS cluster node is a domain controller for the root domain of this dedicated AD DS forest. Deploying an HGS cluster is different from a traditional domain controller deployment, and you don't go through the traditional AD DS promotion process. Instead, you take a standalone server, install the HGS role on that server, and then run the *Initialize-HgsServer* PowerShell cmdlet.

Running the *Initialize-HgsServer* PowerShell cmdlet triggers the process of configuring the server as a domain controller for the newly created single-domain AD DS forest. The trusted Hyper-V hosts need to either be a member of the HGS forest or a member of a domain that has a one-way trust relationship with the forest. This second configuration must be even more secure because a compromised virtualization host won't be able to make any modifications to the HGS forest because the trust relationship goes only in the opposite direction.

TPM and admin-trusted attestation

Trusted attestation allows a shielded VM to verify that it

is running on a virtualization host that participates in a guarded fabric. There are two methods that you can use to provide trusted attestation. The method that you provide depends on the hardware that you have available. Table 19-4 lists the different trusted attestation types.

Table 19-4 Attestation types

Attestation type	Functionality and benefits
TPM-trusted	• Provides strongest protection. • Requires TPM 2.0 and UEFI 2.3.1 with secure boot. • Guarded hosts are approved based on TPM ID, measured boot sequence, and code integrity policies.

	• Each guarded fabric virtualization host's TPM must be registered with the guardian service. To do this, use the *Get-PlatformIdentifier* cmdlet to generate a special file containing the TPM endorsement key's public and private keys, and then use the *Add-HgsAttestationTpmHost* cmdlet in registering the TPM with the guardian service. • It's necessary to generate a baseline code-integrity policy for each virtualization host hardware SKU. This policy allows validation of the integrity of the software and hardware environment of each guardian fabric virtualization host. • The best practice is to use Hardware Security Module to store HSM-backed certificates, though this is not mandatory.
Host key attestation	• Doesn't require TPM 2.0. • Compatible with common server hardware. • Guarded hosts are approved through possession of a cryptographic key. • A guarded host forwards the request to the HGS using an encrypted REST API.

Technet24

Admin-trusted	• Deprecated from Windows Server 2019 onward. • Doesn't require TPM 2.0. • Compatible with common server hardware. • Guarded hosts are approved by Host Guardian Service based on membership of specific AD DS security group. • Guarded host forwards the request to the HGS using an encrypted REST API. • Use the *Add-HgsAttestationHostGroup* cmdlet to authorize the specific security group's SID with the attestation service. • Guarded fabric virtualization hosts must have a one-way trust relationship to HGS domain.

More Info

Guarded fabric

You can learn more about guarded fabric at https://docs.microsoft.com/en-us/windows-server/security/guarded-fabric-shielded-vm/guarded-fabric-and-shielded-vms#attestation-modes-in-the-guarded-fabric-solution.

Encryption-supported VMs

An encryption-supported VM is a VM that can run in a semi-protected manner on a guarded fabric but doesn't have the stringent security requirements of a shielded VM. A virtualization administrator has a greater level of access to an encryption-supported VM. This makes the deployment less secure than a shielded VM. Encryption-supported VMs may be more appropriate for organizations where the VM administrator would normally have access to the VMs anyway, but for compliance reasons VMs need to be run in a manner in which the contents of virtual hard disks are encrypted.

Table 19-5 outlines how encryption-supported VMs differ from shielded VMs.

Table 19-5 Differences between encryption-supported and shielded VMs

Security setting	Encryption-supported VM	Shielded VM
Secure Boot	Required but configurable	Required and enforced
vTPM	Required but configurable	Required and enforced

Technet24

Encrypted VM state	Required but configurable	Required and enforced
Live migration traffic encryption	Required but configurable	Required and enforced
Integration components including PowerShell Direct	Configurable by fabric administrator with PowerShell Direct and Data Exchange allowed	Data Exchange and PowerShell Direct blocked
Virtual Machine Connection	Allowed	Disabled (Can be enabled with Windows Server version 1803 and later)
Keyboard, mouse, HID devices	Allowed	Disabled
COM/Serial ports	Supported	Disabled
Attach debugger to VM process	Supported	Disabled

Shielding data file

The shielding data file (also called the provisioning data

file, or PDK file) is a special encrypted file that stores the following secrets:

- Administrator credentials.
- Answer file (*unattend.xml*).
- Security policy that determines whether VMs created using this shielded data file will be fully shielded VMs or encryption-supported VMs.
- An RDP certificate. This allows secure RDP communication with the VM.
- A volume signature catalog. This contains a list of trusted and signed template disk signatures from which a VM can be created. A shielded template disk is a template virtual hard disk that has a digital signature. This ensures that the template hard disk from which a new VM is deployed is trustworthy. During provisioning, the signature of the disk is computed and checked against a list of trusted and shielded template disk signatures stored in the catalog. If the signatures match, deployment continues. If the signatures do not match—indicating that the shielded template disk has been modified—the deployment fails.
- A Key Protector. This specifies which guarded fabrics a shielded VM generated from this shielding data file is authorized to run on.

Tenant keys protect this file, and the tenant uploads the shielding data file to the secure fabric. Because tenant keys protect the file, the virtualization administrator doesn't have access to the contents of the shielding data file. This means that the virtualization administrators do not have access to the local Administrator password used with virtual machines deployed on the fabrics that they manage.

The shielding data file ensures that any VMs are explicitly created with the parameters specified by the tenant. The virtualization administrator can't swap in their own template hard disks, answer files, or administrator passwords. Should a template hard disk be swapped out or modified by the virtualization administrator, it won't match the cryptographic signature of any authorized template hard disk listed in the shielding data file.

Chapter 20

Security systems and services

Security Compliance Toolkit

Attack Surface Analyzer

Credential Guard

Windows Defender Application Control

Virtualization-based security

Controlled Folder Access

Exploit Protection

Windows Defender

Windows Defender SmartScreen

Windows Server 2019 includes a variety of new and existing technologies that you can use to secure the operating system. While implementing each of these technologies increases the security of your Windows

Technet24

Server 2019 deployment, it's important to understand implementing no single technology or set of technologies guarantees that your deployment can't be compromised by the most determined attackers. By implementing each technology, you will make Windows Server just a bit more difficult to compromise, thereby increasing the security of your organization's Windows Server deployment.

Inside OUT

Assume breach

Assume breach is a security philosophy in which you design a security implementation around the hypothesis that an attacker already has compromised and has privileged access to at least some of your organization's protected resources. You should not rely upon the assumptions that hostile actors only exist outside the perimeter network and that critical protected systems have not already been compromised. You should design your organization's security to mitigate the risk of a privileged insider acting maliciously, either deliberately or because their credentials have been compromised by a nefarious, external third party. Don't assume that just because a host sits on a protected internal subnet that it won't be subject to malicious network intrusion attempts.

SECURITY COMPLIANCE TOOLKIT

The Security Compliance Toolkit (SCT) is a collection of tools that allow you to manage security configuration baselines for Windows Server and other Microsoft products. The SCT includes the Policy Analyzer tool and the Local Group Policy Object (LGPO) tool.

SCT provides security baselines for:

- Windows Server 2019

- Windows Server 2016
- Windows Server 2012 R2
- Windows 10
- Microsoft Office
- Microsoft Edge

> More Info
>
> **Security Compliance Toolkit**
>
> You can learn more about the Security Compliance Toolkit at https://docs.microsoft.com/en-us/windows/security/threat-protection/security-compliance-toolkit-10.

Policy Analyzer tool

The Policy Analyzer tool allows you to compare Group Policy Objects (GPOs). You can use the Policy Analyzer tool to perform the following tasks:

- Determine which existing GPO settings are redundant or are internally consistent
- Determine the differences between different versions or collections of group policies
- Compare GPO settings against a computer's local policy and registry settings

When downloading the SCT, you can also download security baseline GPOs. To perform an analysis of a system against a baseline, perform the following steps:

1. Open the Policy Analyzer and click Add to add downloaded baseline policies. You can specify a folder that contains multiple exported GPOs. When you do this, all exported GPOs under the selected folder path will

be imported. Figure 20-1 shows the Policy File Importer with policies related to Windows Server selected for importing. After selecting which policies to import, click Import and save the selection of files as a policy rule.

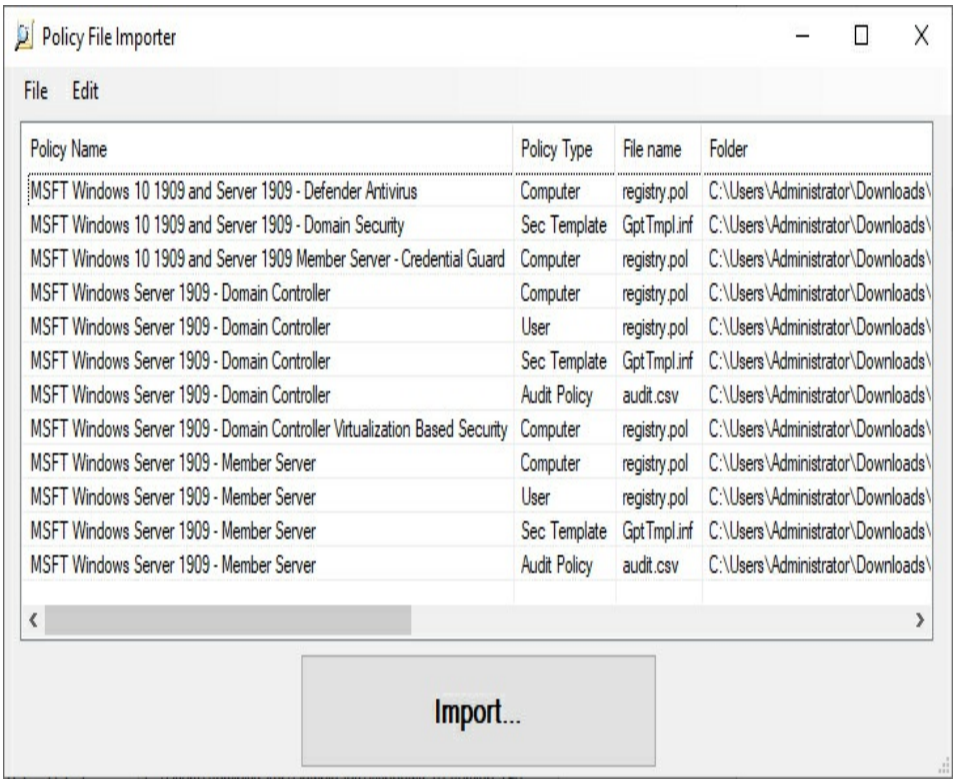


Figure 20-1 Policy File Importer

2. You can then use the View/Compare button to compare policies and determine where they differ, as shown in Figure 20-2.

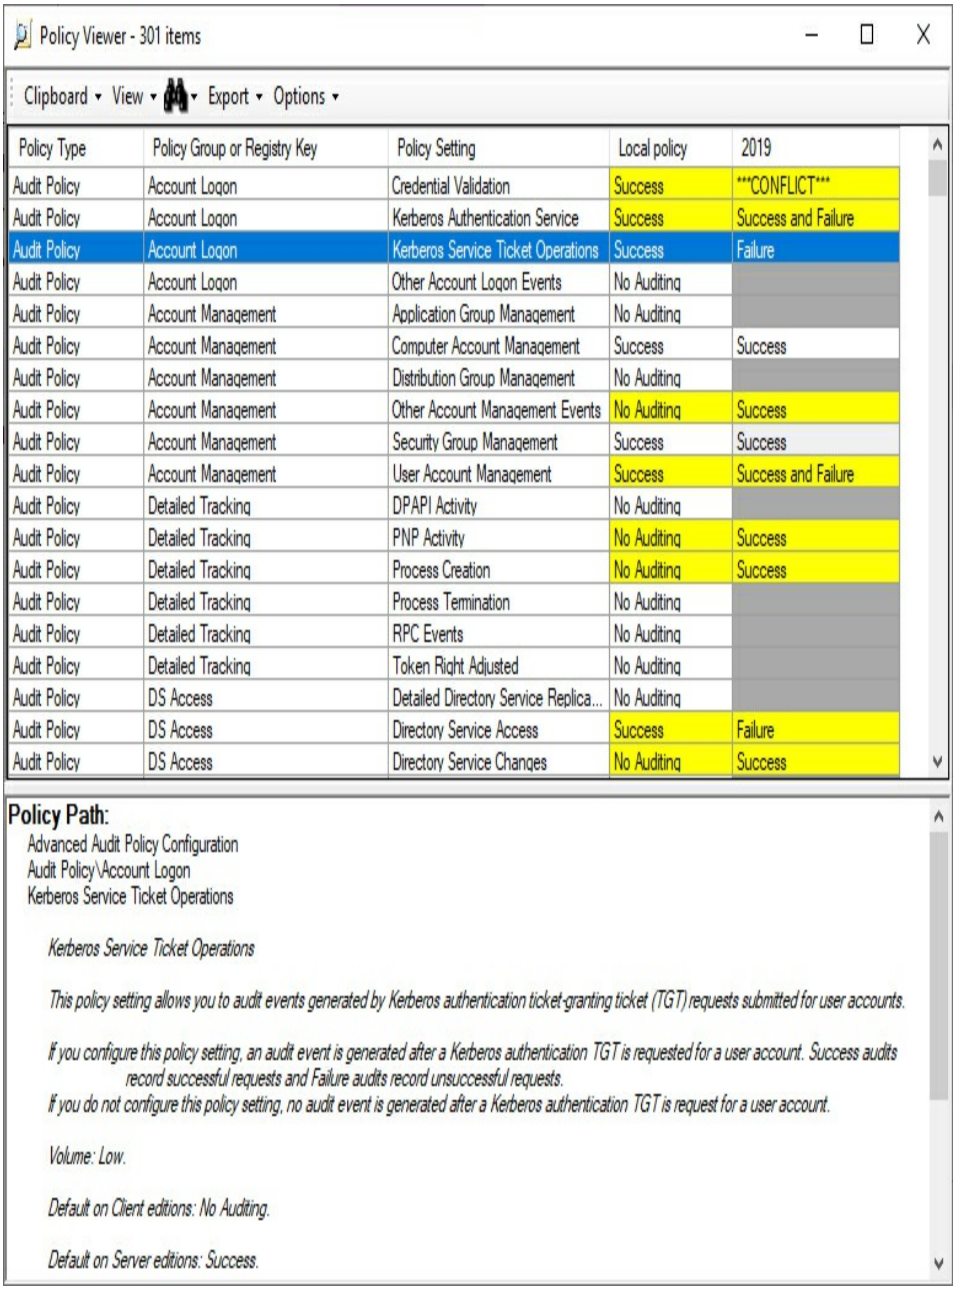


Figure 20-2 Security Configuration Toolkit Policy Viewer

You can use the Policy Analyzer to export this report into a format that can be viewed in Excel.

The Policy Analyzer is a reporting tool, and you cannot

use it to directly apply policies. To apply policies, you can manually update GPO policies within your environment incrementally until they match the secure baseline. Taking an incremental approach allows you to test whether there are any unexpected consequences by applying new policies.

You should avoid simply importing the security baseline policies into Active Directory and applying them unless you want to spend the next few weeks or months figuring out exactly which new policy setting you applied caused problems with your organization's existing workloads.

Local Group Policy Object tool

The local GPO tool is a command-line utility that allows you to perform local group policy operations against domain-joined and nondomain-joined computers. You can use the local GPO tool to perform the following tasks:

- Import settings into a computer's local group policy store from GPO backups
- Import settings into a computer's local Group Policy store from component files including Registry Policy (*registry.pol*), security templates, and advanced auditing CSV files
- Export a computer's local policy settings to a GPO backup
- Enable the use of Group Policy client-side extensions when processing local policy

- Extract a Registry Policy (*registry.pol*) file into a readable text format that can then be edited and built into a new *registry.pol* file with different settings

ATTACK SURFACE ANALYZER

The Attack Surface Analyzer is a tool that allows you to locate possible security vulnerabilities by tracking changes made to the following:

- File System
- User Accounts
- System Services
- Network Ports (listeners)
- System Certificate Stores
- Windows Registry

To use the Attack Surface Analyzer, you first scan a system to determine the baseline settings of the system. You then perform another scan after you install software or suspect an unauthorized change has been made. The Attack Surface Analyzer will generate a report detailing any modifications made between the baseline and subsequent scans.

Once you've downloaded the Attack Surface Analyzer's files from the project's GitHub page, you run the Attack Surface Analyzer by opening an elevated command prompt and running the following command, which launches a web server that you can connect to at the

address *http://127.0.0.1:5000*, as shown in Figure 20-3:

```
Asa.exe gui
```

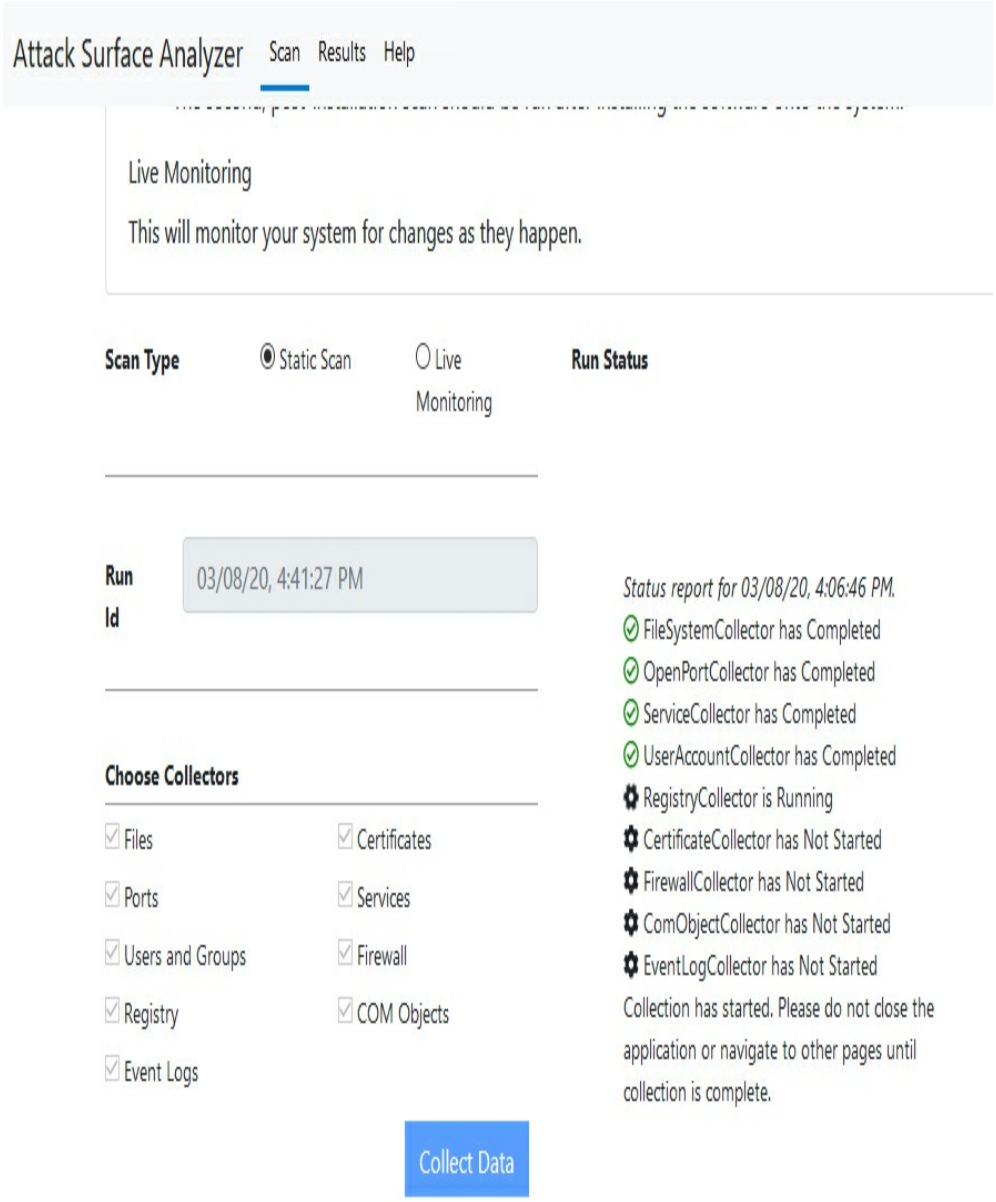


Figure 20-3 Attack Surface Analyzer

More Info

Attack Surface Analyzer

You can learn more about the Attack Surface Analyzer at https://github.com/microsoft/attacksurfaceanalyzer.

CREDENTIAL GUARD

Credential Guard allows you to leverage virtualization-based security to isolate secrets, such as cached user credentials, in a special separate virtualized operating system. The special separate virtualized operating system is configured so that only specific processes and memory in the host operating system can access this secret data. The processes running in the separate virtualized operating system are termed *trustlets*.

Credential Guard is primarily a response to pass-the-hash or pass-the-ticket attacks. Should a host that has credential guard be compromised by an attacker, that attacker won't be able to successfully run a pass-the-hash attack tool to extract cached credentials and then use them to access other computers on the network.

Credential Guard includes the following features and solutions:

- Stores derived domain credentials in a virtualized environment that is protected from the running operating system.
- You can manage Credential Guard by using Group Policy, Windows Management Instrumentation (WMI), or Windows PowerShell.

Credential Guard does not allow:

- Unconstrained Kerberos delegation
- NT LAN Manager version 1 (NTLMv1)
- Microsoft Challenge Handshake Authentication Protocol (MS-CHAPv2)
- Digest Authentication
- Credential Security Support Provider (CredSSP)
- Kerberos DES encryption

Credential Guard can be used in conjunction with the Protected Users group in a layered approach to the protection of highly privileged accounts. The Protected Users group remains useful because your organization might not have computers that support Credential Guard. You can deploy Credential Guard only on computers that meet certain hardware requirements. Credential Guard is primarily useful for Privileged Access Workstations, but you should implement it eventually on any computer where IT operations personnel use privileged domain accounts.

Credential Guard has the following requirements:

- Windows Server 2019, Windows Server 2016, or Windows 10 Enterprise
- UEFI firmware version 2.3.1 or higher
- Secure Boot
- Intel VT-x or AMD-V virtualization extensions
- Second Level Address Translation
- x64 processor architecture

- A VT-d or AMD-Vi IOMMU input/output memory management unit
- TPM 1.2 or 2.0
- Secure firmware update process
- Firmware updated to support Secure MOR implementation

To enable Credential Guard on an appropriately configured computer, you need to configure the Turn On Virtualization Based Security policy, which is located in the *Computer Configuration\Administrative Templates\System\Device Guard* node of a GPO. This is the same policy that you also use to configure Windows Defender Application Control (previously known as Device Guard), which you learn about later in this chapter.

While configuring this policy, you must first set the policy to Enabled, and then you must set the platform security level to either Secure Boot or to Secure Boot And DMA Protection. Secure Boot And DMA Protection ensures that Credential Guard is used with Direct Memory Access protection.

Once this is done, you need to then set the Credential Guard Configuration option to Enabled With UEFI Lock or Enabled Without Lock. If you set the Enabled With UEFI Lock, Credential Guard cannot be remotely disabled and can only be disabled by having someone with local Administrator privileges sign on and disable

Technet24

Credential Guard configuration locally. The Enabled Without Lock option allows Credential Guard to be remotely disabled.

More Info

Credential Guard

You can learn more about Credential Guard at https://docs.microsoft.com/en-us/windows/security/identity-protection/credential-guard/credential-guard.

Inside OUT

Don't enable Credential Guard on domain controllers

Microsoft does not support enabling Credential Guard on domain controllers. This is because the domain controller hosts authentication services that must integrate with processes that would be isolated by virtualization-based security when credential guard is enabled. If you do this, crashes might occur.

WINDOWS DEFENDER APPLICATION CONTROL

Windows Defender Application Control (WDAC), known in previous versions of Windows Server as Device Guard and Configurable Code Integrity (CCI) policies, is a hardware- and software-based security system that restricts the execution of applications to those that are explicitly trusted. WDAC uses virtualization-based security to isolate a special service named the code integrity service from the Windows kernel. Because the code integrity service is running as a trustlet in a special virtualized location, compromising the service is

difficult, if not impossible, even if an attacker has complete control of the operating system. The WDAC name isn't present in all locations of the operating system, such as group policy, so occasionally you will still see references to Device Guard in items such as Group Policy.

WDAC includes the following features:

- **Virtual Secure Mode.** This is a special virtual container that isolates the *LSASS.exe* process from the operating system.
- **Configurable Code Integrity.** This is the rules engine that, in conjunction with Virtual Secure Mode Protected Code Integrity, validates code that Windows Server attempts to enact.
- **Virtual Secure Mode Protected Code Integrity.** This uses two components to enforce Configurable Code Integrity policy:
  - User Mode Code Integrity manages whether user mode code can execute.
  - Kernel Mode Code Integrity manages whether kernel mode code can execute.
- **Platform and UEFI Secure Boot.** This ensures that the boot loader code and firmware are validated and prevents malware from modifying the boot environment.

To enable Virtual Secure Mode, perform the following steps:

1. Enable Secure Boot and Trusted Platform Module (TPM).
2. Enable Hyper-V.
3. Enable Isolated User Mode.
4. Configure the Turn On Virtualization Based Security Policy, located in

the *Computer Configuration\Administrative Templates\System\Device Guard* node.

5. Configure the BCD store to start Virtual Secure Mode. You do this by running the following command:

   Click here to view code image

   ```
   bcdedit /set vsmlaunchtype auto
   ```

By default, WDAC runs in Audit Mode. This allows you to tune a policy to ensure that the software that you want to run can run and won't be blocked. You have several options when it comes to controlling which software can run on a computer protected by Device Guard. These are as follows:

- **Only allow software that is digitally signed by a trusted publisher to run on the server.** If you are using internally written code, you can use your organization's code signing certificate to digitally sign code. You use the *New-CIPolicy* cmdlet to create an XML file with all relevant details about signed files on a system, and then use the *ConvertFrom-CIPolicy* cmdlet to convert this XML file into a binary file that is placed in the *C:\Windows\System32\CodeIntegrity* folder and can be used by Device Guard to determine which software can run on the protected system.
- **Use Package Inspector to create a catalog of all deployed and executed binary files for applications that you trust.** You can use this when you need to deal with third-party applications that you trust but are not digitally signed. Package Inspector is included with the operating system and is located in the *C:\Windows\System32* directory. Package Inspector creates a catalog of hash files for each binary executable. Once the catalog file is created, you can digitally sign the catalog using *signtool.exe*,

which is available in the Windows Software Development Kit (SDK) that you can download from Microsoft's website. Ensure that the signing certificate is added to the code integrity policy. WDAC blocks any software that is not in the signed catalog from running. You can use Group Policy and Configuration Manager to deploy code integrity policies and catalog files.

You can deploy WDAC on a test machine using the Device Guard and Credential Guard Readiness Tool. (The tool still uses the previous name for WDAC.) This tool assesses a computer's readiness to be configured for WDAC and allows you to perform a test deployment. Test deployments are important because you don't want to deploy WDAC and then find that business-critical software can no longer execute.

Once you've verified that a computer is ready for WDAC, you configure the Turn On Virtualization Based Security policy, located in the *Computer Configuration\Policies\Administrative Templates\System\Device Guard* section of a GPO and configure the policy so that the Virtualization Based Protection of Code Integrity is set to either Enabled With UEFI Lock, or Enabled Without Lock. The difference is that if you use the Enabled With UEFI Lock, as shown in Figure 20-4, the policy can only be disabled by signing in directly to the server.

Technet24

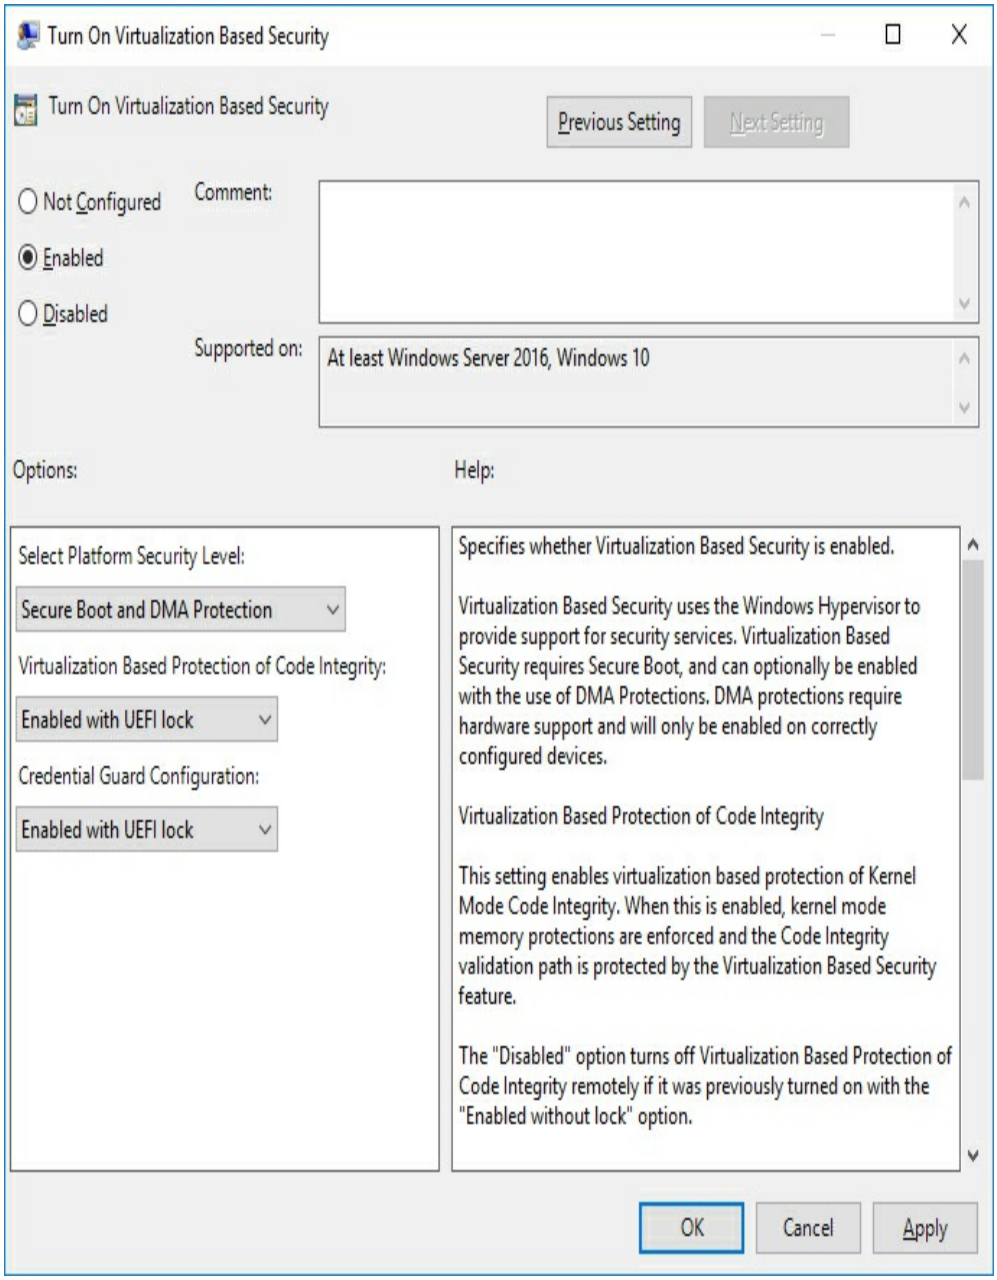


Figure 20-4 Turn On Virtualization Based Security

More Info

WDAC Deployment Guide

You can learn more about deploying WDAC at https://docs.microsoft.com/en-us/windows/security/threat-protection/windows-defender-application-

control/windows-defender-application-control-deployment-guide.

VIRTUALIZATION-BASED SECURITY

Virtualization-based security allows Windows Server to use hardware virtualization to isolate and protect an area of computer memory from typical operating system processes. This “virtual secure mode” blocks malicious code from inserting itself into other high integrity processes. You can enable virtualization-based security in the Windows Security app, as shown in Figure 20-5.

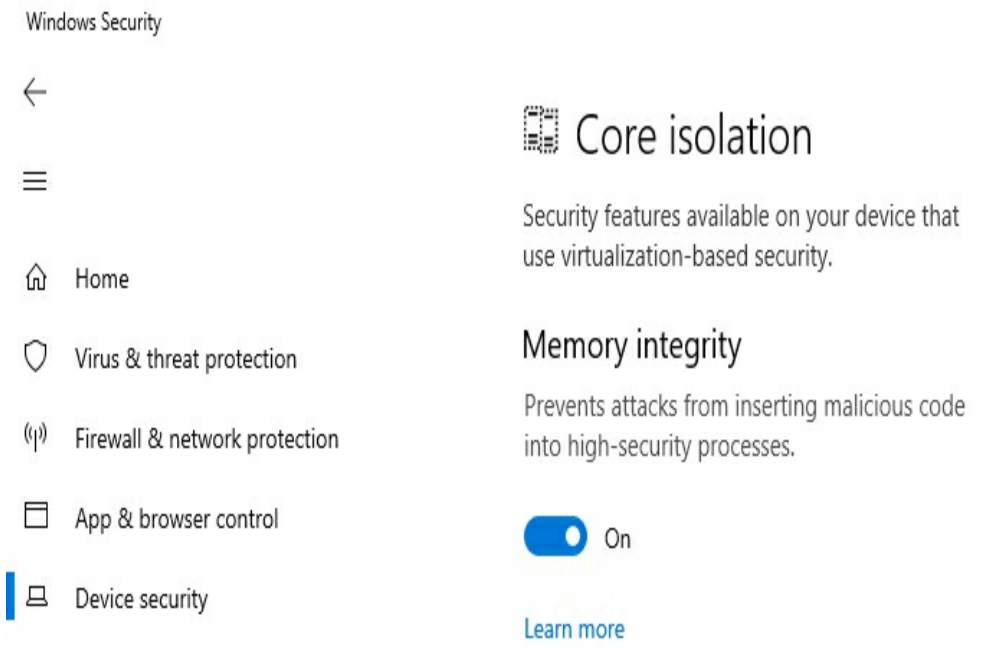


Figure 20-5 Enable virtualization-based security

More Info

Virtualization-based security

You can learn more about virtualization-based security at https://docs.microsoft.com/en-us/windows-hardware/design/device-experiences/oem-vbs.

CONTROLLED FOLDER ACCESS

Controlled Folder Access is a security technology that allows you to protect specific folders on a Windows Server computer against malicious software. Controlled Folder Access is a useful tool for preventing ransomware from encrypting the folders that host file shares because you can use it to restrict which software can interact with specific paths on a computer running Windows Server 2019. Controlled Folder Access will allow trusted applications, such as those that are included with the operating system, to interact with protected folders.

Inside OUT

Protect file servers from ransomware

Controlled Folder Access can be used to protect file servers from ransomware. You should enable it on both file servers as well as any other computer that has access to the file share. You should also ensure that your file servers are regularly backed up to a safe location, such as an Azure Recovery Services Vault, so that you are able to recover data in the event that the ransomware is clever enough to get past other precautions you have put in place.

On computers running Windows Server with Desktop Experience, you can use the Virus & Threat Protection area of the Windows Security control panel to configure

which folders are protected and which applications can interact with protected folders, as shown in Figure 20-6.

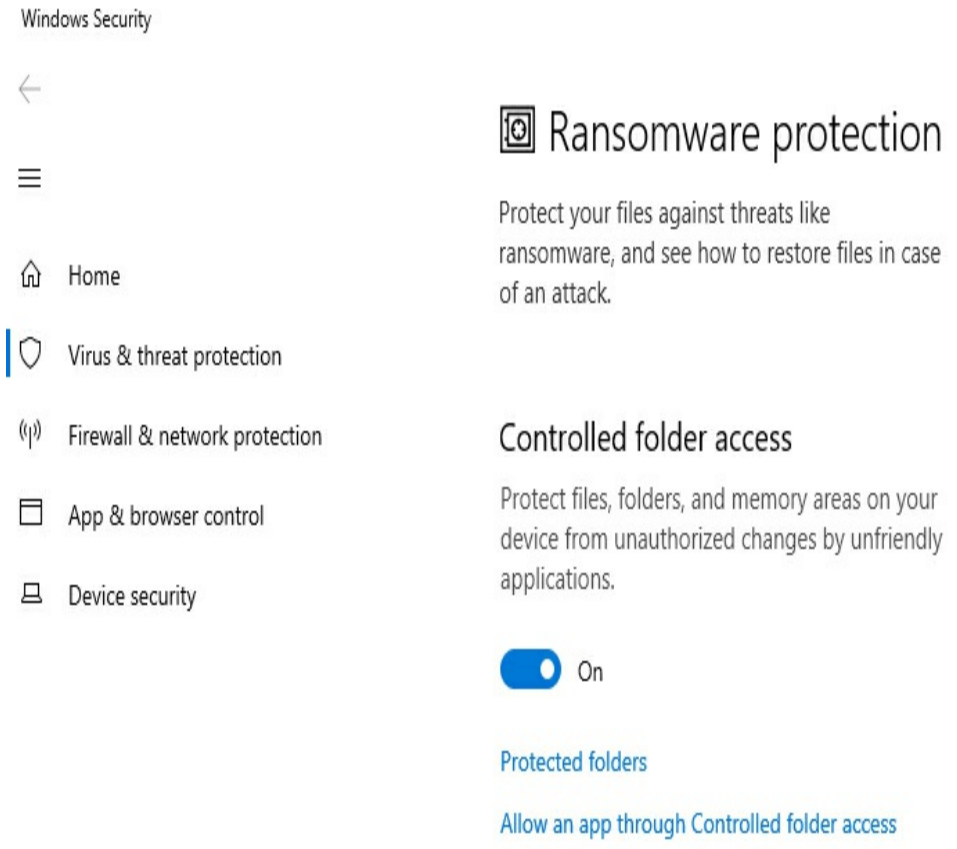


Figure 20-6 Controlled Folder Access

You can configure Controlled Folder Access from PowerShell using the *Set-MpPreference* cmdlet. To enable controlled folder access, run this command:

Click here to view code image

```
Set-MpPreference -EnableControlledFolderAccess Enabled
```

To disable Controlled Folder Access, run this command:

Click here to view code image

```
Set-MpPreference -EnableControlledFolderAccess Disabled
```

To add a new location for Controlled Folder Access to monitor, run this command:

Click here to view code image

```
Add-MpPreference -ControlledFolderAccessProtectedFolders
"E:\fileshare"
```

To remove a location from Controlled Folder Access monitoring, run this command:

Click here to view code image

```
Remove-MpPreference -ControlledFolderAccessProtectedFolders
"E:\fileshare"
```

To allow an application to interact with a protected folder, run this command:

Click here to view code image

```
Add-MpPreference -ControlledFolderAccessAllowedApplications
"c:\application\app.exe"
```

You can also use the *Remove-MpPreference* cmdlet to revoke an application's ability to interact with a protected location.

More Info

Controlled Folder Access

You can learn more about Controlled Folder Access at https://docs.microsoft.com/en-us/windows/security/threat-protection/microsoft-defender-atp/evaluate-controlled-folder-access.

EXPLOIT PROTECTION

Exploit Protection allows you to configure extra security settings for Windows Server, such as Control Flow Guard (CFG) and Data Execution Prevention (DEP). In past versions of Windows Server, you would use the Enhanced Mitigation Experience Toolkit (EMET), a separate product that you download and install to perform this task. Figure 20-7 shows the Exploit Protection settings in the Windows Security Control Panel app.

Technet24

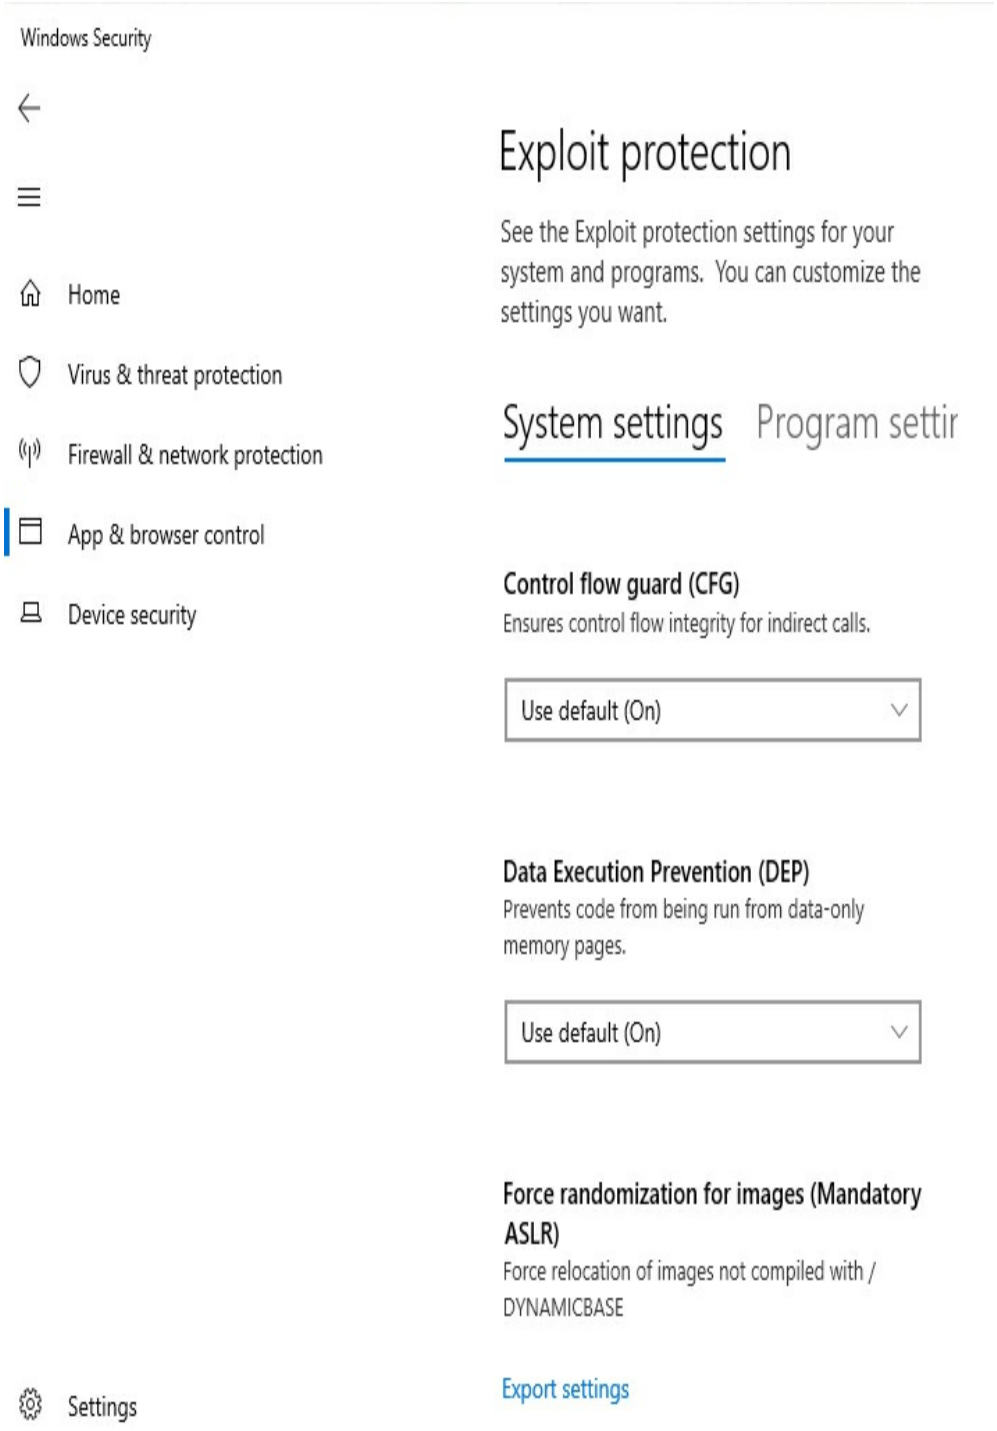


Figure 20-7 Exploit Protection

Some Exploit Protection mitigation settings are configurable at the system level, and some are only configurable on a per-app basis. The settings that can be configured at both the system and app level are as follows:

- **Control Flow Guard.** Configurable at the system and program level. Ensures control flow integrity for indirect calls.
- **Data Execution Prevention.** Configurable at the system and program level. Prevents code from being executed from data-only memory pages.
- **Force Randomization For Images (Mandatory ASLR).** Configurable at the system and the program level. Forces reallocation of memory images that haven't been compiled with the DYNAMICBASE option.

The settings that you can only configure at the application level are as follows:

- **Arbitrary Code Guard.** Prevent non-image backed executable code and code page modification
- **Block Low Integrity Images.** Block the loading of images marked with low integrity
- **Block Remote Images.** Block the loading of images from remote devices
- **Block Untrusted Fonts.** Block the loading of any GDI-based fonts not installed in the system Fonts directory
- **Code Integrity Guard.** Restrict the loading of images to those that have been digitally signed by Microsoft
- **Disable Extension Points.** Disable extensibility mechanisms that allow DLL injection into all processes
- **Disable Win32k System Calls.** Block programs from using the

Win32K system call table

- **Do Not Allow Child Processes.** Prevent programs from spawning child processes
- **Export Address Filtering (EAF).** Filter dangerous exported functions that are being resolved by malicious code
- **Import Address Filtering (IAF).** Filter imported functions being resolved by malicious code
- **Randomize Memory Allocations (Bottom-Up ASLR).** Randomize locations for virtual memory allocation
- **Simulate Execution (SimExec).** Ensure that calls to sensitive functions are returned to legitimate callers
- **Validate API Invocation (CallerCheck).** Ensure that sensitive APIs can only be invoked by legitimate callers
- **Validate Exception Chains (SEHOP).** Ensures the integrity of an exception chain during dispatch
- **Validate Handle Usage.** Raises an exception on any invalid handle references
- **Validate Heap Integrity.** Terminates a program when heap corruption is detected
- **Validate Image Dependency Integrity.** Enforces code signing for Windows image dependency loading
- **Validate Stack Integrity (StackPivot).** Ensures that the stack has not been redirected for sensitive functions

You configure exploit protection from PowerShell using the *Set-ProcessMitigation* cmdlet. For example, to enable Data Execution Prevention (DEP) at the system level, you would run the command:

Click here to view code image

```
Set-ProcessMitigation -System -Enable DEP
```

You can view which exploit protection settings are enabled by running the following command:

Click here to view code image

```
Get-ProcessMitigation -System
```

You can use the following command to export an Exploit Guard configuration:

Click here to view code image

```
Get-ProcessMitigation -RegistryConfigFilePath exportedconfig.xml
```

You can import an exported Exploit Guard configuration by using the following command:

Click here to view code image

```
Set-ProcessMitigation -PolicyFilePath exportedconfig.xml
```

More Info

Exploit Protection

You can learn more about deploying Exploit Protection at https://docs.microsoft.com/en-us/windows/security/threat-protection/microsoft-defender-atp/enable-exploit-protection.

WINDOWS DEFENDER

Windows Server 2019 includes Windows Defender, the Microsoft antimalware solution. Windows Defender is enabled by default when you deploy Windows Server

2019. To disable Windows Defender, you must remove it, either by using the Add Roles And Features Wizard or by using the following *Uninstall-WindowsFeature* command:

Click here to view code image

```
Uninstall-WindowsFeature -Name Windows-Server-Antimalware
```

You can use the following PowerShell cmdlets to manage Windows Defender on computers running the GUI or the Server Core version of Windows Server 2019:

- *Add-MpPreference*. Modifies Windows Defender settings
- *Get-MpComputerStatus*. View the status of antimalware software on the server
- *Get-MpPreference*. View the Windows Defender preferences for scans and updates
- *Get-MpThreat*. View the history of malware detected on the computer
- *Get-MpThreatCatalog*. Get known threats from the definitions catalog. You can use this list to determine whether a specific threat is known to Windows Defender
- *Get-MpThreatDetection*. Lists active and past threats that Windows Defender has detected on the computer
- *Remove-MpPreference*. Removes an exclusion or default action
- *Remove-MPThreat*. Removes an active threat from the computer
- *Set-MpPreference*. Allows you to configure scan and update preferences
- *Start-MpScan*. Starts an antimalware scan on a server

- *Start-MpWDOScan*. Triggers a Windows Defender offline scan
- *Update-MpSignature*. Triggers an update for antimalware definitions

More Info

Windows Defender

You can learn more about Windows Defender at https://docs.microsoft.com/en-us/windows-server/security/windows-defender/windows-defender-overview-windows-server.

WINDOWS DEFENDER SMARTSCREEN

Windows Defender SmartScreen, shown in Figure 20-8, provides you with a way of checking unrecognized files that have been downloaded to Windows Server against the properties of known malicious files that are stored in the Microsoft Graph security database. A good general security rule is to avoid downloading application files directly from the Internet and installing them on Windows Server; instead, you should download them to a separate location, test them, and then copy them remotely to Windows Server for installation.

Technet24

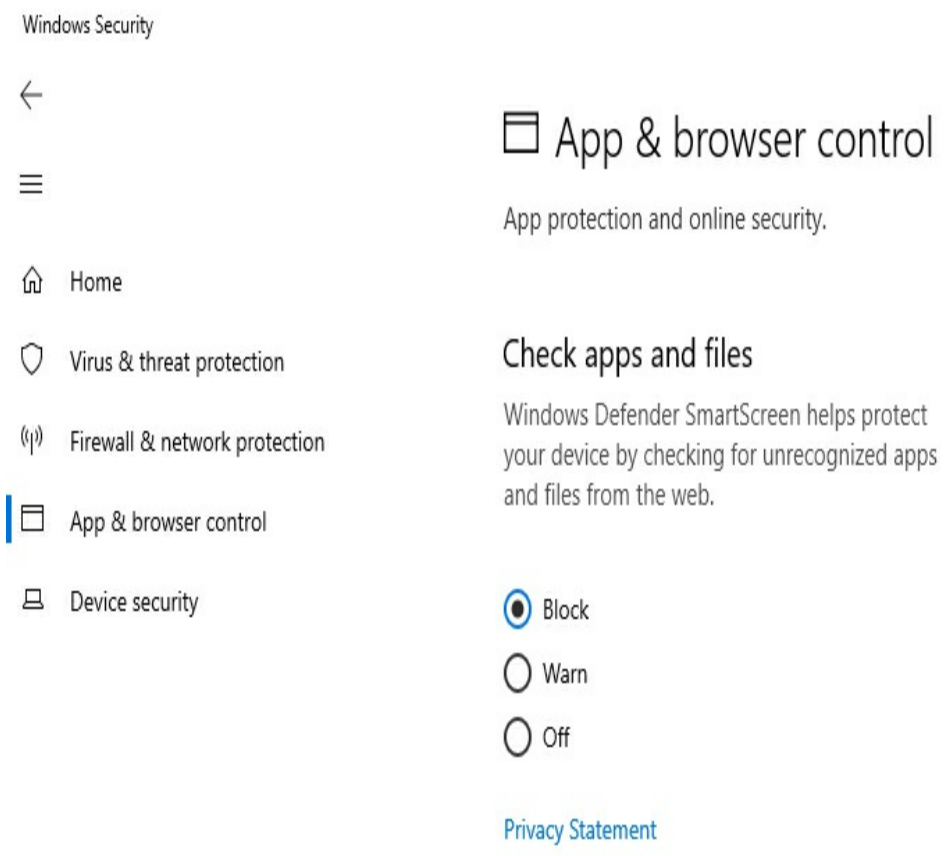


Figure 20-8 Windows Defender SmartScreen

More Info

Windows Defender Smart Screen

You can learn more about Windows Defender Smart Screen at https://docs.microsoft.com/en-us/windows/security/threat-protection/windows-defender-smartscreen/windows-defender-smartscreen-overview.

Chapter 21

Maintenance and monitoring

Data collector sets

Alerts

Event Viewer

Network monitoring

Azure Monitor

Windows Server Backup

Azure Backup

Vssadmin

Windows Server Update Services

Azure Update Management

Monitoring and maintenance related PowerShell cmdlets

WSUS related PowerShell cmdlets

Technet24

Just like any other form of precision equipment, servers need to be monitored and maintained over their operational lifespans to ensure that they continue to function as well as possible. Monitoring and maintenance not only involves regularly checking performance and event logs but also ensuring that servers and the data that those servers host are backed up in such a way that all important data and settings can be recovered in the event of a catastrophic failure. In this chapter, you learn how to monitor Windows Server performance, manage event logs, configure advanced auditing, as well as use the built-in backup and recovery tools.

DATA COLLECTOR SETS

Data collector sets enable you to gather performance data, system configuration information, and statistics into a single file. You can use Performance Monitor or other third-party tools to analyze this information to determine how well a server is functioning against an assigned workload.

You can configure data collector sets to include the following:

- **Performance counter data.** The data collector set not only includes specific performance counters but also the data generated by those counters.
- **Event trace data.** Enables you to track events and system

activities. Event trace data can be useful when you need to troubleshoot misbehaving applications or services.

- **System configuration information.** Enables you to track the state of registry keys and record any modifications made to those keys.

Windows Server 2219 includes the following built-in data collector sets:

- **Active Directory diagnostics.** Available if you have installed the computer as a domain controller and provides data on Active Directory health and reliability.
- **System diagnostics.** Enables you to troubleshoot problems with hardware, drivers, and STOP errors.
- **System performance.** Enables you to diagnose problems with sluggish system performance. You can determine which processes, services, or hardware components might be causing performance bottlenecks.

To create a data collector set, perform the following steps:

1. Open Performance Monitor from the Tools menu of the Server Manager console.
2. Expand Data Collector Sets.
3. Click User Defined. On the Action menu, click New and then click Data Collector Set.
4. You are given the option to create the data collector set from a template, which enables you to choose from an existing data collector set or to create a data collector set manually. If you choose to create a data collector set manually, you have the option to create a data log, which can include a performance counter, event trace data, and system configuration information, or to create a performance counter alert.
5. If you select Create Data Logs and select Performance Counter, next

choose which performance counters to add to the data collector set. You also specify how often Windows should collect data from the performance counters.

6. If you choose to include event trace data, you need to enable event trace providers. A large number of event trace providers are available with Windows Server 2219. You use event trace providers when troubleshooting a specific problem. For example, the Microsoft Windows-AppLocker event trace provider helps you diagnose and troubleshoot issues related to AppLocker.
7. If you choose to monitor system configuration information, you can select registry keys to monitor. Selecting a parent key enables you to monitor all registry changes that occur under that key while the data collector set is running.
8. Next, specify where you want data collected by the data collector set to be stored. The default location is the *%systemdrive%\PerfLogs\Admin* folder. If you intend to run the data collector set for an extended period of time, you should store the data on a volume that is separate from the one hosting the operating system because this ensures that you don't fill up your system drive by mistake.
9. The final step in setting up a data collector set is to specify the account under which the data collector set runs. The default is Local System, but you can configure the data collector set to use any account for which you have credentials.

You can schedule when a data collector set runs by configuring the Schedule tab of a data collector set's Properties dialog box.

ALERTS

Performance counter alerts enable you to configure a task to run when a performance counter, such as available disk space or memory, falls under or exceeds a

specific value. To configure a performance counter alert, you create a new data collector set, choose the Create Manually option, and select the Performance Counter Alert option.

You add the performance counter, threshold value, and whether the alert should be triggered if the value exceeds or falls below the set value. By default, when you create an alert, it only adds an event to the event log when it is triggered. This is useful if you have tools that allow you to extract this information from the event log and use the information as a way of tracking when alerts occur. You can also configure an alert to run a scheduled task when it is triggered. You can do this by editing the Properties dialog box for the alert and by specifying the name of the scheduled task on the Task tab.

EVENT VIEWER

Event Viewer enables you to access recorded event information. Windows Server 2219 not only offers the application, security, setup, and system logs, but it also contains separate application and service logs. These logs are designed to provide information on a per role or per application basis, rather than having all application and role service-related events funneled into the application log. When searching for events related to a specific role service, feature, or application, check to see whether that role service, feature, or application has its

own application log.

Event log filters

Filters and event logs enable you to view only those events that have specific characteristics. Filters apply only to the current Event Viewer session. If you constantly use a specific filter or set of filters to manage event logs, you should instead create a custom view. Filters apply only to a single event log. You can create log filters based on the following properties:

- **Logged.** Enables you to specify the time range for the filter.
- **Event Level.** Enables you to specify event levels. You can choose the following options: Critical, Warning, Verbose, Error, and Information.
- **Event Sources.** Enables you to choose the source of the event.
- **Event IDs.** Enables you to filter based on event ID. You can also exclude specific event IDs.
- **Keywords.** Enables you to specify keywords based on the contents of events.
- **User.** Enables you to limit events based on user.
- **Computer.** Enables you to limit events based on the computer.

To create a filter, perform the following steps:

1. Open Event Viewer and select the log that you want to filter.
2. Determine the properties of the event that you want to filter.
3. On the Actions pane, click Filter Current Log.
4. In the Filter Current Log dialog box, specify the filter properties.

Event log views

Event log views enable you to create customized views of events across any event log stored on a server, including events in the forwarded event log. Rather than looking through each event log for specific items of interest, you can create event log views that target only those specific items. For example, if there are certain events that you always want to look for, create a view and use the view rather than combing through the event logs another way. By default, Event Viewer includes a custom view named Administrative Events. This view displays critical, warning, and error events from a variety of important event logs such as the application, security, and system logs.

Views differ from filters in the following ways:

- **Persistent.** You can use a view across multiple Event Viewer sessions. If you configure a filter on a log, it is not available the next time you open the Event Viewer.
- **Include multiple logs.** A custom view can display events from separate logs. Filters are limited to displaying events from one log.
- **Exportable.** You can import and export event log views between computers.

The process for creating an event log view is similar to the process for creating a filter. The primary difference is that you can select events from multiple logs, and you give the event log view a name and choose a place to save it. To create an event log view, perform the

following steps:

1. Open Event Viewer.
2. Click the Custom Views node and then click Create Custom View from the Actions menu.
3. In the Create Custom View dialog box, select the properties of the view, including:
    1. When the events are logged
    2. The event level
    3. Which event log to draw events from
    4. Event source
    5. Task category
    6. Keywords
    7. User
    8. Computer
4. In the Save Filter To Custom View dialog box, enter a name for the custom view and a location to which to save the view. Click OK.
5. Verify that the new view is listed as its own separate node in the Event Viewer.

You can export a custom event log view by selecting the event log view and clicking Export Custom View. Exported views can be imported on other computers running Windows Server 2219.

Event subscriptions

Event log forwarding enables you to centralize the collection and management of events from multiple

computers. Rather than having to examine each computer's event log by remotely connecting to that computer, event log forwarding enables you to do one of the following:

- Configure a central computer to collect specific events from source computers. Use this option in environments where you need to consolidate events from only a small number of computers.
- Configure source computers to forward specific events to a collector computer. Use this option when you have a large number of computers that you want to consolidate events from. You configure this method by using Group Policy.

Event log forwarding enables you to configure the specific events that are forwarded to the central computer. This enables the computer to forward important events. It isn't necessary, however, to forward all events from the source computer. If you discover something from the forwarded traffic that warrants further investigation, you can log on to the original source computer and view all the events from that computer in a normal manner.

Event log forwarding uses Windows Remote Management (WinRM) and the Windows Event Collector (wecsvc). You need to enable these services on computers that function as event forwarders and event collectors. You configure WinRM by using the *winrm quickconfig* command, and you configure wecsvc by using the *wecutil qc* command. If you want to configure

subscriptions from the security event log, you need to add the computer account of the collector computer to the local Administrators group on the source computer.

To configure a collector-initiated event subscription, configure WinRM and Windows Event Collector on the source and collector computers. In the Event Viewer, configure the Subscription Properties dialog box with the following information:

- **Subscription Name.** The name of the subscription.
- **Destination Log.** The log where collected events are stored.
- **Subscription Type And Source Computers: Collector Initiated.** Use the Select Computers dialog box to add the computers that the collector retrieves events from. The collector must be a member of the local Administrators group or the Event Log Readers group on each source computer, depending on whether access to the security log is required.
- **Events To Collect.** Create a custom view to specify which events are retrieved from each of the source computers.

If you want to instead configure a source computer-initiated subscription, you need to configure the following Group Policies on the computers that act as the event forwarders:

- **Configure Forwarder Resource Usage.** This policy determines the maximum event forwarding rate in events per second. If this policy is not configured, events are transmitted as soon as they are recorded.
- **Configure Target Subscription Manager.** This policy enables you to set the location of the collector computer.

Both of these policies are located in the *Computer Configuration\Policies\Administrative Templates\Windows Components\Event Forwarding* node. When configuring the subscription, you must also specify the computer groups that hold the computer accounts of the computers that are forwarding events to the collector.

Event-driven tasks

Event Viewer enables you to attach tasks to specific events. A drawback to the process of creating event-driven tasks is that you need to have an example of the event that triggers the task already present in the event log. Events are triggered based on an event having the same log, source, and event ID.

To attach a task to a specific event, perform the following steps:

1. Open Event Viewer. Locate and select the event on which you want to base the new task.
2. On the Event Viewer Actions pane, click Attach Task To This Event. The Create Basic Task Wizard opens.
3. On the Create A Basic Task page, review the name of the task that you want to create. By default, the task is named after the event. Click Next.
4. On the When An Event Is Logged page, review the information about the event. This lists the log that the event originates from, the source of the event, and the event ID. Click Next.
5. On the Action page, you can choose the task to perform. The Send An E-Mail and Display A Message tasks are deprecated, and you receive an

error if you try to create a task using these actions. Click Next.

6. On the Start A Program page, specify the program or script that should be automatically triggered, as well as additional arguments.
7. After you create the task, you can modify the task to specify the security context under which the task executes. By default, event tasks only run when the user is signed in, but you can also choose to configure the task to run regardless of whether the user is signed in.

NETWORK MONITORING

Network monitoring enables you to track how a computer interacts with the network. Through network monitoring, you can determine which services and applications are using specific network interfaces, which services are listening on specific ports, and the volume of traffic that exists. There are two primary tools that you can use to perform network monitoring on computers running Windows Server 2219:

- Resource Monitor
- Message Analyzer

Resource Monitor

Resource Monitor enables you to monitor how a computer running the Windows Server 2219 operating system uses CPU, memory, disk, and network resources. Resource Monitor provides real-time information. You can't use Resource Monitor to perform a traffic capture and review activity that occurred in the past, but you can use Resource Monitor to view activity that is currently

occurring.

Resource Monitor provides the following information that is relevant to network monitoring:

- **Processes With Network Activity.** This view lists processes by name and ID and provides information on bits sent per second, bits received per second, and total bits per second.
- **Network Activity.** Lists network activity on a per-process basis, but also lists the destination address, sent bits per second, received bits per second, and total bits per second.
- **TCP Connections.** Provides information on connections based on local address, port, and remote address and port.
- **Listening Ports.** Lists the ports and addresses on which services and applications are listening. This option also provides information about the firewall status for these roles and services.

Message Analyzer

Microsoft Message Analyzer is the successor to Network Monitor. You can use Message Analyzer to perform network traffic capture and analysis. Message Analyzer also functions as a replacement for LogParser, which enables you to manage system messages, events, and log files. When performing a capture, select the scenario that best represents the type of event that you are interested in capturing traffic for. For example, the LAN scenario enables you to capture traffic on local area network (LAN) interfaces.

When performing certain types of network traffic captures, you need to run Message Analyzer using an

account that is a member of the local Administrators group. After the capture is performed, you can analyze the content of each message. By applying appropriate filters, you can locate network traffic that has specific characteristics, such as traffic that uses a particular TCP port, source, or destination address. Figure 21-1 shows the Microsoft Message Analyzer.

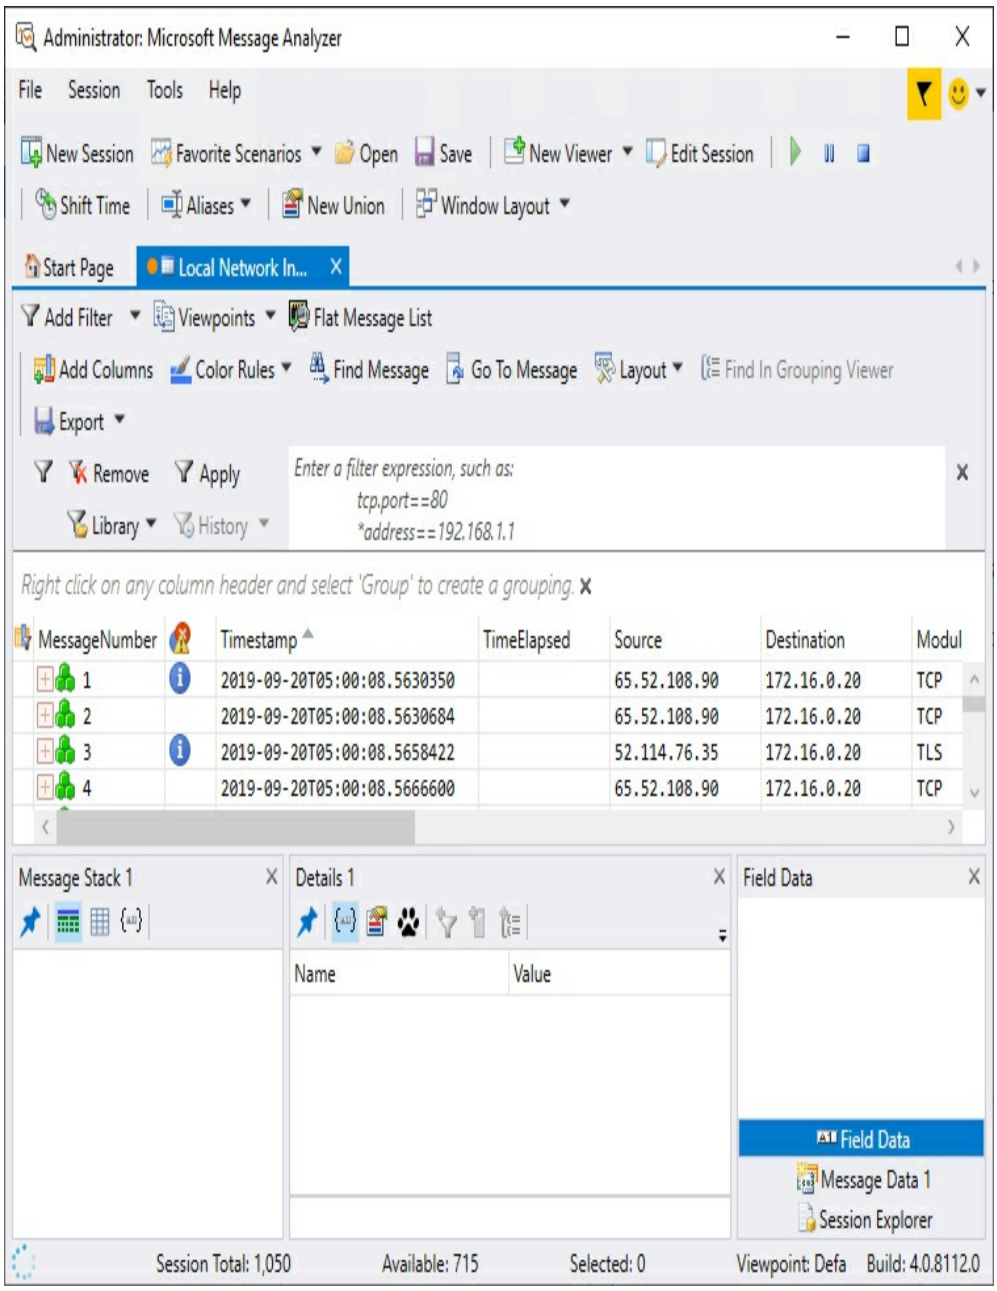


Figure 21-1 Microsoft Message Analyzer

AZURE MONITOR

Azure Monitor allows you to centrally collect

performance counter and event log data for storage and analysis in Azure. You can use the automation functionality to perform tasks on the local server when specific events or thresholds are met. While not nearly as fully featured as a System Center Operations Manager deployment, Azure Monitor allows you to leverage the analytics capabilities of Azure at a scale that's not possible with an on-premises monitoring solution.

To install Azure Monitor on a computer running Windows Server 2219, connect to the server using Windows Admin Center and perform the following steps:

1. In Windows Admin Center, click Azure Hybrid Services and then click Discover Azure Services. If you haven't connected the Windows Admin Center instance to an Azure subscription, you'll be prompted to do so.
2. On the list of Azure Resources, click Set Up in the section on Azure Monitor, as shown in Figure 21-2.

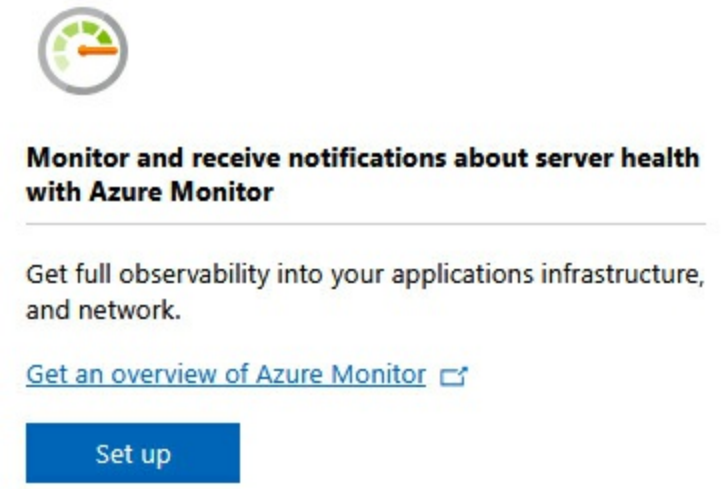


Figure 21-2 Setup Azure Monitor

3. If an appropriate resource group and log analytics workspace already exists within the subscription, the Azure Monitor setup process will detect these automatically. If these workspaces are not present, you will be prompted to create them.

4. Once the connection is configured, you can review analytics information in the appropriate log analytics workspace in the Azure portal, as shown in Figure 21-3

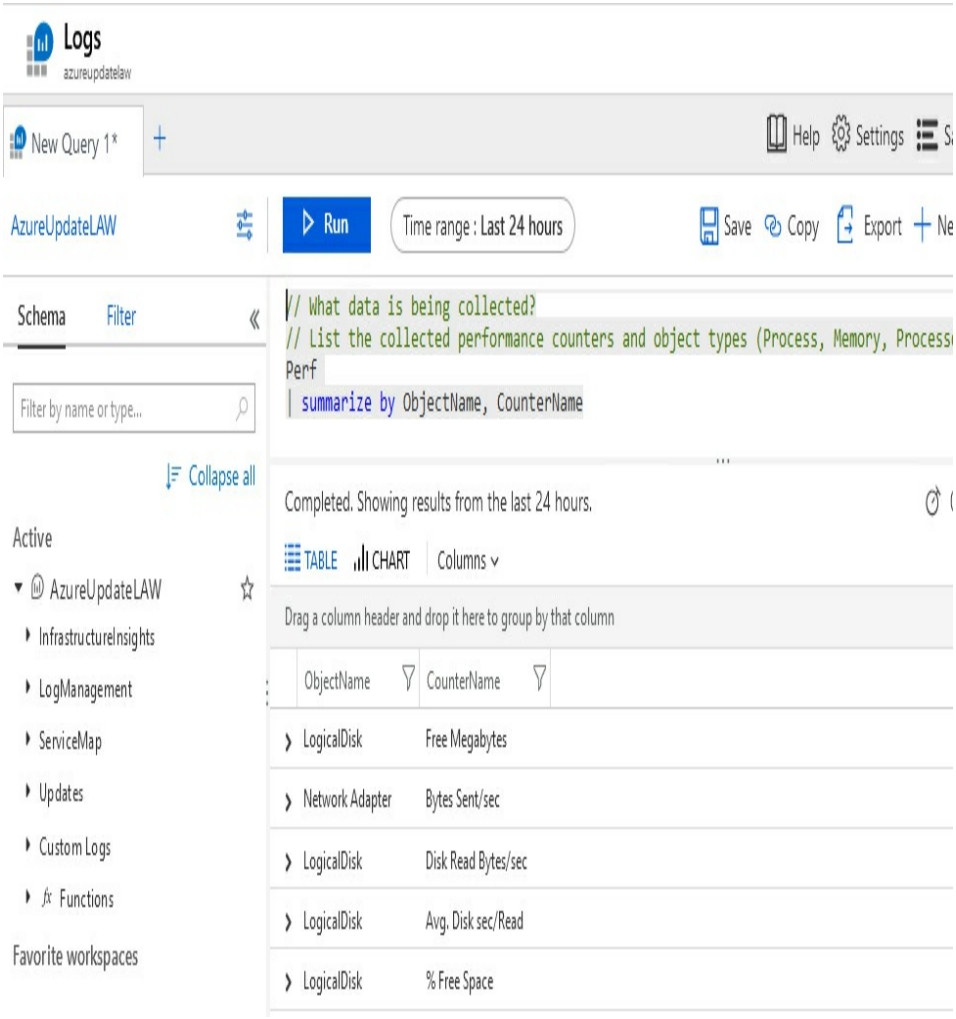


Figure 21-3 Log Analytics Workspace

Technet24

More Info

Azure Monitor

You can learn more about Azure Monitor at https://docs.microsoft.com/en-us/azure/azure-monitor/overview.

WINDOWS SERVER BACKUP

Windows Server Backup is the default backup application included with Windows Server 2219. Windows Server Backup is a basic backup solution and only enables you to back up to disk or a network share. Tasks for long-term retention, like export to tape, require a more sophisticated solution, such as System Center Data Protection Manager. Windows Server Backup has a minimal level of reporting functionality and has no native functionality that makes it possible to send alerts to an administrator through email if a backup fails.

Windows Server Backup enables you to back up and recover the following:

- Full server (all volumes)
- Specific volumes
- Specific folders
- System State data
- Individual Hyper-V hosted virtual machines
- Volumes exceeding 2 terabytes (TB) in size
- Volumes with 4 kilobyte (KB) sector size

- Cluster Shared Volumes

Although you can connect from the Windows Server Backup console to other computers to remotely manage Windows Server Backup, you can only use an instance of Windows Server Backup to back up the local computer. For example, whereas you can back up network shared folders that the computer is able to access, such as a mapped network drive, you can't configure Windows Server Backup on one computer to do a full-volume or System State backup of another computer.

You can configure exclusions for backup jobs run on Windows Server Backup. Exclusions are specific file types that you want to exempt from the backup process. You can configure exclusions based on file type, or you can choose to exclude the contents of entire folders and their subfolders. For example, you can configure an exclusion that stops files with the *.tmp* extension in the *C:\shared-docs* folder and its subfolders from being written to a backup.

Users who are local Administrators or members of the Backup Operators group can use Windows Server Backup to back up the entire computer, volumes, files or folders, and the System State. You can grant other security principals this right by editing the Back Up Files And Directories Group Policy item. By default, Windows

Server Backup does not encrypt backups, which means that the data stored in those backups might be read if the backup media falls into the wrong hands. When you are backing up data, you can configure access control so that only users with specific credentials can access the backup, but this still does not encrypt the backup.

Inside OUT

Backup security

The security of backups is important. Anyone who has access to backed-up data can restore it to a separate location where they have complete access to that data. The security of backed-up data is just as important as the security of the servers being backed up.

Backup locations

Windows Server Backup enables you to back up to any locally attached disk, to a volume, to the storage area network (SAN), or to any network folder. When configuring a scheduled backup, you should specify a destination volume or disk that is empty. If the disk is local or connected to the SAN, Windows Server Backup performs a full backup at least once every 14 days and incremental backups on subsequent backups. Incremental backups in Windows Server 2219 use block-level backups rather than file-level backups, meaning that the incremental backups are much smaller. Rather than backing up all the files that have changed since the last backup, only the data blocks that have changed on

the hard disk are backed up. For example, if you change one image in a 25-megabyte (MB) PowerPoint file after it is backed up, only the data blocks associated with that image are backed up next time, not the whole 25-MB file.

Inside OUT

Backup location

Although you can write backup data to a different volume on the same disk, you should back up data to a different disk, so if your disk fails, you don’t lose both the data being protected and the backup itself.

There is an exception to the rule about how data is written during scheduled automatic full and incremental backups. This exception occurs when the backup data is written to a network folder. When you back up to a network folder, as opposed to a SAN-connected disk, which appears as local to Windows Server, each backup is a full backup and the previous full backup that was stored on the network share is erased. Because only one backup at a time can be stored on a network share, if you choose to back up to this location you can't recover data from any backup other than the most recently performed one.

You can modify the performance of full system and full volume backups by using the Optimize Backup Performance dialog box, and you can increase backup performance by using the Faster Backup Performance

option. The drawback of selecting this option is that it reduces disk performance on the disks that host the volumes you are backing up.

Backing up data

You can back up data using a variety of methods with Windows Server Backup. You can configure a scheduled backup, which means a backup occurs on a scheduled basis. You can also perform a one-off backup. When you perform a one-off backup, you can either use the existing scheduled backup settings or you can configure a separate set of settings for one-off backups. For example, you can configure Windows Server Backup to perform a full server backup twice a day to a locally attached disk. You can connect at any time and perform a one-off backup where you select only specific files and folders and have them written to a location on the network.

Role- and application-specific backups

Most Windows Server 2219 roles and features store data in locations that are backed up when you perform a System State backup. System State data is automatically backed up when you perform a full server backup or select it for backup. Depending on the roles and features installed on the computer running Windows Server 2219, the System State can contain the following data:

- Registry
- Local users and groups
- COM+ Class Registration database
- Boot files
- Active Directory Certificate Services (AD CS) database
- Active Directory database (*Ntds.dit*)
- SYSVOL directory
- Cluster service information
- System files under Windows Resource Protection
- Internet Information Services (IIS) settings

Some applications also register themselves with Windows Server Backup. This means that when you perform a full server backup, you can recover data that is only relevant to the application without having to perform a full system restore. Support for application registration depends on the application. You can't select a specific application for backup using Windows Server Backup, but you can restore applications that have registered themselves with Windows Server Backup as long as you've previously performed a full server backup.

Restore from backups

Windows Server Backup enables you to restore data that has been backed up on the local computer, or data that was backed up on another computer using Windows

Server Backup that is accessible to the local computer. You can use Windows Server Backup to do the following:

- You can use Windows Server Backup to restore files and folders as well as applications.
- You can use Windows Server Backup to restore the System State data. After the System State data is restored, you need to restart the computer.
- You can use Windows Server Backup to restore any volume except the one that hosts the operating system. If you want to restore the volume that hosts the operating system, you need to boot into the Windows Recovery Environment.
- You can use the Windows Recovery Environment to perform a full server restore, also known as a bare metal recovery. When you do this, all existing data and volumes on the server are overwritten with the backed-up data.

If multiple backups of the data you want to restore exist, you need to select which version to restore. If you are unsure which date holds the data you want to restore, you should restore to multiple alternative locations and then perform the comparison. This process saves you from restoring, figuring out you've restored the wrong data, performing another restore, and then figuring out that restore isn't right either.

Windows Server Backup writes backups to *.vhdx* files, the same type of file that is used with Hyper-V and when creating disks for Internet Small Computer System Interface (iSCSI) targets and disks in storage

pools. Windows Server 2219 also allows you to mount the contents of a virtual hard disk (VHD), which allows you to examine those contents without having to perform a full restoration using Windows Server Backup.

Restore to an alternative location

When you are performing data restoration, you can choose to restore data to the original location or to an alternative location. It is not uncommon when restoring data to the original location for backup administrators to unintentionally overwrite good, live data with older, restored data. If you restore to an alternative location, it's possible to compare the restored data against the current data. It's also important when restoring data that you retain permissions associated with data.

If you choose to restore to the original location, you can configure Windows Server Backup to perform one of the following tasks:

- Automatically create copies if an original exists
- Overwrite existing versions
- Do not recover any item that exists in the recovery destination

AZURE BACKUP

The Microsoft Azure Recovery Services (MARS) Agent allows you to back up data directly to a Microsoft Azure

backup recovery services vault service. You install the MARS agent client on the Internet-connected computer server that you want to back up, and the backup data is stored in an Azure Recovery Vault. The MARS agent functions as an off-site data-storage and recovery service and can be used in conjunction with Windows Server Backup. If a site is lost, you can still recover this important data from Azure.

You can run both the Windows Server Backup and Azure Backup Agent services in parallel. This approach enables you to perform local backups and backups to Azure on the same server. You can then do frequent full server backups locally, with infrequent critical data backups to Azure. When you employ this strategy, you mostly restore data from your local backups. It's only when something goes drastically wrong that you need to restore data from Azure.

The key to understanding what data to back up to Azure is looking at what types of information your organization might lose if a site has some type of disaster. You can always reinstall operating systems and applications from media that you can easily obtain again. The data stored on your servers, however, such as documents and settings, is something that you can't generate from installation media. By storing data in the cloud, you can recover it in the event of a disaster after you've rebuilt your server infrastructure.

Preparing Azure Backup

The entire process of configuring and managing Azure Backup can be performed from Windows Admin Center. This includes the creation of a Recovery Services Vault that is used to store backup data in Azure.

To configure Azure Backup with Windows Admin Center, perform the following steps:

1. When you are connected to the server you wish to configure with Azure Backup, click Backup in Windows Admin Center, as shown in Figure 21-4.

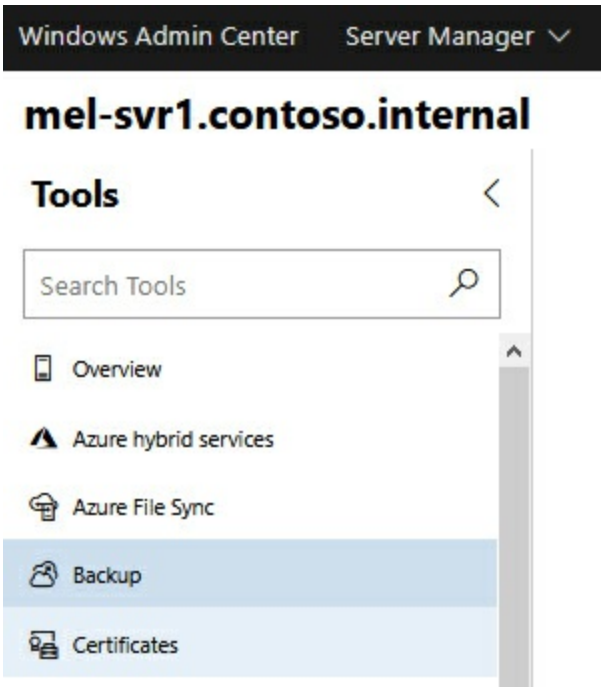


Figure 21-4 Backup item in Windows Admin Center

2. With Backup selected, click Set Up Azure Backup.
3. If you aren't already signed in to an account with Azure administrator permissions, you'll be required to sign in to Azure.

Technet24

4. By default, an Azure datacenter and recovery services vault will be selected for creation. You also have the option of creating your own recovery services vault, and you can specify a Resource Group and datacenter Location, as shown in Figure 21-5.

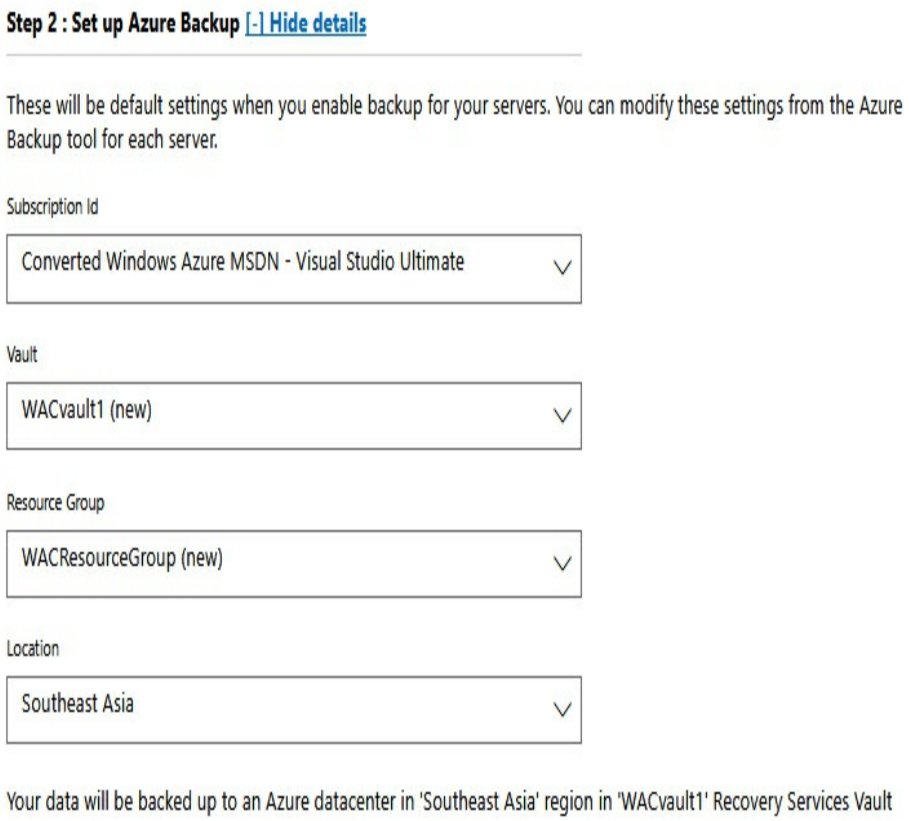


Figure 21-5 Configure Recovery Vault

5. By default, the System State data and all volumes will be protected. By default, backups will occur weekly and are retained for four weeks. You can alter what is protected using the Azure Backup Utility that will be installed on your server.

6. You will be prompted to provide an encryption passphrase. You will need to ensure that you keep a record of this passphrase somewhere other than a text file on the desktop of the server, as without the passphrase you will not be able to perform restore operations. Microsoft cannot perform a restoration for you if you lose the passphrase, as they are

never provided with a copy and just store the backup data encrypted with the passphrase. The passphrase must be a minimum of 16 characters long, and you can save it to an external location, such as a USB storage device.

7. Once you click OK, the infrastructure in Azure will be created for you, and the selected backup schedule will be enacted.

Backing up data to Azure Backup Agent

Modifying an Azure Backup Schedule and selections is very similar to the Schedule Backup Wizard in Windows Server Backup. After the Azure Backup Agent is installed by Windows Admin Center, you can run the tool to:

- **Select which items to back up.** Item selection is file and folder based. Although you can select a volume to back up, you don't use Azure Backup to perform full volume recovery in the same way you would use Windows Server Backup. When selecting items to back up, you can configure exclusions for file types, folders, or folder trees.
- **Select a backup schedule.** This option determines how often a synchronization occurs. You can configure Azure Backup to synchronize up to three times per day. You can also configure bandwidth throttling. Throttling enables you to limit the utilization of bandwidth and ensures that your organization's Internet connection isn't choked with backup traffic replicating to the recovery vault on Azure during business hours.
- **Configure backup retention.** The retention setting, which you configure on the Specify Retention Setting page, determines how long backup data is stored in Azure before it is deleted. You can configure retention for Windows Server Backup when creating a policy in Windows PowerShell.

Restore from Azure Backup

Azure Backup enables you to restore files and folders that are stored in your Azure recovery vault. You can't perform a full server recovery, System State recovery, or volume recovery using Azure Backup; you can only restore files and folders. Of course, if you've backed up all the files and folders on a volume, you can either restore all of them individually or all at one time. Azure Backup also enables you to restore data from another server that you've backed up to Azure. Recovering data using Azure Backup involves performing the following steps:

- Select the server from which you want to recover data.
- Select whether you want to browse or search for files to recover.
- Select the volume that hosts the data you want to recover and the specific backup date and time from which you want to recover data.
- Select the items that you want to recover. You can recover files, folders, and folder trees.
- Select Recovery Options, including whether to restore to the original location, to an alternative location, or to create copies or overwrite original files. You can also use this page to choose whether to restore the original permissions.

To remove all backup data from a recovery vault using Azure Backup, you'll need to go into the Azure Console and generate a security PIN, as shown in Figure 21-6. The PIN is valid only for a short period of time.

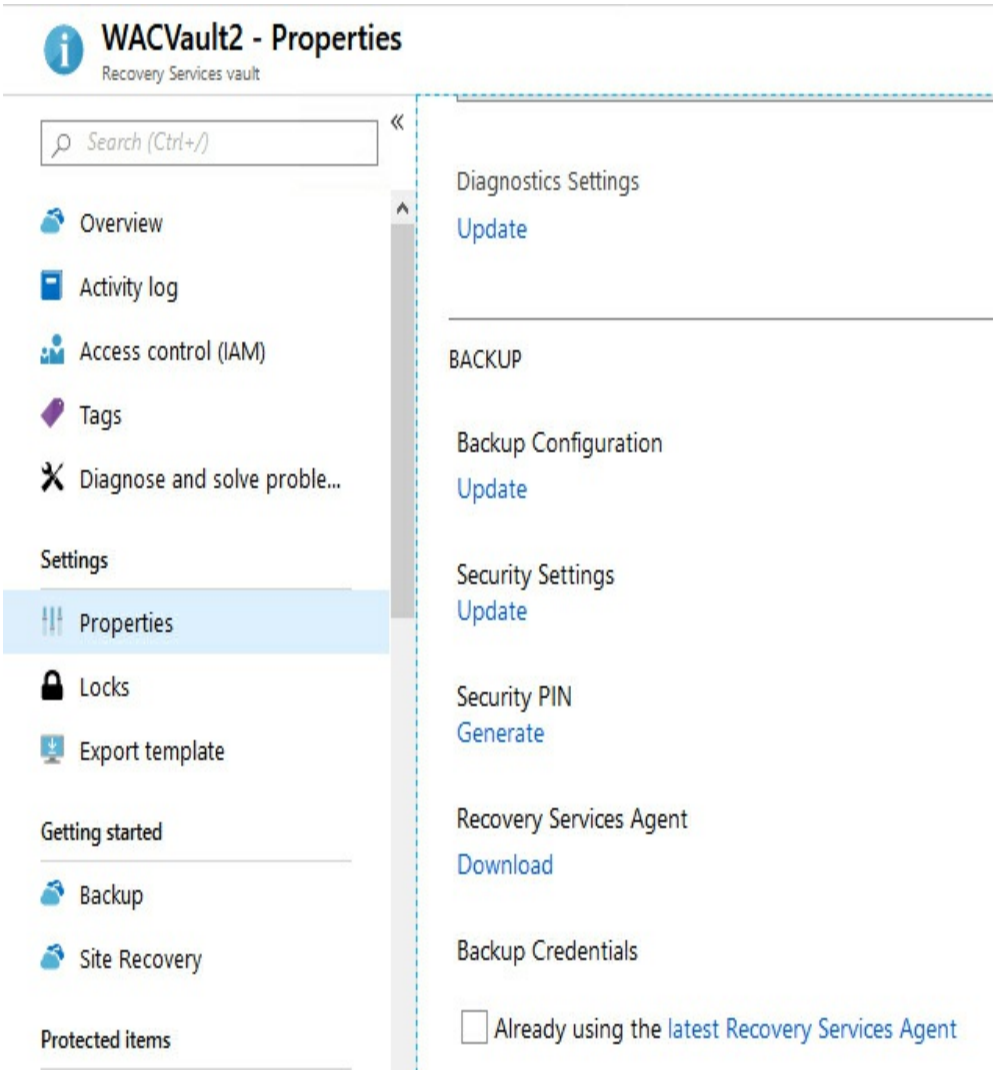


Figure 21-6 Configure Recovery Vault

More Info

Azure Backup

You can learn more about Azure Backup at https://docs.microsoft.com/en-us/azure/backup/backup-overview.

VSSADMIN

Volume Shadow Copy Services (VSS) is a technology that provides a snapshot of data on a volume as it existed at a specific point in time. VSS enables you to make a consistent backup of a file that is in use, such as a mailbox database or SQL Server database. Prior to the introduction of VSS, you might have needed to take such a database offline to ensure that the database's backup was consistent. Consistency issues arise when it takes so long to back up a large file or system that the configuration of the system or the contents of the file change during the backup. Windows Server Backup, Azure Backup Agent and other backup products, such as Data Protection Manager, use VSS to ensure that the data backed up is consistent and represents the state of the backed up data as it was at the point when the backup started without having to take in-use files offline.

Vssadmin is a command-line utility that enables you to manage volume shadow copy snapshots. You can use VSS admin to perform the following tasks:

- Configure the location of shadow copy storage.
- Create a shadow copy.
- Delete a shadow copy.
- Delete shadow copy storage.
- View existing shadow copies.
- View existing shadow copy storage.

- View volumes that are configured with shadow copies.
- View subscribed shadow copy writers and providers (special software that creates and manages shadow copies).
- Resize shadow copy storage.

You can also view shadow copy status on a per volume basis through the Previous Versions tab. When used with file shares, the VSS snapshots exposed through the Previous Versions functionality enable users to recover previous versions of files and folders without having to restore from backup. To do this, right-click the parent folder or volume and click Restore Previous Versions. You can then select previous versions of files that correspond to existing VSS snapshots.

Although Vssadmin allows you to create and manage VSS snapshots, you can't use Vssadmin to configure a schedule to automatically create VSS snapshots. You can configure a schedule to create VSS snapshots on a per-volume basis by right-clicking a volume and clicking Configure Shadow Copies. After you enable shadow copies, you can configure a schedule in the Settings dialog box. By default, when you enable Shadow Copies, a shadow copy is created at 07:00 and noon every weekday, but you can modify the schedule so that copies are created more often. When doing this, remember that after the initial allocated space used to store shadow copies is consumed, older shadow copies are removed to make space to store new versions. The amount of space

needed to store shadow copies and the retention period depends on the properties of the data stored on the volume.

WINDOWS SERVER UPDATE SERVICES

You can install Windows Server Update Services (WSUS) as a role on Windows Server 2219 in both the Server Core and GUI configurations. Although WSUS does include Windows PowerShell support, not all WSUS functionality has been replicated in Windows PowerShell, and you'll likely need to manage a WSUS server deployed on Server Core using the RSAT tools.

When you install WSUS, you can choose between using a local Windows Internal Database (WID) or a SQL Server instance. The advantage of using a SQL Server instance is that it's easier to back up and you can run more complex reports. The majority of WSUS deployments use the built-in WID database. When you install WSUS on Windows Server 2219, all prerequisite components are also installed.

Products, security classifications, and languages

During setup, you are asked to choose which update you want to download based on product name, security classification, and languages. Although you can choose

to download updates for all product categories for all classifications in all languages, you'll minimize the amount of configuration required later if you download updates only for products used on your organizational network.

When WSUS synchronizes, it may update the list of available product names to reflect newly released software. If your organization deploys a new product, if it retires an old product, or if you simply want to alter which updates are synchronized, you can modify which products are updated using the Products And Classifications dialog box, available through Options in the Update Services console, and shown in Figure 21-7.

Technet24

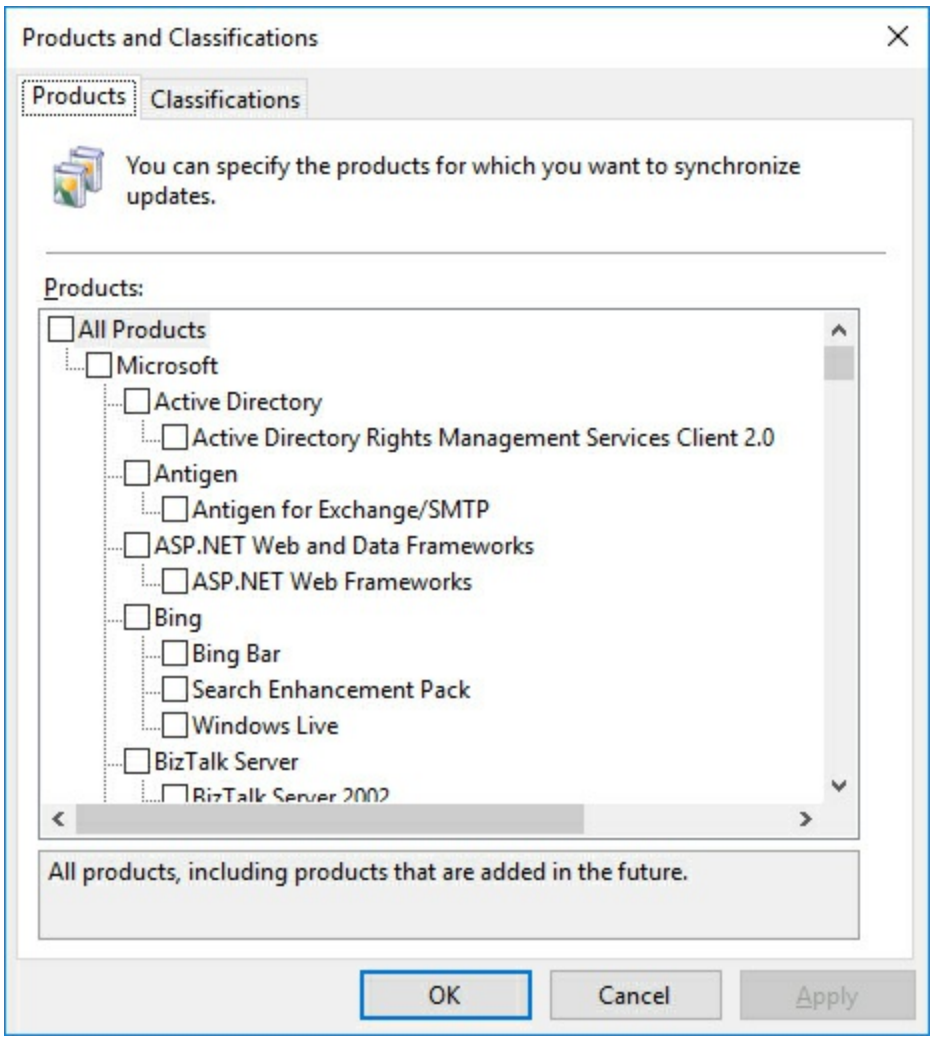


Figure 21-7 WSUS Products

Autonomous and replica modes

A single WSUS server can support approximately 25,000 clients. Large organizations often have multiple WSUS servers because it's better to have a local WSUS server at each large site, rather than having clients pull updates and approvals across wide area network (WAN) links. Instead of administrators performing the same

approvals on each WSUS server in the organization, you can configure a WSUS server as a replica of another server. When you configure a WSUS server as a replica, the downstream server copies all update approvals, settings, computers, and groups from its parent. You can configure the Update Source settings as well as specify information that enables WSUS to use a proxy server, through the Update Source And Proxy Server item in Options, in the Update Services console, and shown in Figure 21-8.

Technet24

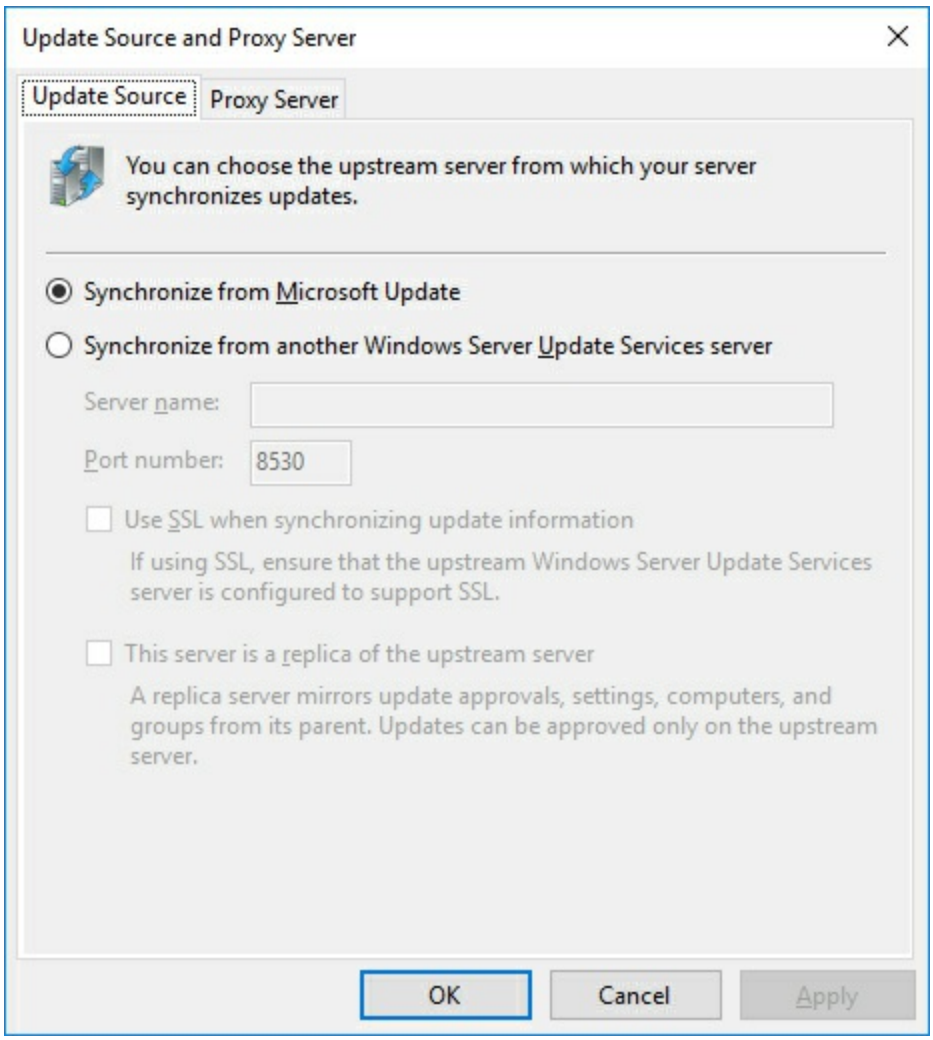


Figure 21-8 Update Source

Update files

One of the benefits of deploying WSUS is that clients on the local network download their updates from the WSUS server rather than downloading updates from the Microsoft Update servers on the Internet. You can configure update storage location settings using the Update Files And Languages item in the Options area of

the Update Services console. You can configure the following options, as shown in Figure 21-9:

- **Store Update Files Locally On This Server.** When you choose this option, you can choose whether to download files only after they have been approved; download express installation files, which install quicker on clients; or download files from Microsoft Update. With the last option, you can configure a server as a replica server but have update files downloaded from Microsoft Update rather than from the upstream replica server.
- **Don't Store Update Files Locally; Computers Install From Microsoft Update.** When you configure this option, clients use WSUS for update approvals but retrieve the updates from the Microsoft Update servers on the Internet. This option is most appropriate when you are providing update approvals to clients located outside of the organizational network.

Technet24

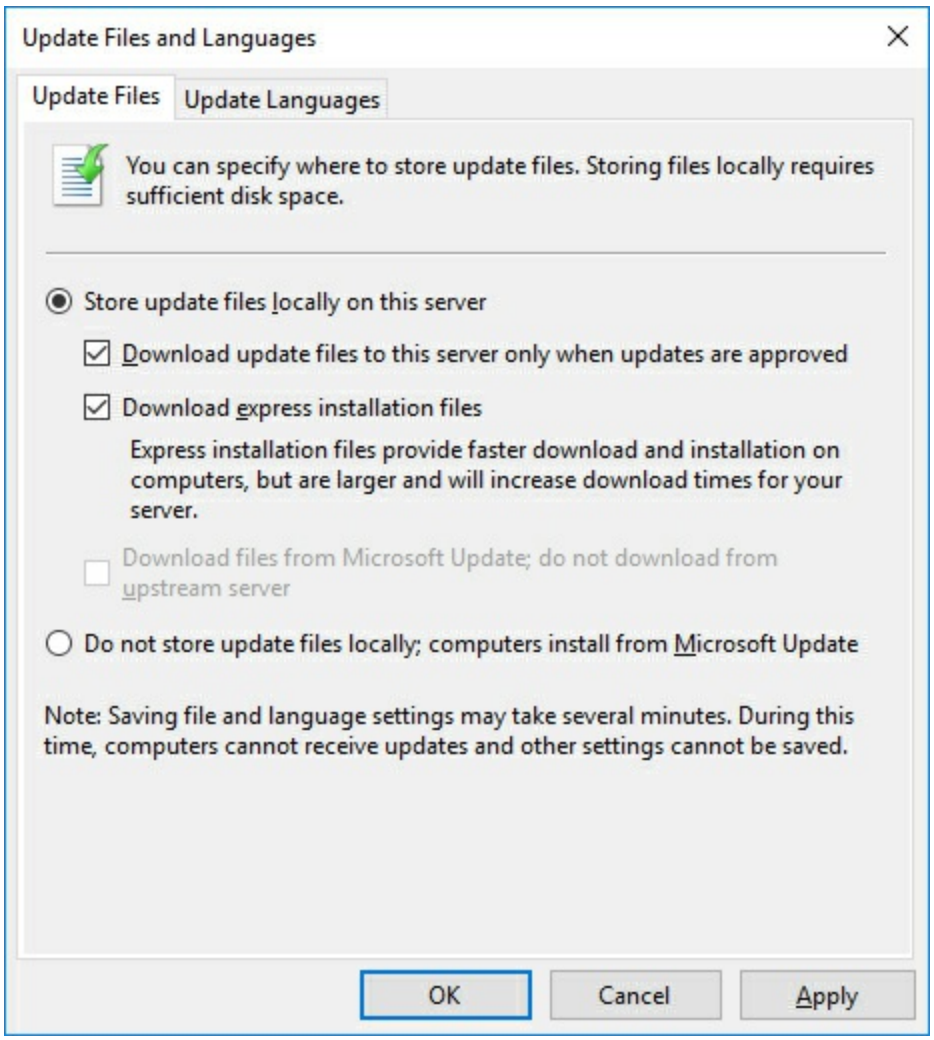


Figure 21-9 Update Files

WSUS security roles

In large organizations, you are more likely to separate the roles of server administrator and update administrator. When you install WSUS, two local security groups are created. By adding users to these groups, you grant users the permission to perform the tasks assigned with these roles. The roles are as follows:

- **WSUS Administrators.** Users who are added to the local WSUS Administrators group can perform any WSUS administration task. These tasks include approving updates, managing computer groups, configuring automatic approval rules, and modifying the WSUS server's update source.
- **WSUS Reporters.** Users who are members of this role can run reports on the WSUS server. These reports detail the update compliance status on the basis of update and computer. For example, a user who is a member of this group can run a WSUS report and determine which computers are missing a specific critical update.

WSUS groups

You can use WSUS groups to organize computers for the purpose of deploying updates. For example, you might have a WSUS group for servers in Sydney and another WSUS group for servers in Melbourne. A computer can be a member of multiple WSUS groups, and WSUS groups can exist in parent–child relationships. For example, the Australia WSUS group might have both the Melbourne and Sydney WSUS groups as members. Updates approved for the Australia group are automatically approved for members of the Melbourne and Sydney groups unless overridden.

You can assign computers to WSUS groups manually or through Group Policy. Computers can be assigned to WSUS groups through Group Policy only if the computer groups already exist on the WSUS server. To assign a computer manually, the computer must have

already reported to the WSUS server. Computers that have reported to the WSUS server but have not been assigned to a group are members of the Unassigned Computers group.

An administrator must create WSUS groups. To create a WSUS group, perform the following steps:

1. Open the Update Services console.
2. Click the group you want to have as the parent group. The Computers/All Computers group is the parent group for all groups.
3. From the Action menu, click Add Computer Group.
4. Specify the computer group name and click Add.

WSUS policies

You can configure most WSUS client options through Group Policy. Many of these policies are related to the experience that users of client operating systems have when updates are installed and are not directly applicable to updating server operating systems. Windows Update policies are located in the *Computer Configuration\Policies\Administrative Templates\Windows Components\Windows Update* node of a standard GPO, as shown in Figure 21-10.

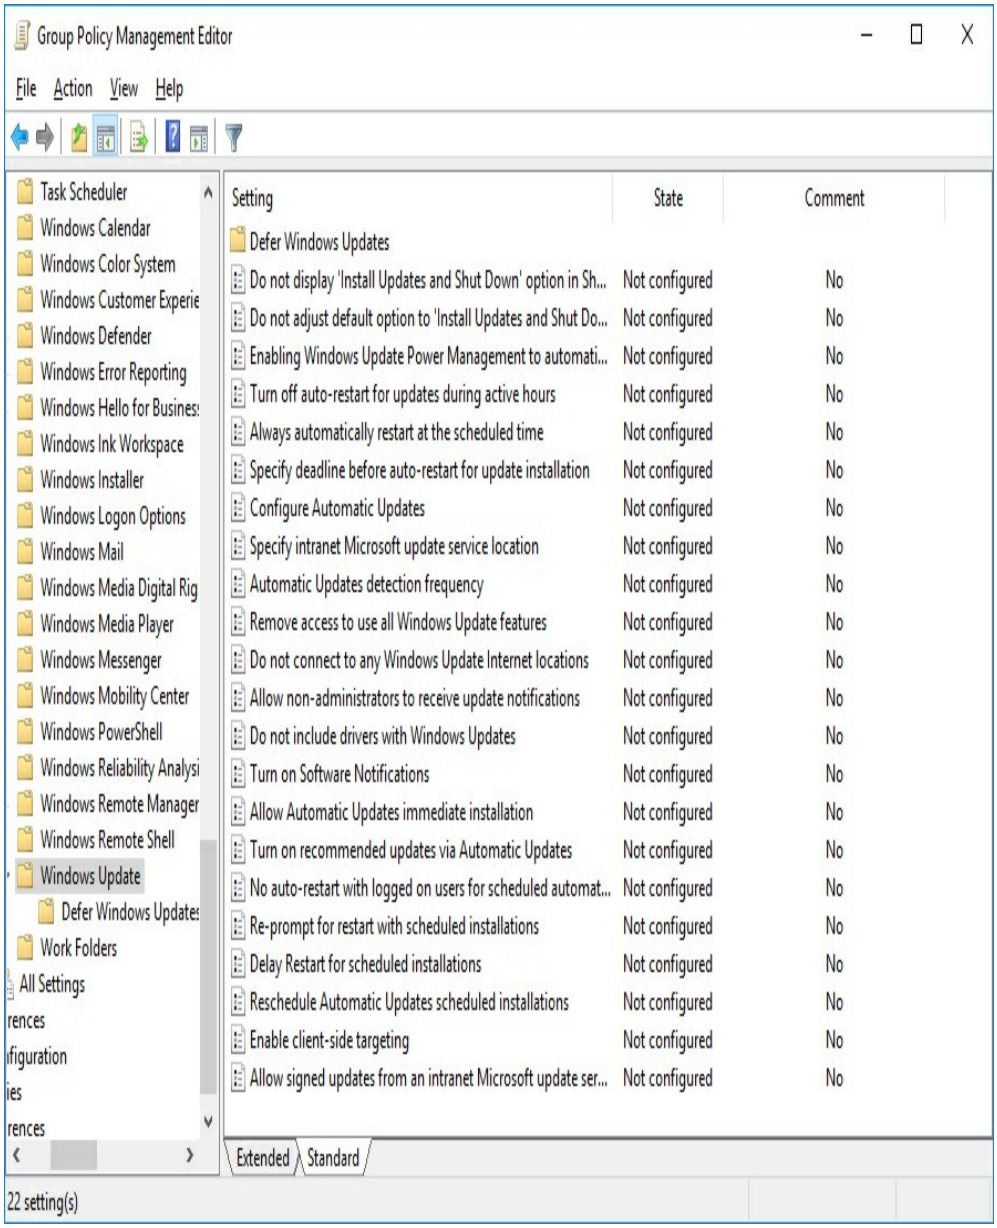


Figure 21-10 Update Policies

The most important policies from the perspective of the server administrator are as follows:

- **Configure Automatic Updates.** You can enable automatic updating, specify a day for update installations, and you can specify

a time for update installation to occur. It's usually not a good idea to have this one policy to apply to all servers in your organization. Having all servers install and reboot at the same time can cause substantial disruptions.

- **Specify Intranet Microsoft Update Service Location.** You can specify the location of the WSUS server and the statistics server. (The statistics server receives information on successful update installation and is usually the same as the WSUS server.)
- **Automatic Update Detection Frequency.** Determines how often the computer checks for updates.
- **Enable Client-Side Targeting.** Use this policy to specify which WSUS groups computers should be a member of. If the names do not match, those computers end up in the Unassigned Computers group.

Deploying updates

When you deploy updates, you decide whether to deploy the update, to which computer groups you deploy the update, and what deadline should apply to the deployment. You can deploy an update multiple times to different groups, so you can deploy an update to a test group and then, if no issues arise with the update, deploy the update more generally. Prior to deploying updates, you should perform a synchronization, which ensures that the WSUS server is up to date before choosing whether to deploy updates.

To deploy an update, perform the following steps:

1. Open the Update Services console and select the Updates\All Updates node. You can also select a child node, such as Critical Updates, if you want to view only available critical updates.

2. Set the Approval setting to Unapproved, set the Status to Any, and click Refresh. All unapproved updates are then listed.
3. Click one or more updates and then click Approve in the Actions pane.
4. In the Approve Updates dialog box, select which computer groups the update is approved for. You can choose between the following settings:
    1. **Approved For Install.** Approves the update.
    2. **Approved For Removal.** Removes a previously deployed update.
    3. **Not Approved.** Does not approve the update.
    4. **Keep Existing Approvals.** Inherits the approval from the parent group.
    5. **Deadline.** Specifies an update deployment deadline.

Automatic approval rules

Automatic approval rules enable specifically categorized updates to be automatically approved. For example, you might choose to automatically approve critical updates for the Sydney-Development-Servers WSUS group, as shown in Figure 21-11.

Technet24

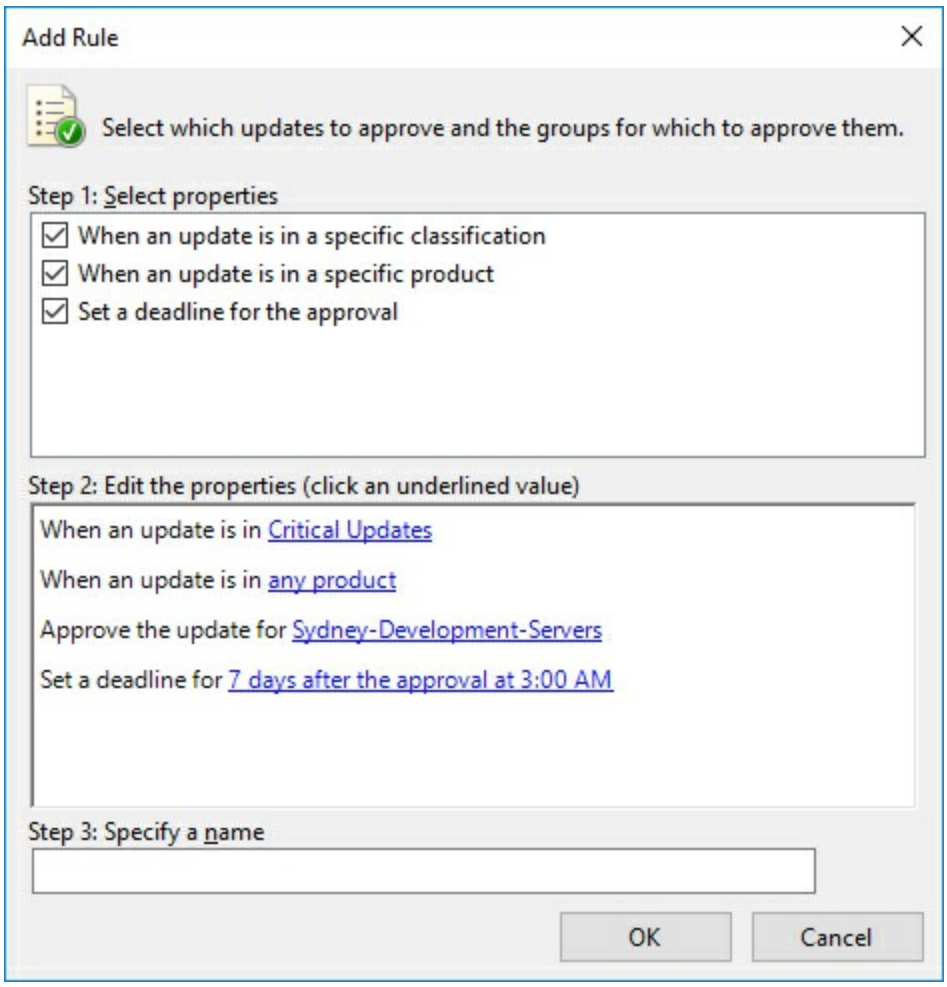


Figure 21-11 Automatic approval rules

To configure an automatic approval rule, perform the following steps:

1. Open the Update Services console. You can do this from the Tools menu of Server Manager or by right-clicking the server in a Server group and clicking Windows Server Update Services.
2. In the Update Services console, click Options, and then click Automatic Approvals.
3. In the Automatic Approvals dialog box, click New Rule.

4. In the Add Rule dialog box, choose the following rule options:

    1. **When An Update Is In A Specific Classification.** You can choose that the rule applies when an update matches a specific classification or number of classifications. Update classifications include Critical Updates, Definition Updates, Drivers, Feature Packs, Security Updates, Service Packs, Tools, Update Rollups, and Updates. Microsoft includes classifications for each software update when it publishes the update.
    2. **When An Update Is For A Specific Product.** You can specify products, either by category, such as Exchange, or by specific product, such as Exchange Server 2213.
    3. **Approve The Update For A Specific Computer Group.** The update can be approved for selected computer groups.
    4. **Set An Approval Deadline.** Sets an installation deadline for the update based on the time and date the update was first approved.

Inside OUT

Canary users

Automatic approval rules aren't suitable for production servers hosting important workloads because it is possible that an update will be installed without being properly tested. Automatic approval rules are suitable for test groups. You should populate your test group with users who are more likely to offer feedback if something goes wrong. Just as a canary in a coal mine was used by miners to detect dangerous gas, "canary users" are likely to raise an alarm when a software update causes problems that indicate it shouldn't be deployed in a production environment. Users who complain are much more valuable as deployment targets for update testing than users who ignore problems and do not provide feedback.

More Info

WSUS

You can learn more about WSUS at https://docs.microsoft.com/en-us/windows-server/administration/windows-server-update-services/deploy/deploy-windows-server-update-services.

AZURE UPDATE MANAGEMENT

Azure Update Management allows you to automate the deployment of updates to computers running both the Windows and Linux operating systems. You can configure on-premises Windows Servers to use Azure Update Management using Windows Admin Center. You can also manage the deployment of updates to those servers using Windows Admin Center. Azure Update Management differs from WSUS in that it requires a connection to Azure to work and cannot be deployed in offline scenarios. It also allows you to force the installation of updates on servers remotely, which is something that you can't directly do using WSUS.

To configure and use Azure Update Management on a Windows Server managed by Windows Admin Center, perform the following steps:

1. In Windows Admin Center, select Azure Hybrid Services node and click Discover Services. In the section describing Azure Update Management, shown in Figure 21-12, click Set Up.

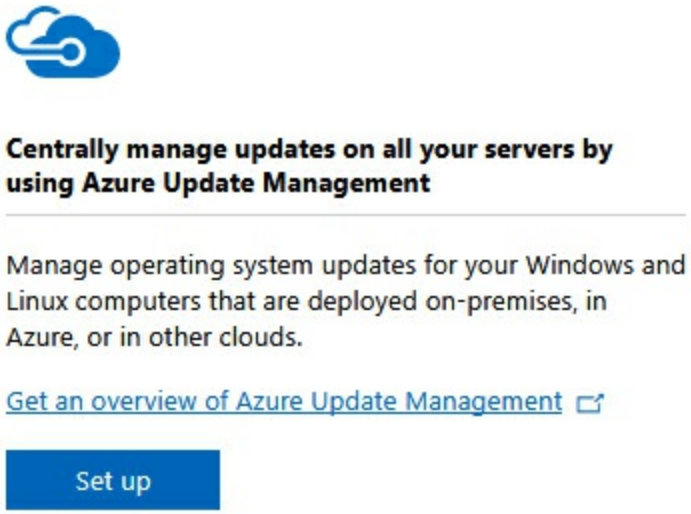


Figure 21-12 Configure Azure Update Management

2. After some time, you'll be asked whether you want to create a new or use an existing resource group to host the Azure elements of Azure Update Management. You'll also be asked whether you want to create a new or use an existing log analytics workspace and whether you want to create a new or use an existing Azure Automation account. Once you've made these decisions, as shown in Figure 21-13, click Set Up.

Technet24

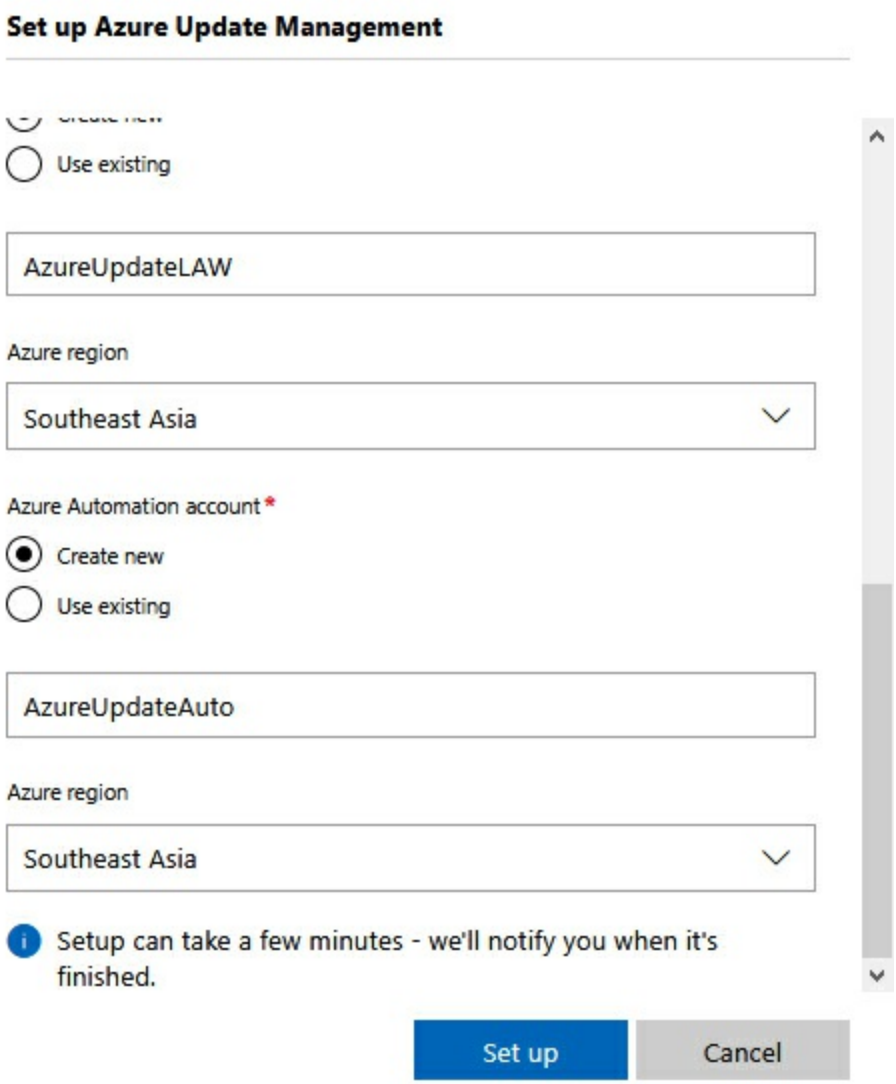


Figure 21-13 Configure Azure Update Management settings in Azure

3. Once Azure Update has been configured, click Azure Hybrid Services, and then click Azure Update Management, as shown in Figure 21-14.

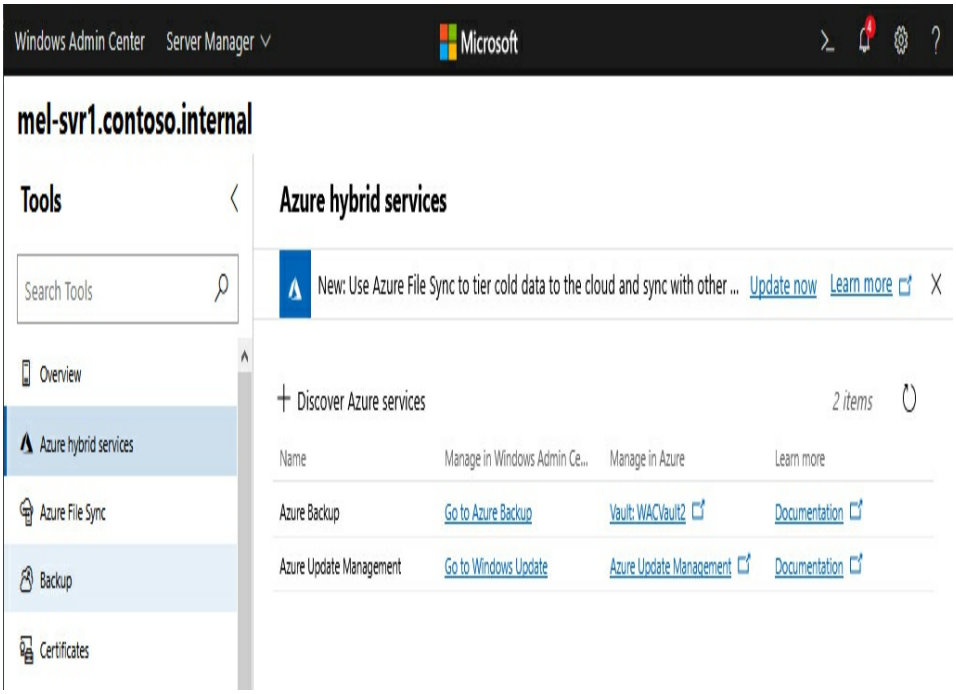


Figure 21-14 Azure Hybrid Services.

4. This will open the Azure Portal, which will provide a report on all servers that are configured to report into the Update Management instance. The Update management instance will provide information on each server associated with it, including which updates are missing and the nature of those updates, as shown in Figure 21-15.

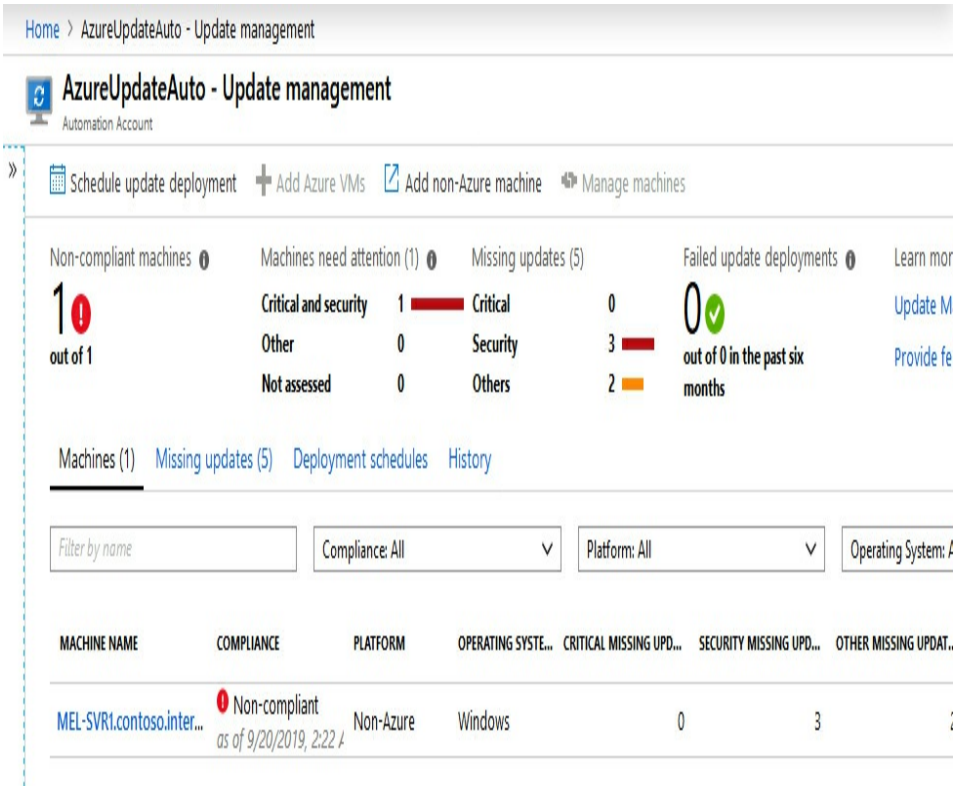


Figure 21-15 Azure Update Management.

5. To deploy updates, click Schedule Update Deployment in the Update Management blade of the Azure console. When configuring a scheduled update, you need to provide the following information:

    1. **Name Of The Update Deployment.** This is especially useful if you configure a recurring schedule.
    2. **Operating System.** Azure Update Management allows you to deploy updates to either computers running Windows or Linux operating systems in a single update deployment, but not to both operating systems in the one deployment.
    3. **Groups.** You can configure query-based groups so that the update deployment targets computers that meet the criteria specified in the query.
    4. **Machines To Update.** Allows you to select specific computers to which the update deployment applies.

5. **Update Classifications.** Rather than select specific updates, you configure an update deployment so that all updates that meet a specific update classification will be deployed (though you do have the option of excluding specific updates by Update ID). For Windows computers, the update classifications are Critical Updates, Security Updates, Update Rollups, Feature Packs, Service Packs, Definition Updates, Tools, and Updates.
6. **Include/Exclude Updates.** Allows you to choose specific updates based on KBIDs. You can find relevant KBIDS in the list of missing updates in the Azure Update Management console.
7. **Schedule Settings.** Allows you to specify when updates will be deployed. You can configure a schedule to recur. When you do this, all updates that meet the classification characteristics specified in the update deployment will be deployed.
8. **Maintenance Window.** The amount of time that can be taken to install updates, with the final 20 minutes of the assigned maintenance window reserved for restart operations. For example, if you set a maintenance window of 120 minutes, update installation will be halted after 100 minutes have elapsed so that restart operations can occur.
9. **Reboot Options.** Allows you to specify whether the server can restart automatically after update deployment or whether this step must be performed manually.

Azure Update Management provides you with more telemetry about which servers need updates than WSUS does. This isn't entirely surprising given that you'd be paying for Azure Update Management, and WSUS can be obtained free of charge.

More Info

Azure Update Management

You can learn more about Azure Update Management at https://docs.microsoft.com/en-us/azure/automation/manage-update-multi.

MONITORING AND MAINTENANCE RELATED POWERSHELL CMDLETS

There are a variety of PowerShell cmdlets that you can use for monitoring and maintenance tasks. These cmdlets are described in Table 21-1.

Table 21-1 Cmdlets used for Event Logs and Windows Server Backup

Noun	Verbs	Function
EventLog	Clear, Get, Limit, New, Remove, Show, Write	Manage Windows Event logs, including viewing events, creating events, and clearing logs
WinEvent	Get, New	View and create events from traditional event logs, as well as files generated by Event Tracing for Windows
WBApplicationRecovery	Start	Start application recovery
WBBackup	Start, Resume	Start or resume an existing backup

WBBackupSet	Get, Remove	Manage backups
WBBackupTarget	Get, Remove, Add, New	Manage backup targets
WBBackupVolumeBrowsePath	Get	View backup volumes
WBBareMetalRecovery	Remove, Get, Add	Configure bare metal recovery
WBCatalog	Restore, Remove	Manage Windows Server Backup catalog
WBDisk	Get	View disks online on the local computer
WBFileSpec	Get, New, Remove, Add	Manage items to include or exclude from a backup
WBHyperVRecovery	Start	Start Hyper-V VM recovery
WBJob	Stop, Get	Manage Windows Backup jobs
WBPerformanceConfiguration	Set, Get	Configure performance options

WBPolicy	Get, Remove, Set, New	Configure Windows Backup policies
WBSchedule	Set, Get	Configure Windows Backup schedule
WBSummary	Get	View backup summary information
WBSystemState	Add, Get, Remove	Manage backup of system state data
WBVirtualMachine	Add, Remove, Get	Manage the backup of virtual machines
WBVolume	Add, Remove, Get	Configure Windows Backup volumes
WBVolumeRecovery	Resume, Start	Manage volume recovery
WBVssBackupOption	Get, Set	Configure VSS backup options

WSUS RELATED POWERSHELL CMDLETS

There are a variety of PowerShell cmdlets that you can use for WSUS tasks. These cmdlets are described in

Table 21-2.

Table 21-2 Cmdlets used for Windows Server Update Services

Noun	Verbs	Function
WsusClassification	Set, Get	Manage which update classifications are synchronized
WsusComputer	Add, Get	Manage computers registered with the WSUS server
WsusDynamicCategory	Add, Remove, Get, Set	Manage WSUS dynamic categories
WsusProduct	Set, Get	Manage which products WSUS will synchronize updates for
WsusServer	Get	View the properties of a WSUS server
WsusServerCleanup	Invoke	Cleans up obsolete computer and updates registered with the WSUS server
WsusServerSynchronization	Set	Configures WSUS server synchronization with Microsoft Update or upstream server
WsusUpdate	Get, Approve, Deny	Manage specific updates

Technet24

Chapter 22

Upgrade and migration

Supported upgrade and migration paths

Active Directory

Active Directory Certificate Services

DNS

DHCP

File and storage servers

In most cases, organizations that are deploying Windows Server 2019 will be deploying the operating system into an existing Windows Server environment. While initial prototype and proof-of-concept deployments are likely to involve a small number of servers, at some point organizations will want to upgrade or migrate their existing Windows Servers and the workloads hosted on those servers to Windows

Technet24

Server 2019.

While some organizations deploy a new server operating system and upgrade existing servers as soon as possible, and other organizations wait until their current servers reach the end of extended support before planning a move, most organizations take a middle-ground approach. New workloads are deployed on the newest version of the operating system, while existing workloads are left on the platform they were deployed to unless there is a pressing reason to upgrade.

In this chapter, you'll learn about upgrading or migrating several core Windows Server roles and features from previous supported versions of the Windows Server operating system to Windows Server 2019.

SUPPORTED UPGRADE AND MIGRATION PATHS

A general rule about supported upgrade and migration paths is that you need to ensure that the product that you are upgrading from is still within Microsoft's extended support window and that you've applied all released service packs and software updates to the workload before attempting migration or upgrade. The difference between migration and upgrade is as follows:

- **Migration.** More common than an upgrade, a migration involves

moving roles and data from an existing server running a supported version of Windows Server to a freshly deployed server running Windows Server 2019. Migration is Microsoft's preferred option and fits better with the "treat servers as cattle and not as pets" philosophy of systems administration. This philosophy dictates that installation and configuration be as automated as possible, so instead of painstakingly crafting the configuration of an individual server, administrators can apply complex configurations quickly so that server upgrade or replacement is considered an uncomplicated routine operation rather than one that requires painstaking attention to arcane and not entirely understood processes.

- **Upgrade.** Sometimes called an in-place upgrade, this usually involves running the Windows Server 2019 setup routine from within an existing Windows Server 2012 R2 or Windows Server 2016 deployment. You cannot perform a direct upgrade to Windows Server 2019 from Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008. You perform upgrades when you need the new OS to use the same hardware or VM as the source OS, and you want to keep the roles and data that are on the source instance for the upgraded instance. Upgrades are only supported for certain roles. Upgrades are most often used in smaller organizations where the small number of servers means that the time required to build automated deployment and configuration tools and processes exceeds the time needed to deploy the organization's infrastructure manually.

Inside OUT

End of Support

When Windows Server 2003 reached the end of support, there were still millions of servers running the operating system. Even though organizations had ample warning of the approaching deadline, many simply were unable to assign the resources to transition from the now-unsupported operating system. While it was nice to hope that IT departments would be better prepared for the end-of-support deadline for Windows Server 2008 and Windows Server 2008 R2, this wasn't really the case either. Alarge number of organizations have likely adopted the option of extended security updates, but this option is expensive, coming in at around 75 percent of the cost of a Windows Server license per system, per year for a

maximum of three years.

Given the festival of unpreparedness that characterized the end of support for Windows Server 2003/2003R2 and Windows Server 2008/2008R2, I'm going to go out on a limb here and suggest that the end-of-support deadline for Windows Server 2012 and Windows Server 2012 R2 will also be reached with tens of millions of servers still running the operating system. Windows Server 2012 and Windows Server 2012 R2 are going to reach their end-of-support date on the 10th of October 2023. Windows Server 2016 reaches its end of extended support on the 12th of January 2027. Windows Server 2019 reaches end of support on January 9, 2029.

An in-place upgrade is supported for certain editions of Windows Server 2012 R2 and Windows Server 2016 to Windows Server 2019. When you perform an in-place upgrade, all the applications, data, roles, and features that were present on the source server remain available on the upgraded server. While the applications, data, roles, and features will remain present, that doesn't mean that they will all be functional. Prior to attempting an in-place upgrade, you should check whether applications that are hosted on the source operating system are compatible with Windows Server 2019.

Whether you can perform an upgrade depends on the edition of Windows Server that is the source and the target of the upgrade. Upgrading the operating system is supported in the cases listed in Table 22-1.

Table 22-1 Supported edition upgrades

Source Edition	Can be upgraded to

Windows Server 2012 R2 Standard	Windows Server 2019 Standard or Datacenter
Windows Server 2012 R2 Datacenter	Windows Server 2019 Datacenter
Windows Server 2016 Standard	Windows Server 2019 Standard or Datacenter
Windows Server 2016 Datacenter	Windows Server 2019 Datacenter
Windows Server 2012 R2 Essentials	Windows Server 2019 Essentials
Windows Server 2016 Essentials	Windows Server 2019 Essentials

Upgrading roles and features

Table 22-2 lists which server roles and features can be directly upgraded and which server roles and features, if migrated, require the role or feature be taken offline for a period of time.

Table 22-2 Upgrades and migration

Role	Direct upgrade from Windows Server 2012 R2 or Windows Server 2016	Migration involves downtime?

AD CS	Yes	Yes
AD DS	Yes	No
AD FS	No	Yes
AD LDS	Yes	No
AD RMS	Yes	Yes
DHCP	Yes	No
DNS	Yes	No
Failover Clusters	Cluster rolling upgrade supported for Windows Server 2012 R2 and Windows Server 2016	No
File and Storage Servers	Yes	No
Hyper-V	When a member of a Windows Server 2012 R2 failover cluster or a Windows Server 2016 failover cluster	No for Server 2012 R2 and Server 2016 with the scale-out file server role

Print and Fax Services	No	Yes
Web Server (IIS)	Yes	Yes
WSUS	Yes	Yes
Work Folders	Yes	No

Upgrades are constrained by the following additional limitations:

- Upgrades from one language to another are not supported. For example, you can't upgrade a French edition of Windows Server 2012 R2 to a German edition of Windows Server 2019.
- You cannot upgrade a Server Core installation to a Server with a Desktop installation. If the source operating system is running Windows Server 2012 R2, you can get around this by changing from Server Core to Server with a Desktop prior to performing the upgrade.
- You cannot upgrade a Server with a Desktop installation to a Server Core installation. If the source operating system is running Windows Server 2012 R2, you can get around this by converting to Server Core prior to performing the deployment.
- You cannot upgrade from a licensed version of a Windows Server

operating system to an evaluation edition. Microsoft recommends that you only install evaluation editions as clean installations.

Converting evaluation version to licensed version

Although it is far from best practice, many organizations that deploy the 180-day evaluation version of Windows Server as part of a project prototype find, several months later, that the prototype has morphed into a critical component in the organization's workflow. Rather than reinstall Windows Server and migrate the prototype deployment to a properly licensed, long-term production platform, it is instead possible to convert an evaluation version of Windows Server 2019 to a properly licensed version by performing the following steps:

1. From an elevated command prompt, run the command *slmgr.vbs /dlv*. This will display the current computer's licensing information as shown in Figure 22-1.

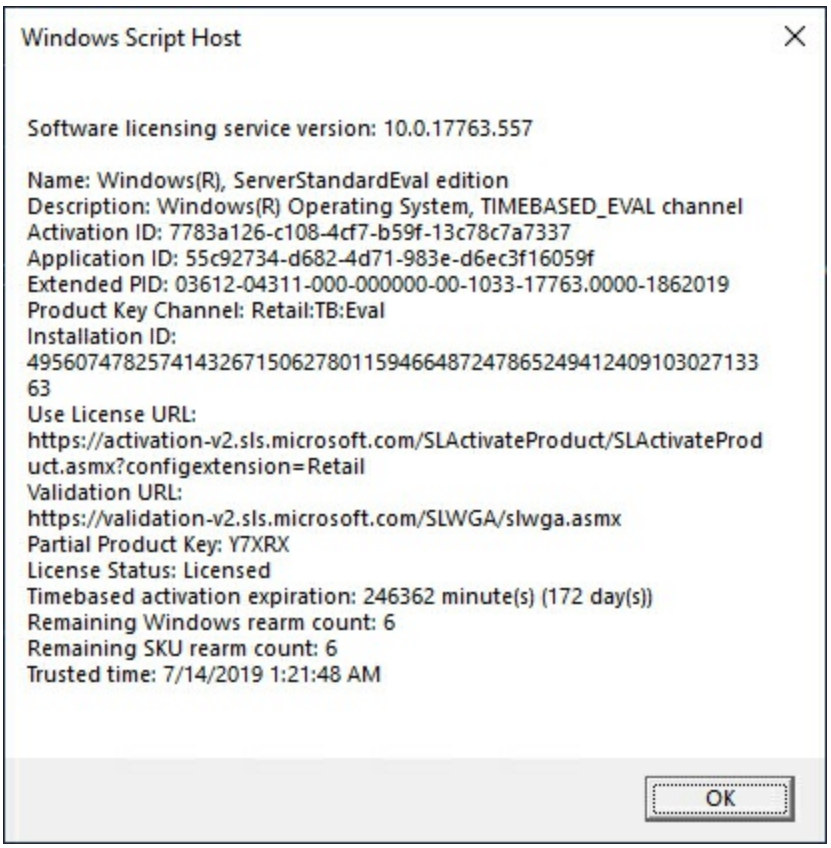


Figure 22-1 Software licensing information

2. Determine the current edition by running the command *DISM /online /Get-CurrentEdition*. This will display the edition ID, such as *ServerStandard* or *ServerStandardCore*. You will need this edition ID for the next step.

3. Once you've determined the edition, run the command *DISM /online /Set-Edition:<editionID> /ProductKey:XXXXX-XXXXX-XXXXX-XXXXX-XXXXX /AcceptEula*.

Inside OUT

Evaluation domain controllers

You can't switch a computer that is running as a domain controller on an evaluation version of Windows Server 2019 to a retail version of the operating system. If you have a domain controller running on an evaluation version that you wish to convert, Microsoft's recommendation is that you install an additional Windows Server 2019 domain controller running a non-evaluation and licensed version of Windows

Server 2016, remove AD DS from the server running the evaluation version, perform the switch, and then reinstall AD DS.

Upgrading editions

You can also use DISM to upgrade from one edition of Windows Server 2019 to another. For example, if you want to upgrade from the Standard edition of Windows Server 2019 to the Datacenter edition, specify the ServerDatacenter or ServerDatacenterCore product edition and enter the appropriate product key using the *DISM /online /Set-Edition:<editionID> /ProductKey: <productkey>* command.

You can use *slmgr.vbs* to switch between a retail license and a volume-license. To do this from an elevated prompt use the command *slmgr.vbs /ipk <key>* where *<key>* is the volume license key. If you are trying to convert an evaluation version of Windows Server 2019 to a volume licensed version, you must switch to a retail version, as described earlier.

Windows Server Migration Tools

The Windows Server Migration Tools (WSMT) is a Windows Server 2019 feature that you can use to migrate a specific set of server roles, features, shared folders, operating system settings, and data from a source computer running the Windows Server 2008, Windows Server 2008 R2, Windows Server 2012,

Windows Server 2012 R2, or Windows Server 2016 operating systems to a computer running Windows Server 2019.

You can use the WSMT to migrate the following roles and features:

- Active Directory Certificate Services
- Active Directory Federation Services
- Active Directory Rights Management Services
- File and storage services
- Hyper-V
- Network Policy Server
- Remote Desktop Services
- Windows Server Update Services
- Cluster Roles
- DHCP Server
- Web Server (IIS)

> **Inside OUT**
>
> **Storage Migration Services**
>
> While WSMT supports migration of file servers, Storage Migration Services is a new and free tool available from Microsoft that provides a more comprehensive way of migrating file servers. You'll learn more about Storage Migration Services later in the chapter.

The WSMT supports:

- Migration from physically deployed operating systems to virtually deployed operating systems.

- Server Core and Server with a GUI installation option on Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, or Windows Server 2019 as the source operating system.
- Server Core and Server are supported with a GUI installation option for Windows Server 2019 as the destination operating system.
- You can't use WSMT with Server Core on Windows Server 2008 because the .NET framework is not available on the Server Core installation option of Windows Server 2008.
- WSMT can only be used if the source and destination operating system are using the same UI language. The system UI language is the one used when setting up the operating system.

Inside OUT

Where possible, Server Core

The Server Core installation option is the Windows Server 2019 default. When migrating to Windows Server 2019, strongly consider shifting as many workloads to hosts running the Server Core installation option rather than the Desktop Experience installation option. Unless you are in an environment with a small number of servers, it's likely that you already use remote consoles and PowerShell to manage servers (and rarely, RDP) to the server desktop. You won't be able to do it with all servers, especially those running applications that require desktop components, but you will be able to do it with a good number of servers, such as those hosting boring roles such as functioning as file servers, DHCP servers, DNS servers, and Domain Controllers. With the ongoing improvements in Windows Admin Center, a Server Core first-deployment strategy makes more sense. Though having said all that, every time I ask audiences at an event if anyone is running Server Core in their environments, only a fraction of the audience puts up their hands. Hopefully, Windows Admin Center changes this.

The WSMT has the following requirements:

- Must be deployed using an account that has local administrator privileges on the source and destination server.
- The source server must have at least 23 MB of free space to store

the WSMT deployment folder.

- The source server must have the .NET Framework 3.51 installed.
- Windows PowerShell must be present on the source server. By default, PowerShell is included with Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, and Windows Server 2016. By default, PowerShell is not included on servers running the Windows Server 2008 operating system, though it's likely that it has been installed at some point prior to the deployment of that operating system in your environment.
- You should always use the most recent version of the WSMT.

Installing the WSMT

You can install the WSMT on the destination server using the *Install-WindowsFeature Migration* PowerShell command or by installing the role using the Add Roles And Features Wizard, as shown in Figure 22-2. When you do this, the WSMT is deployed in the *%WinDir%\System32\ServerMigrationTools* folder on the destination server.

Technet24

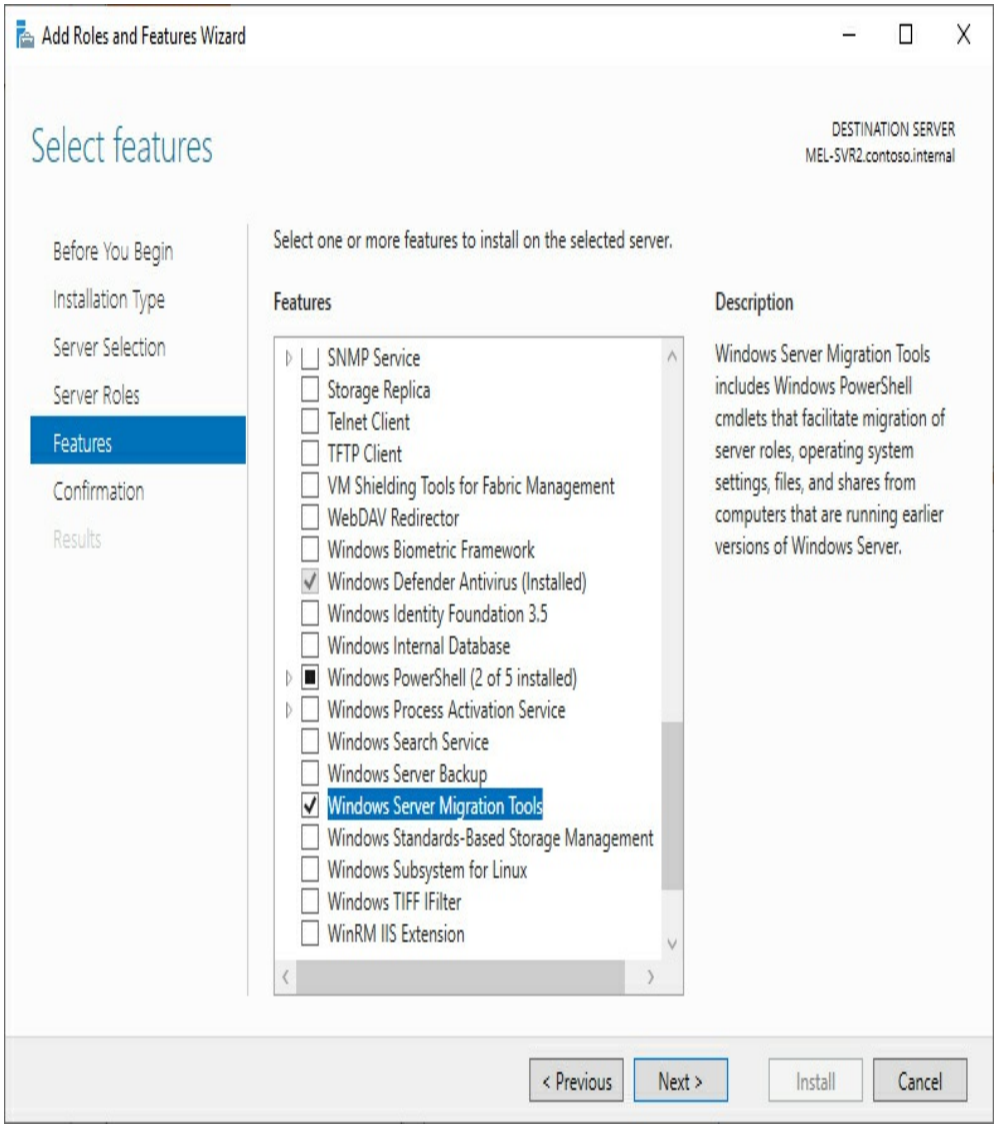


Figure 22-2 Windows Server Migration Tools

Once you've installed the WSMT, you need to configure a deployment folder from which you will deploy the appropriate version of the WSMT. Rather than deploying a generic version of the tools, deploy a version tailored for the source operating system. For example, to create a set of tools for a source computer running

Windows Server 2012 R2, issue the command as shown in Figure 22-3:

Click here to view code image

```
Smigdeploy.exe /package /architecture amd64 /os WS12R2 /path
<path to share>
```

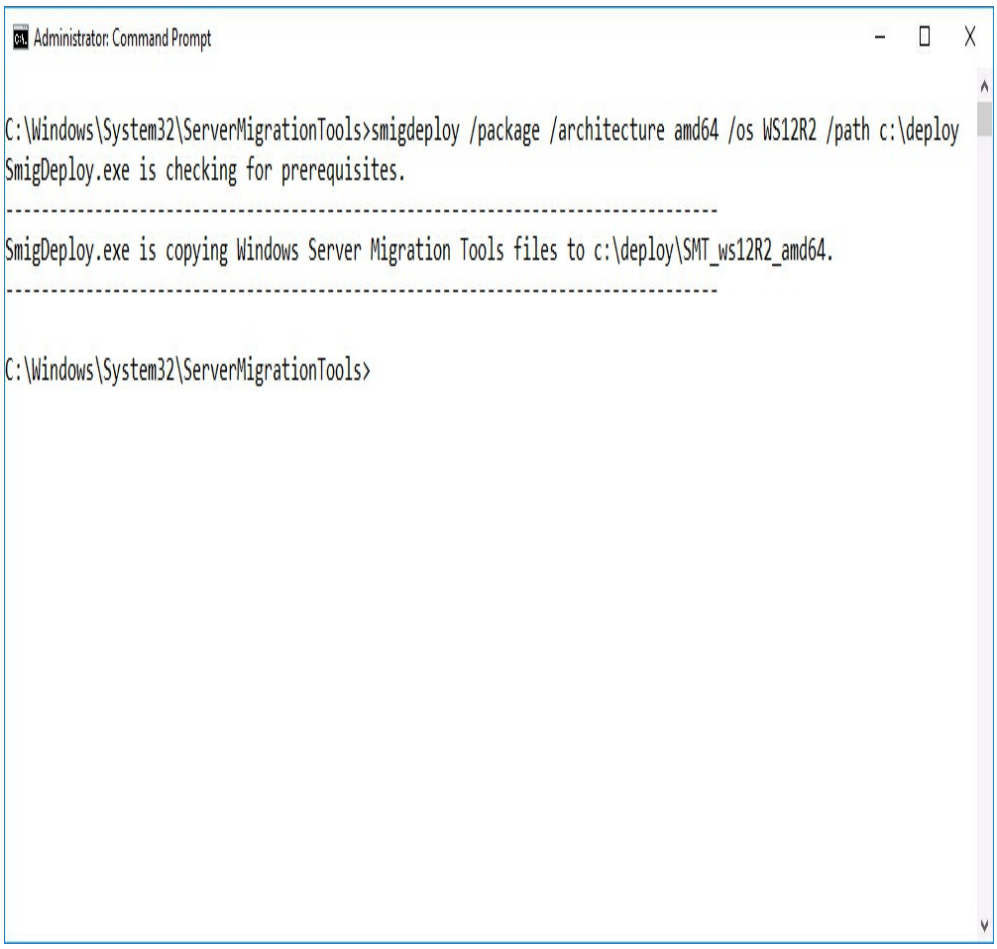


Figure 22-3 Create migration tools package

You can either copy the contents of the deployment folder to the source computer or enact the tools remotely over a network connection by running

smigdeploy.exe from an elevated command prompt. Installing the tools on a source server running the GUI option will give you a special migration tools PowerShell session, as shown in Figure 22-4.

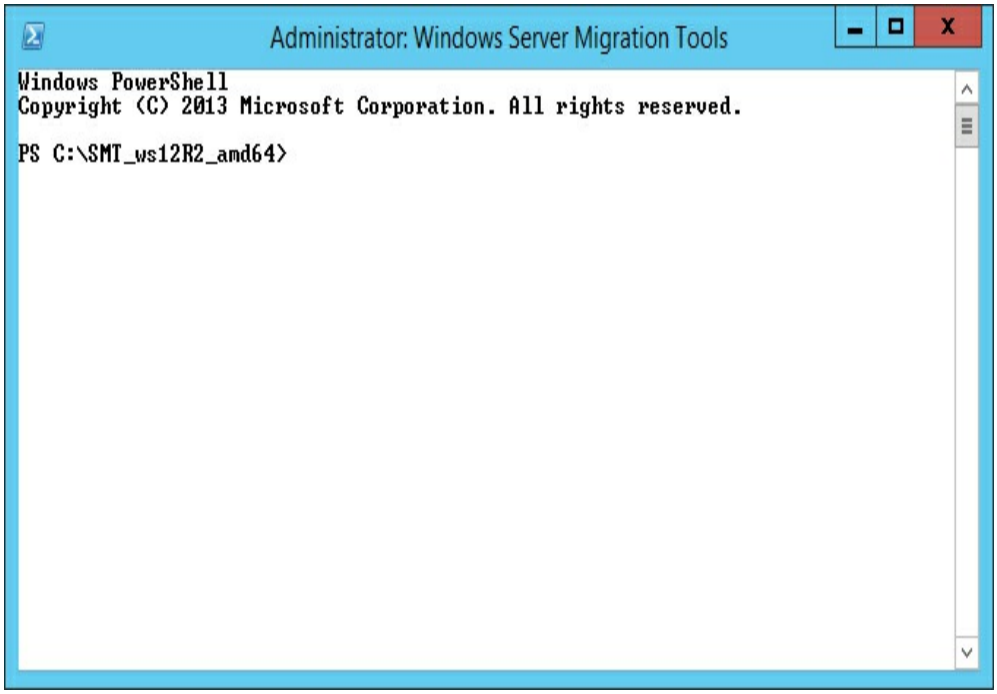


Figure 22-4 Server Migration Tools running on Windows Server 2012 R2

If you are running the tools on a computer that is configured with the Server Core installation option, open an elevated PowerShell session, and run the following command to add the WSMT cmdlets:

Click here to view code image

```
Add-PSSnapin Microsoft.Windows.ServerManager.Migration
```

ACTIVE DIRECTORY

Microsoft's recommended strategy for upgrading domain controllers in a domain or forest to Windows Server 2019 is to

- Add member servers running Windows Server 2019 to a domain
- Promote those member servers to domain controller
- Transfer FSMO roles across to the domain controllers running Windows Server 2019 from the domain controllers running previous versions of the Windows Server operating system
- Demote those domain controllers running the older version of the operating system

In some circumstances, you might need to perform an upgrade of a domain controller rather than using the recommended strategy. Prior to attempting the upgrade, manually run *adprep /forestprep* and *adprep /domainprep* from the *\support\adprep* folder of the Windows Server 2019 installation media.

You should upgrade the domain controllers in the forest root domain and then upgrade domain controllers in any child domains. Once you've upgraded all the domain controllers in your organization, you can then raise the domain functional level of each domain in the forest to Windows Server 2016 using the Active Directory Domains and Trusts console, as shown in Figure 22-5. Remember that unlike previous versions of Windows Server, the highest domain and forest functional level

Technet24

that you can configure for Windows Server 2019 domain controllers is Windows Server 2016.

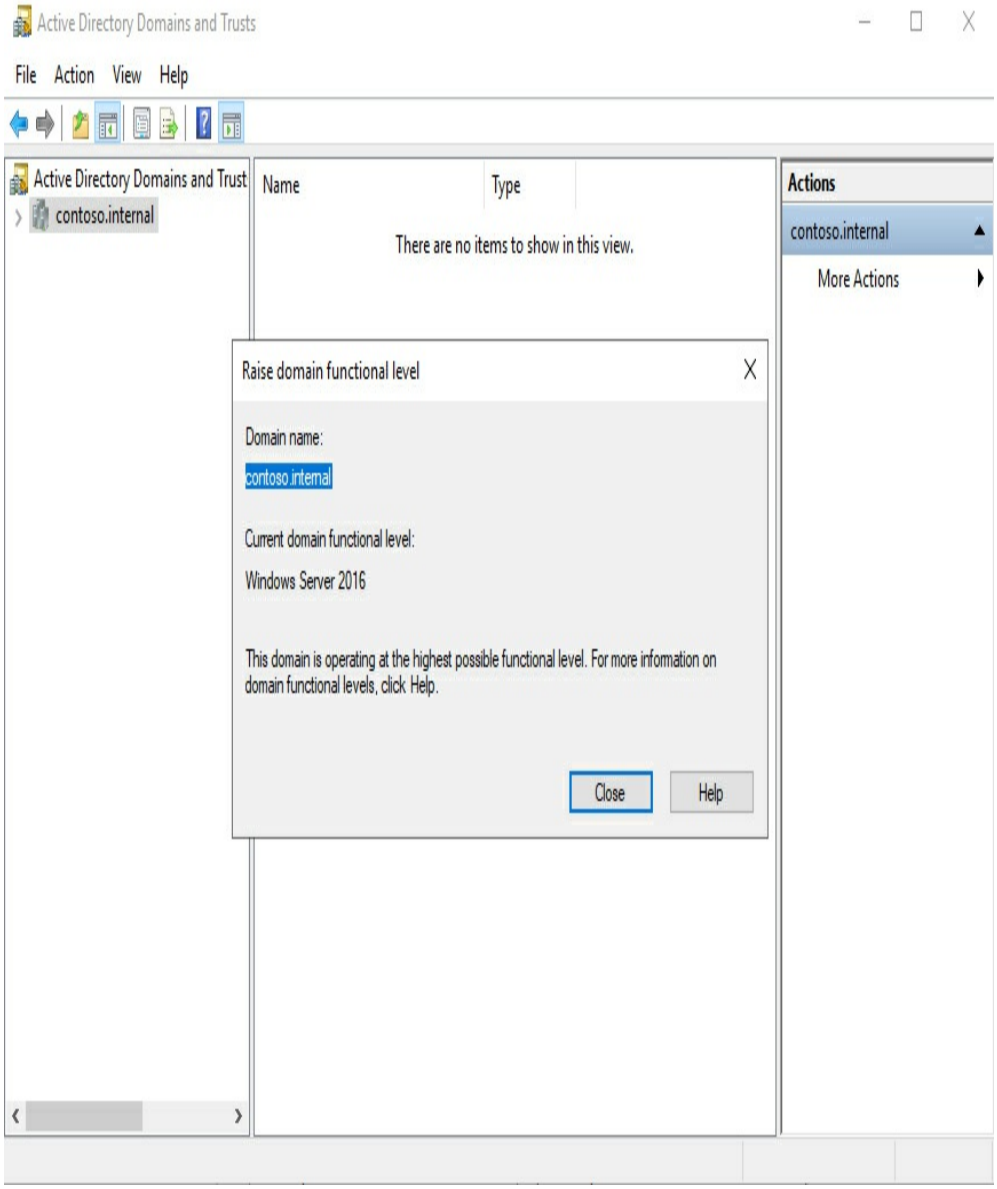


Figure 22-5 Domain functional level

Once you have raised the domain functional level of each domain in the forest to Windows Server 2016, you

can raise the forest functional level to Windows Server 2016 using the Active Directory Domains and Trusts console, as shown in Figure 22-6. Remember, there is no Windows Server 2019 functional level.

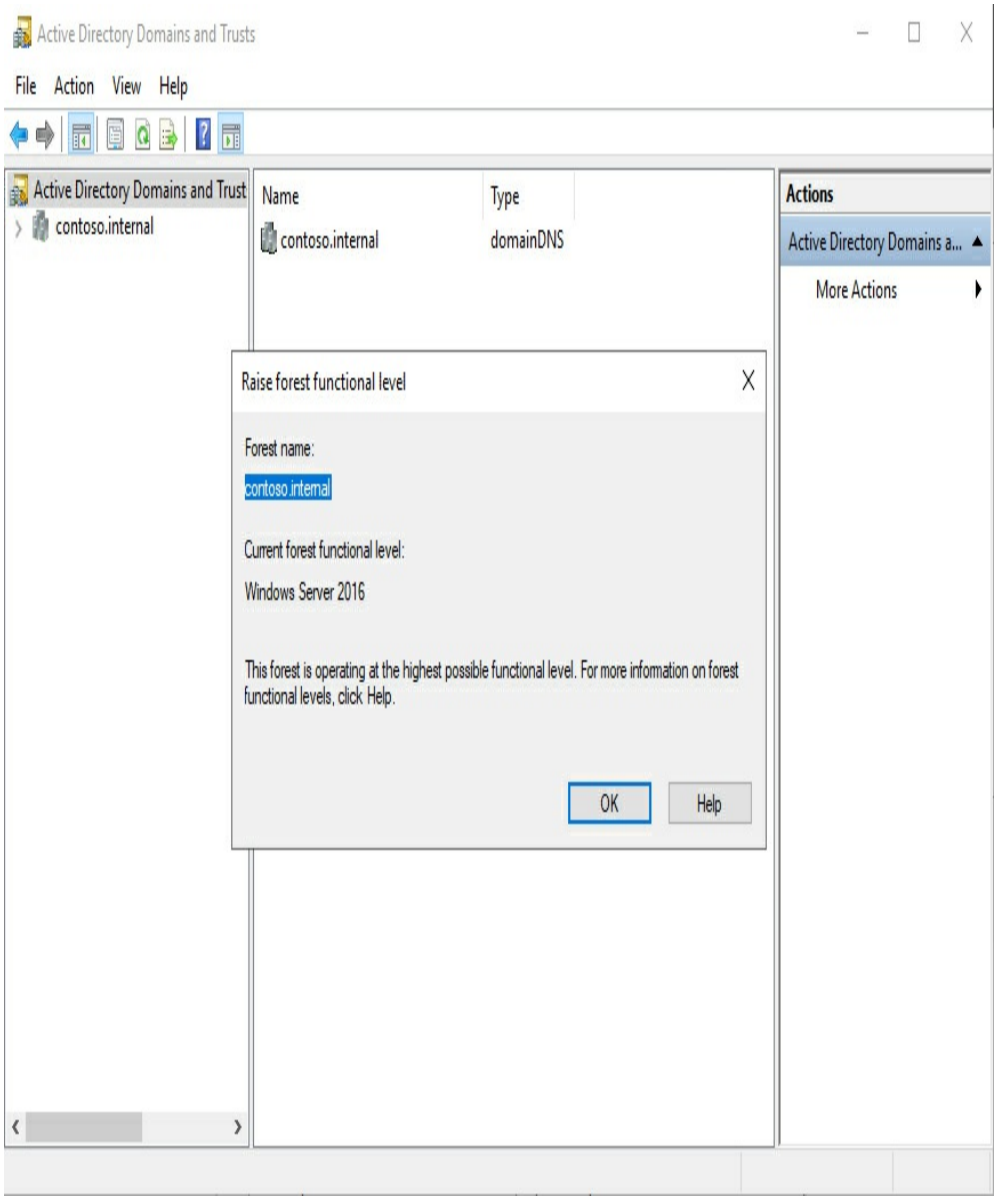


Figure 22-6 Forest functional level

More Info

Upgrading domain controllers

You can learn more about upgrading domain controllers to Windows Server 2019 at https://techcommunity.microsoft.com/t5/ITOps-Talk-Blog/How-to-Migrate-Windows-Server-2008-R2-FSMO-roles-to-Windows/ba-p/538377.

FRS to DFSR migration

Depending on the age of your organization's Active Directory deployment—especially if a domain once functioned at the Windows Server 2003 functional level—and how diligently Active Directory upgrades were performed, it's possible that SYSVOL replication is configured to use the older inefficient File Replication Service, rather than the more efficient Distributed File System.

You can check whether SYSVOL replication is configured to use FRS or DFS using the *dfsrmig.exe* utility with the */globalstate* parameter. If the output state of the command shows Eliminated, as in Figure 22-7, then SYSVOL already uses DFS for replication.

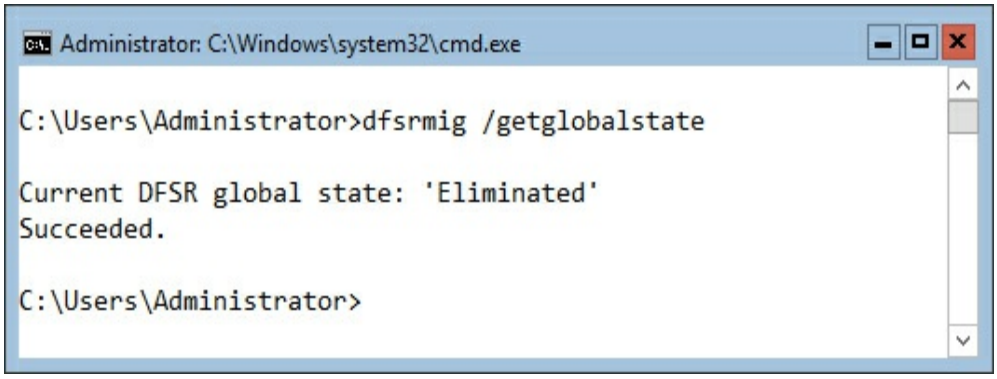


Figure 22-7 SYSVOL migration state

If the output of dfsrmig.exe instead shows that migration has yet to be initialized, you should update the SYSVOL replication. You should only use FSR for SYSVOL replication if your organization still has domain controllers running the Windows Server 2003 or Windows Server 2003 R2 operating systems. At this stage, it really shouldn't be running those OS versions, should it?

To upgrade SYSVOL replication to DFS from FRS, perform the following steps:

1. Sign in to a domain controller with credentials that are a member of the Domain Admin or Enterprise Admin groups.
2. Start PowerShell and then run the command *dfsrmig /setglobalstate 1*. Depending on the number of domain controllers in the domain, it can take more than an hour to reach the prepared state. You can verify that the prepared state has been reached by running the command *dfsrmig /getmigrationstate*. When preparation is completed, the utility will report that the global state is Prepared.
3. Next, set the DFSR state to Redirected. You do this using the command *dfsrmig /setglobalstate 2*. You use *dfsrmig /getmigraitonstate* to verify that the DFSR state is set to Redirected.
4. Once the DFSR state is set to Redirected, run the command *dfsrmig /setglobalstate 3*. This will finalize the transition of SYSVOL to the Eliminated state, and DFS will be used for replication instead of FRS.

Migrating to a new forest

Many organizations have had their current Active

Technet24

Directory environment for more than a decade. The older an Active Directory environment is, the more likely it is to have an inefficient domain and OU structure, derelict user and computer accounts, unutilized sites, and forgotten group policy objects and groups. Just as Microsoft recommends performing fresh installations rather than in-place upgrades for individual computers, it may make more sense to migrate to a new forest than it does to remain with an existing Active Directory forest. This is especially true if an organization is considering moving toward a secure administrative model where the privileged accounts in a production forest are unprivileged standard accounts in the secure administrative forest.

The Active Directory Migration Tool (ADMT) is a special tool that you can use to migrate Active Directory objects to and from any Active Directory environments, from moving accounts within a forest, to migrating accounts to a new Active Directory forest.

You can use ADMT to perform an interforest restructure or an intraforest restructure. An intraforest restructure allows you to modify your current domain and forest structure to reduce complexity. An interforest restructure involves creating a brand new, simplified forest structure using Active Directory objects from your existing forest.

Table 22-3 lists the differences between each restructure type.

Table 22-3 Migration differences

Consideration	Interforest restructure	Intraforest restructure
Object preservation	Objects, such as accounts, will be cloned. The original object remains in the source forest.	User and group objects are migrated and are no longer present in the source location. Computer and managed service account objects are copied and remain in the source location.
SID history maintenance	SID history maintenance is optional.	SID history is mandatory for user, group, and computer accounts. It is not required for managed service accounts.
Password retention	Password retention is optional.	Passwords are always retained.

Technet24

Profile migration	Local profiles can be migrated using ADMT.	Local profiles are automatically migrated, and account GUID is preserved.

The Active Directory Migration Tool requires that you have access to a SQL Server instance. This can be a full SQL Server deployment, or a local SQL Server express instance. The ADMT includes the following wizards, shown in Figure 22-8.

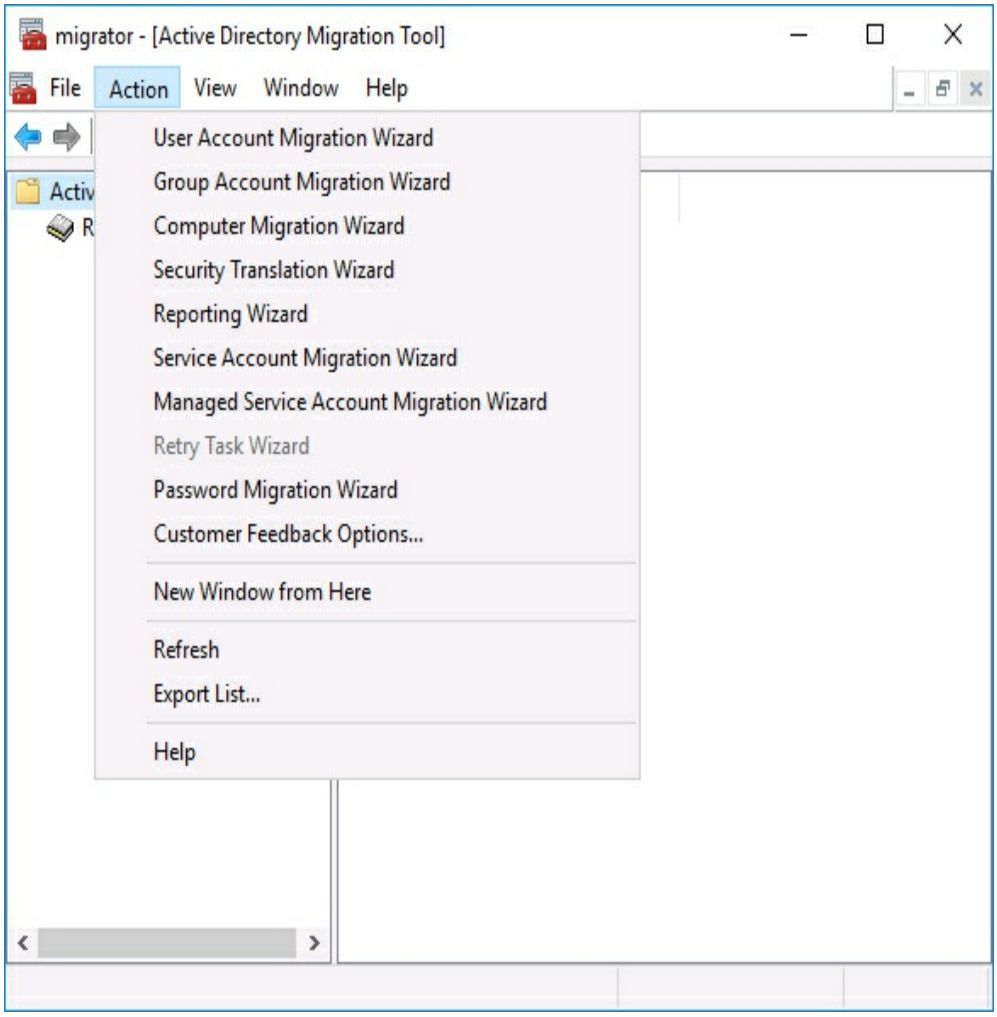


Figure 22-8 Active Directory Migration Tool

- **User Account Migration Wizard.** Allows you to perform an intraforest or interforest migration of user account objects.
- **Group Account Migration Wizard.** Allows you to perform an intraforest or interforest migration of group account objects, including group membership information.
- **Computer Migration Wizard.** Allows you to perform an intraforest or interforest migration of computer account objects.
- **Security Translation Wizard.** Wizard that allows you to transfer security permissions during object migration.

- **Reporting Wizard.** Generates the following reports:
  - **Migrated Users, Groups And Managed Service Accounts report.** Summarizes result of the migration of user, group, and managed service accounts.
  - **Expired Computers Report.** List of computers found that have expired computer account passwords.
  - **Account References Report.** List of accounts that have been assigned permissions for resources on a specific computer.
  - **Name Conflicts Report.** A list of user accounts and groups that are present in both the source and target domains.
- **Service Account Migration Wizard.** Allows you to perform intraforest or interforest service account migration.
- **Managed Service Account Migration Wizard.** Allows you to perform intraforest or interforest managed service account migration.
- **Password Migration Wizard.** Allows you to perform intraforest or interforest password migration.

More Info

Active Directory Migration Tool

You can learn more about the Active Directory Migration Tool at https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/cc974332(v=ws.10).

ACTIVE DIRECTORY CERTIFICATE SERVICES

You can use WSMT to migrate an AD CS deployment from one server to another. The WSMT method works best when you are migrating a standalone root or

subordinate CA. Prior to performing AD CS migration, you need to be aware of the following:

- It will be necessary to either turn off the source CA or stop the source server's certificate services service.
- The destination CA must have the same name as the source CA. The CA name is separate from the computer name.
- While the computer name of the destination CA does not need to match that of the source CA, the destination CA name must not be identical to the name assigned to the destination computer. For example, *ContosoRootCA* is the name of the source CA which is hosted on a computer named *MEL-CA*. When migrated, the destination CA will be called *ContosoRootCA* and might be hosted on a computer named *MEL-NEW-CA*. The destination computer just cannot be named *ContosoRootCA*.
- During migration, the CA will be unable to issue certificates or publish CRLs.
- To ensure that revocation status checking can occur, ensure that a CRL is published prior to migration that will remain valid to a point in time after the expected completion of the migration.
- Ensure that at least one of the CRL distribution points specified for existing certificates remains available during migration. Enterprise CAs publish to Active Directory, but standalone CAs may only publish to the local server. If the CRL distribution point is unavailable, clients will be unable to perform revocation checks.
- The Authority Information Access (AIA) and CRL distribution point extensions of previously issued certificates may reference the name of the source CA. This is the default setting and will only be different if you explicitly changed it.
- If migrating an enterprise CA or standalone CA on a computer that is a member of a domain, the account used to perform the migration must be a member of the Enterprise Admins or Domain Admins group.

Inside OUT

AD CS upgrades

You can directly upgrade a computer running Windows Server 2012 R2 or Windows Server 2016 that hosts the AD CS role to Windows Server 2019. It's substantially easier to do it this way if you have a complex configuration that includes an enterprise CA, CAweb enrollment, online responders, network device enrollment, and the certificate enrollment web services— all on the same host,. Given the amount of work involved, you might consider a hop upgrade or Windows Server 2019 if you are hosting this role and the role service on Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012. When it comes to AD CS—especially given AD-integrated CAs—servers are definitely far more in the category of "pets" rather than "cattle."

CA migration can be carried out in three general phases:

1. Preparation
2. Migration
3. Verification and post migration tasks

Preparation

To prepare the source server, perform the following steps:

1. **Export the CA templates list.** This is the list of templates that the CA is configured to issue. The master list of templates is stored within Active Directory. You can export this list from the Certificate Services console by right-clicking the *Certificate Templates* node and selecting Export. You can also do this using the command *certutil.exe -catemplates > catemplates.txt command.*
2. Make a note of the CA's Cryptographic Service Provider (CSP) and which signature algorithm is used. This is visible on the CA's Properties dialog box, as shown in Figure 22-9. You can also do this using the command *certutil.exe -getreg ca\csp* > csp.txt*.

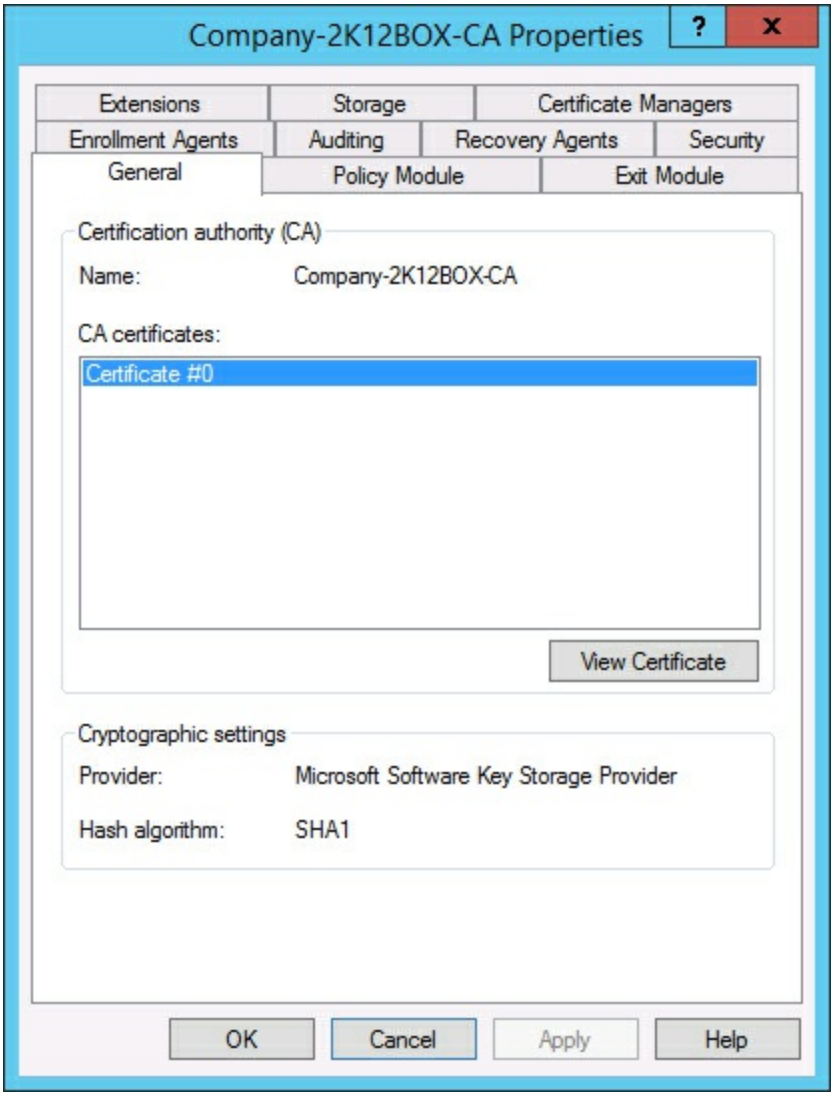


Figure 22-9 CSP and signature algorithm

3. Back up the source server CA, including the private key, CA certificate, Certificate Database, and Certificate Database Log. You can do this using the Certification Authority Backup Wizard, as shown in Figure 22-10. Copy the contents of the backup folder to a location that will be accessible to the destination server.

Technet24

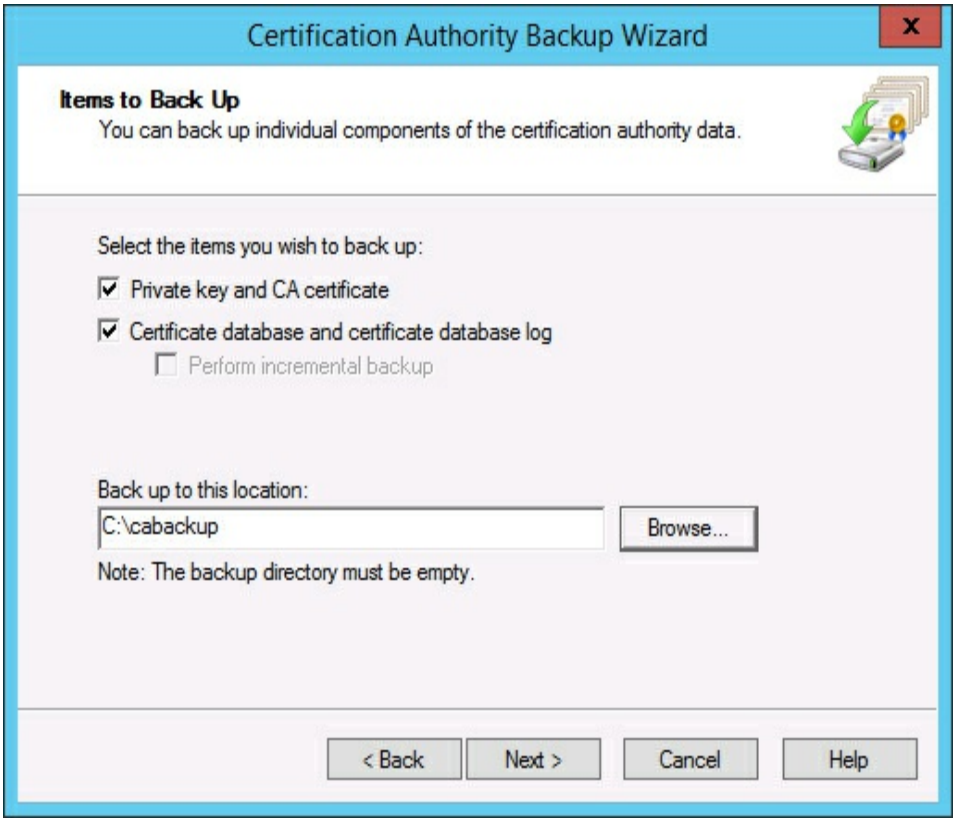


Figure 22-10 CA Backup Wizard

4. Stop the Certificate Services service. You can do this using the Services console, as shown in Figure 22-11, or by using the *Stop-Service certsvc* Windows PowerShell command. This will block the current certificate server from issuing additional certificates.

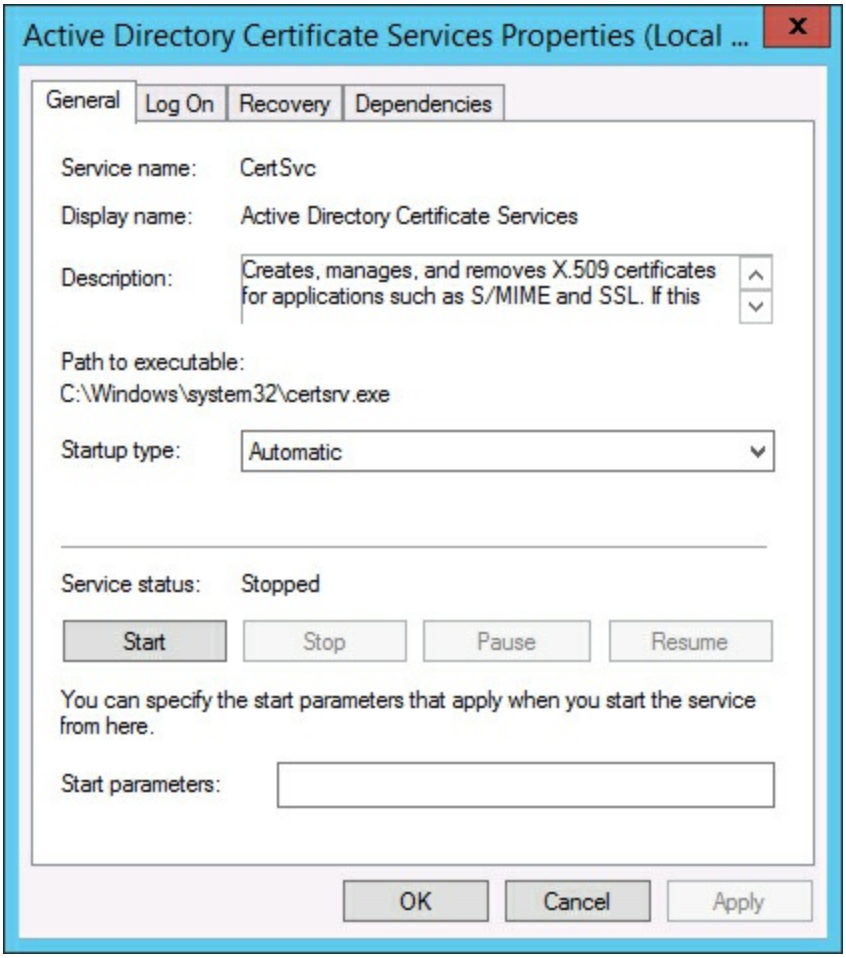


Figure 22-11 Stop certificate services service

5. Back up the registry settings of the CA using the following command, and place it in a location that is accessible to the destination server:

   Click here to view code image

   ```
   reg export
   HKLM\SYSTEM\CurrentControlSet\Services\CertSvc\
   Configuration
   <output file>.reg
   ```

6. If the source CA uses a custom *CAPolicy.inf* file, you should copy this file to a location that is accessible to the destination server.

Migration

Once you have performed the necessary preparation steps, you should perform the following steps to complete the migration of the CA:

1. Once you have backed up the source CA and its registry settings, remove the CA role service from the source server. This is important because removing the role service from the server will also clean up data stored in Active Directory associated with the CA. If you don't perform this clean up, you will encounter problems when introducing the migrated CA that shares the original CA's name.
2. Removing the role service from the CA does not delete the CA database, private key, or CA certificate. This means that if the migration fails, you will be able to restore the CA by reinstalling the certificate services role on the server.
3. Microsoft also recommends manually removing the source server from the domain prior to deploying the destination server. This is especially important if the source and destination server share the same computer name.
4. Change the name of the destination server to match the name of the source server.
5. Join the destination server to the domain.
6. Import the CA certificates into the local computer's Personal certificate store.
7. Once the computer is domain joined, add the CA role service to the destination server. Ensure that you install the server in the same manner, such as installing an Enterprise Root or Standalone Subordinate CA, depending on the properties of the original source CA. During the installation routine on the Set Up Private Key page, utilize the Use Existing Private Key option, and select the existing private key. Complete the setup wizard.
8. Use the Certification Authority snap-in to restore both the certificate

database and the certificate database log from the backup taken on the source server. You will need to restart the certificate services server.

9. Import the registry settings exported from the source CA. If the source CA and the destination CA share the same name and the same file paths, then this can be done without modification. If the source CA has a different name from the destination CA or uses different file paths, you'll need to manually alter these registry settings. Prior to importing the registry, stop the certificate services service. Use the command *reg import backup.reg* to import the registry file. Restart the certificate service.
10. Import the list of CA templates that are issued from the CA using the command *certutil -setcatemplates*, and then provide a list of templates that you backed up from the source server in comma-separated format.

Verification and post migration tasks

To verify that the CA has migrated successfully, verify that you can successfully request and enroll a new certificate. You should also verify that you can successfully publish a CRL. Performing these tasks is covered in Chapter 12, "Active Directory Certificate Services."

More Info

Migrate Active Directory Certificate Services

You can learn more about migrating Active Directory Certificate Services to Windows Server 2019 at https://techcommunity.microsoft.com/t5/ITOps-Talk-Blog/Step-By-Step-Migrating-The-Active-Directory-Certificate-Service/ba-p/697674.

DNS

The simplest way to migrate DNS is where DNS zones

are Active Directory integrated. In that scenario, you add member servers running Windows Server 2019 to an existing domain, promote those servers to domain controllers, and make sure that you add the DNS role. Any existing AD integrated zones that are configured to replicate to other domain controllers in the domain will replicate to the newly promoted domain controller. Remember to add custom partitions if you have AD integrated zones that only replicate to specific domain controllers.

The trick to migrating primary zones from one DNS server to another is to first configure the destination server to host a secondary zone of the primary zone. Once the zone has replicated to the intended destination server, convert the zone from a secondary zone to function as a primary zone. You can do this by performing the following steps:

1. On the General tab of the secondary zone Properties, shown in Figure 22-12, click Change.

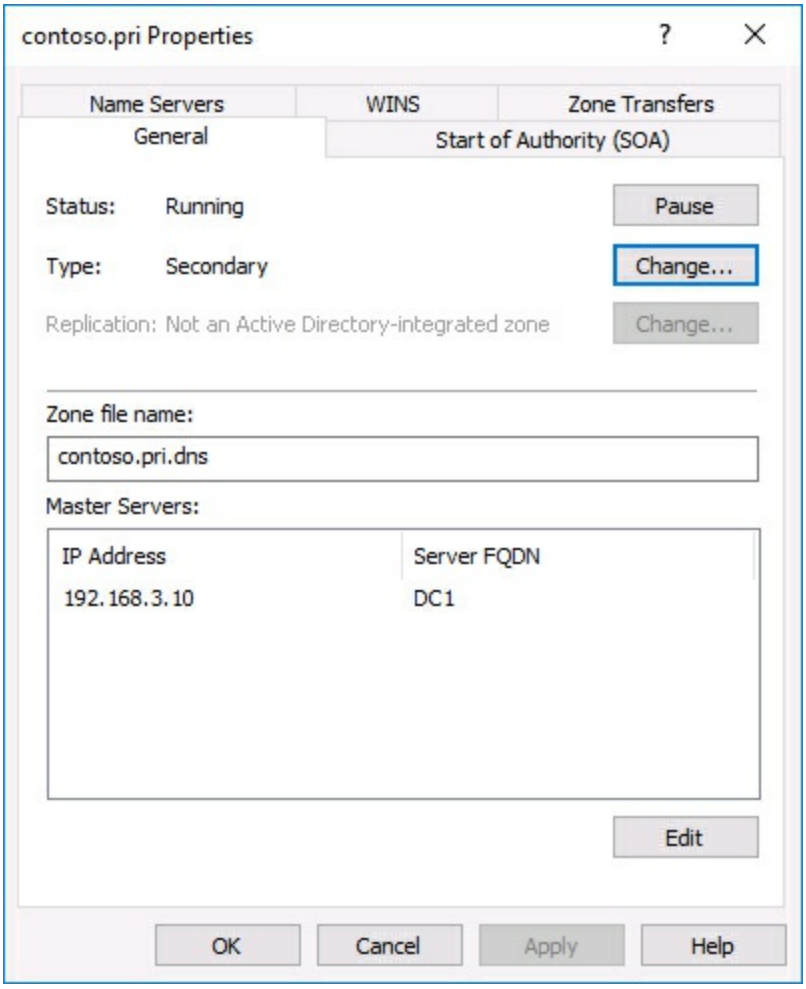


Figure 22-12 Zone General Properties tab

2. In the Change Zone Type dialog box, shown in Figure 22-13, select Primary Zone and then click OK.

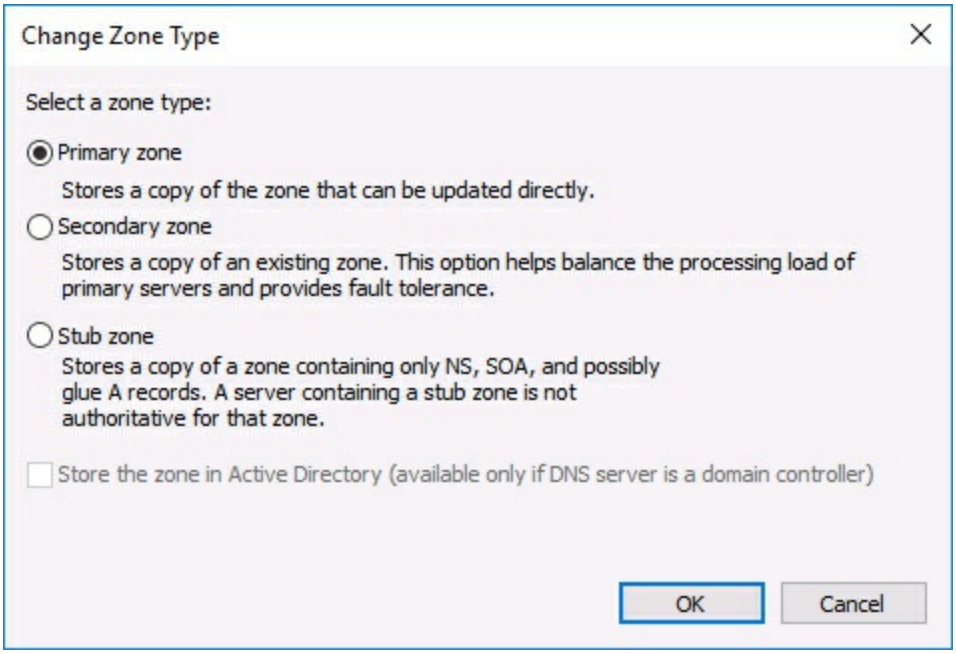


Figure 22-13 Change Zone Type

DHCP

You can use the WSMT to migrate a DHCP server, scopes, and settings from computers running both the Server Core and Server with a GUI installation options of Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019 to Windows Server 2019.

To perform a migration of a DHCP server using WSMT, you need access to accounts with the following permissions:

- Permission to authorize the DHCP server in Active Directory
- Membership of the local Administrator group on the computers

that are both source and destination DHCP servers

- Permission to write data to the WSMT migration store location

DHCP migration follows the following general steps:

1. Preparation
2. Migration
3. Verification and post migration tasks

Preparing to migrate DHCP

Before you start the process of migrating a DHCP server and its associated scopes and settings to Windows Server 2019, you'll need to perform several preparation steps.

1. Determine which current DHCP servers you wish to migrate and which servers running Windows Server 2019 will host the migrated DHCP server roles, associated scopes, and settings.
2. Ensure that all critical updates, update rollups, and service packs have been installed on the source and destination servers.
3. Deploy the Windows Server Migration Tools (WSMT) on the source and destination servers.
4. Install the DHCP role service on the destination server. This service must be present on the destination server prior to beginning the migration.
5. Ensure that the source and destination servers have the same number of network adapters. Determine which subnet each network adapter is bound to on the source server and ensure that the destination server has network adapters bound to the same subnets. This will ensure that the migrated DHCP server can service the same subnets as the source DHCP server.
6. Determine the IP address strategy for the destination server. DHCP servers should use static IP addresses on network interfaces that are used

by the DHCP server service to manage DHCP traffic. Options include:

1. **Assigning the destination server different static IP addresses that are on the same subnets as the static IP addresses assigned to network interfaces on the source server.** Doing this will allow you to perform the migration with both the source and destination server online.
2. **Assign the destination server the same IP addresses that are used for network interfaces for the source server.** While this is possible, it will mean that the source and destination server cannot be online at the same time without causing IP address conflicts, or it would require extensive preparation of network virtualization. Because existing clients will be able to renew their IP address information after migration from a destination server with a new IP address on the same subnet, you should only keep the same IP addresses on the source and destination server if there are substantive reasons for doing so.

7. If the source server does not use the default path for the DHCP server database, such as having the database stored on a separate volume, the destination server will need to have a volume configuration that matches that of the source server. For example, if the DHCP database is stored on volume H: of the source server, the destination server will need a properly formatted volume H: for migration to be successful.
8. Ensure that the migration store location is accessible to both the source and destination servers.
9. Ensure that both the source and destination servers have network access to the AD DS environment.
10. Perform a backup of the DHCP server from the DHCP console, rather than the Windows Server Backup, as shown in Figure 22-14. You can do this by right-clicking the DHCP server node, selecting Backup, and specifying a location for the backup.

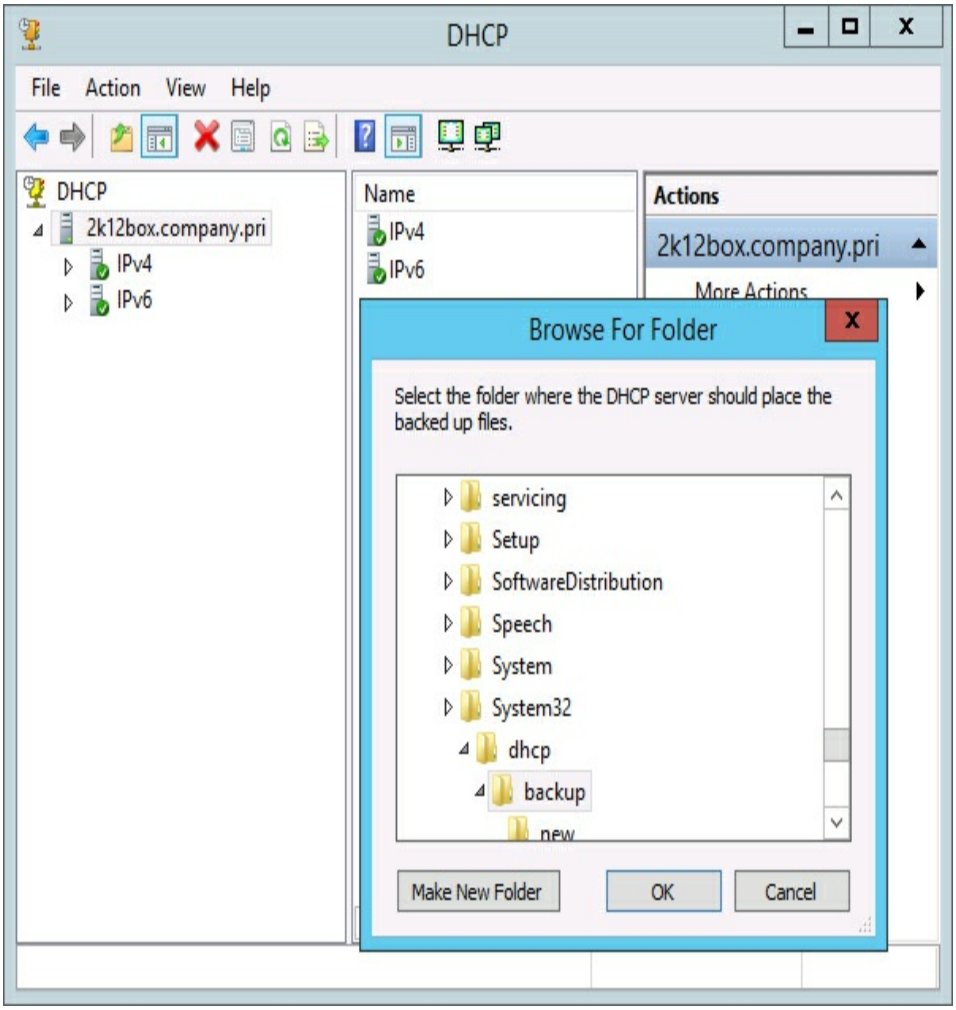


Figure 22-14 Backup DHCP database

11. Once the backup is complete, stop the DHCP service on the source and the destination server, either by using the Services console as shown in Figure 22-15, or by running the *Stop-Service DHCPserver* Windows PowerShell cmdlet.

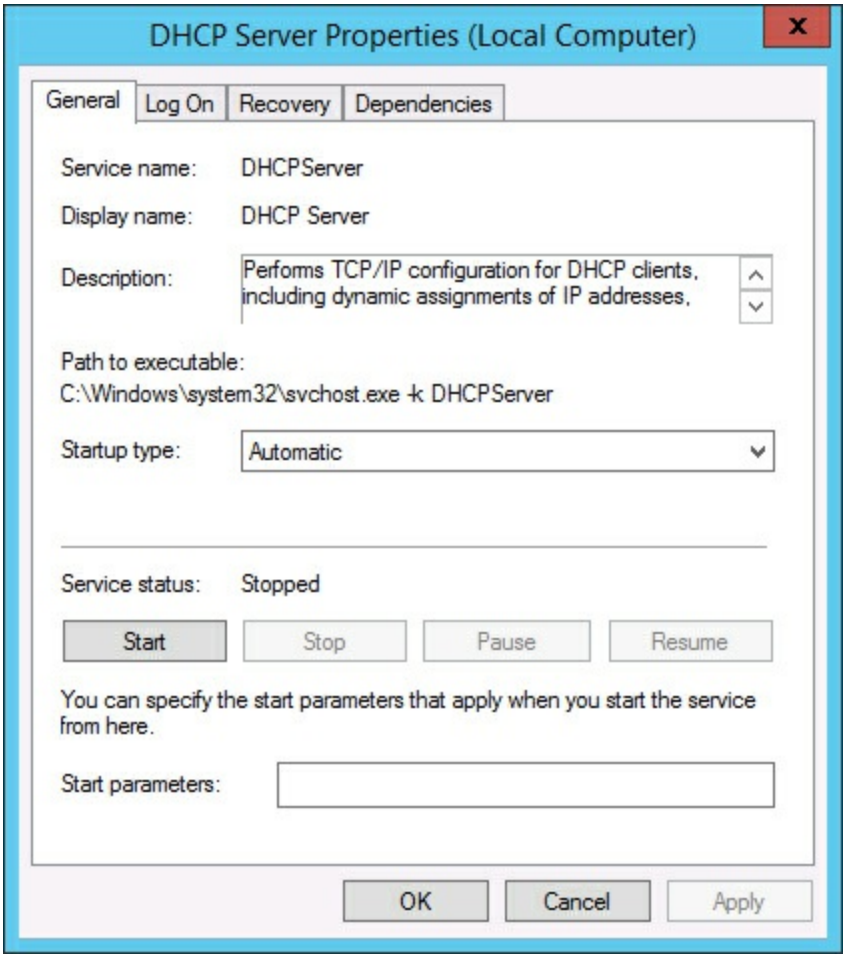


Figure 22-15 Stop DHCP Server service

Migration

Once the source and destination servers are prepared, the next step is to open the Windows Server Migration Tools PowerShell session. You perform the migration using the *Export-SmigServerSetting* PowerShell cmdlet. Specific parameters of the *Export-SmigServerSetting* cmdlet that are relevant to DHCP server migration include:

- Use the *FeatureID* DHCP to specify that the DHCP server role is

what you want to migrate.

- Use the *Users* parameter if the DHCP administrators group includes local user accounts. Use the *Group* parameter to migrate all members of the DHCP administrators group that don't include local user, but does include domain users and groups.
- You can use the *IPConfig* parameter if you want to export the IP address configuration of the source server to the destination server.
- Use the *Path* parameter to specify the location of the *svrmig.mig* file. If you don't specify a network location, you'll need to manually transfer this file to a location that is accessible to the destination server.

Considerations for using the *Import-SmigServerSetting* cmdlet on the destination server include:

- Use the *FeatureID* DHCP to specify that the DHCP server role is what you are importing. While the *Import-SmigServerSetting* cmdlet will install the DHCP server role if it is not already present on the destination server, Microsoft recommends installing the service and then placing it in a stopped state before importing data using the WSMT.
- Use the *Users* parameter if the DHCP Administrators group on the destination server will include local users. Use the *Group* parameter if the DHCP administrators group contains domain users and groups.
- Use the *Path* parameter to specify the location of the *svrmig.mig* file. This file can be stored on an accessible network share or stored locally.
- Once migration is complete, start the DHCP server service using the *Start-Service DHCPServer* cmdlet.
- If you haven't done so already, you'll need to authorize the DHCP server in Active Directory before it will function.

Verification and post migration tasks

Once the migration has occurred, ensure that the DHCP server service on the source server remains in a stopped state, as well as on a client release, and then renew an IP address to verify that the DHCP server service on the destination server is functioning properly. You should also manually check that the scopes and settings that were present on the source server are also present on the destination server. Once you have verified that the migration has completed successfully, remove the DHCP server service from the source computer.

More Info

DHCP migration

Though the guide mentions Server 2012 R2, you can learn about migrating DHCP to Windows Server 2016 at https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/dn495425(v=ws.11).

FILE AND STORAGE SERVERS

You can use both the WSMT and Storage Migration Service to migrate file servers to Windows Server 2019. Storage Migration Service is a newer tool, and you should attempt migration using this tool first and fall back to WSMT only if you encounter a scenario that Storage Migration Service does not support.

Migrate file servers using Storage Migration Service

Storage Migration Service is a new service included in Windows Server 2019 that leverages Windows Admin Center to allow you to inventory the settings and data of a source file server, transfer the files and settings to a destination server, and then configure the destination server with the settings and identity that were present on the source server. The advantage of Storage Migration Service over WSMT is that it can:

- Inventory multiple servers and the data they store
- Quickly transfer files, file shares, and security settings from source servers
- Apply the identity of the source server to the destination server, including name and IP address settings so that users and applications do not have to alter their existing behavior and configuration

While copying files and their permissions has always been something that most administrators could generally figure out how to do fairly quickly, it was often a bit more challenging to copy the settings of the shared folders that hosted those files. Storage Migration Service allows you to migrate all the flags, settings, and security settings of SMB file shares. The flags that Storage Migration Service migrates include:

- Share State
- Availability Type
- Share Type
- Folder Enumeration Mode (also known as Access-Based

Enumeration or ABE)

- Caching Mode
- Leasing Mode
- SMB Instance
- CA Timeout
- Concurrent User Limit
- Continuously Available
- Description
- Encrypt Data
- Identity Remoting
- Infrastructure
- Name
- Path
- Scoped
- Scope Name
- Security Descriptor
- Shadow Copy
- Special
- Temporary

There are some limitations to what you can transfer using Storage Migration Service. These include not allowing you to migrate the following folders and files:

- Windows, Program Files, Program Files (x86), Program Data, Users, *$Recycle.bin*, Recycler, Recycled, System Volume Information, *$UpgDrv$*, *$SysReset*, *$Windows.~BT*, *$Windows.~LS*, Windows.old, boot, Recovery, Documents and

Settings, *pagefile.sys*, *hiberfil.sys*, *swapfile.sys*, *winpepge.sys*, *config.sys*, *bootsect.bak*, *bootmgr*, and *bootnxt*

- Any file or folder on the source server that conflicts with folders assigned the exclusion designation folders on the destination server

You can't use Storage Migration Service to migrate file servers that are members of separate domains. You also can't use Storage Migration Service to migrate file servers between workgroups.

Deploy Storage Migration Service

Storage Migration Service Orchestrator servers manage the migration process. It's important to note, especially if you are building your own test lab, that you can't deploy the Storage Migration Service Orchestrator from a host installed using the Evaluation version of Windows Server 2019. To install the Orchestrator service on a computer, Open Windows Admin center and click on Storage Migration Service as shown in Figure 22-16.

Technet24

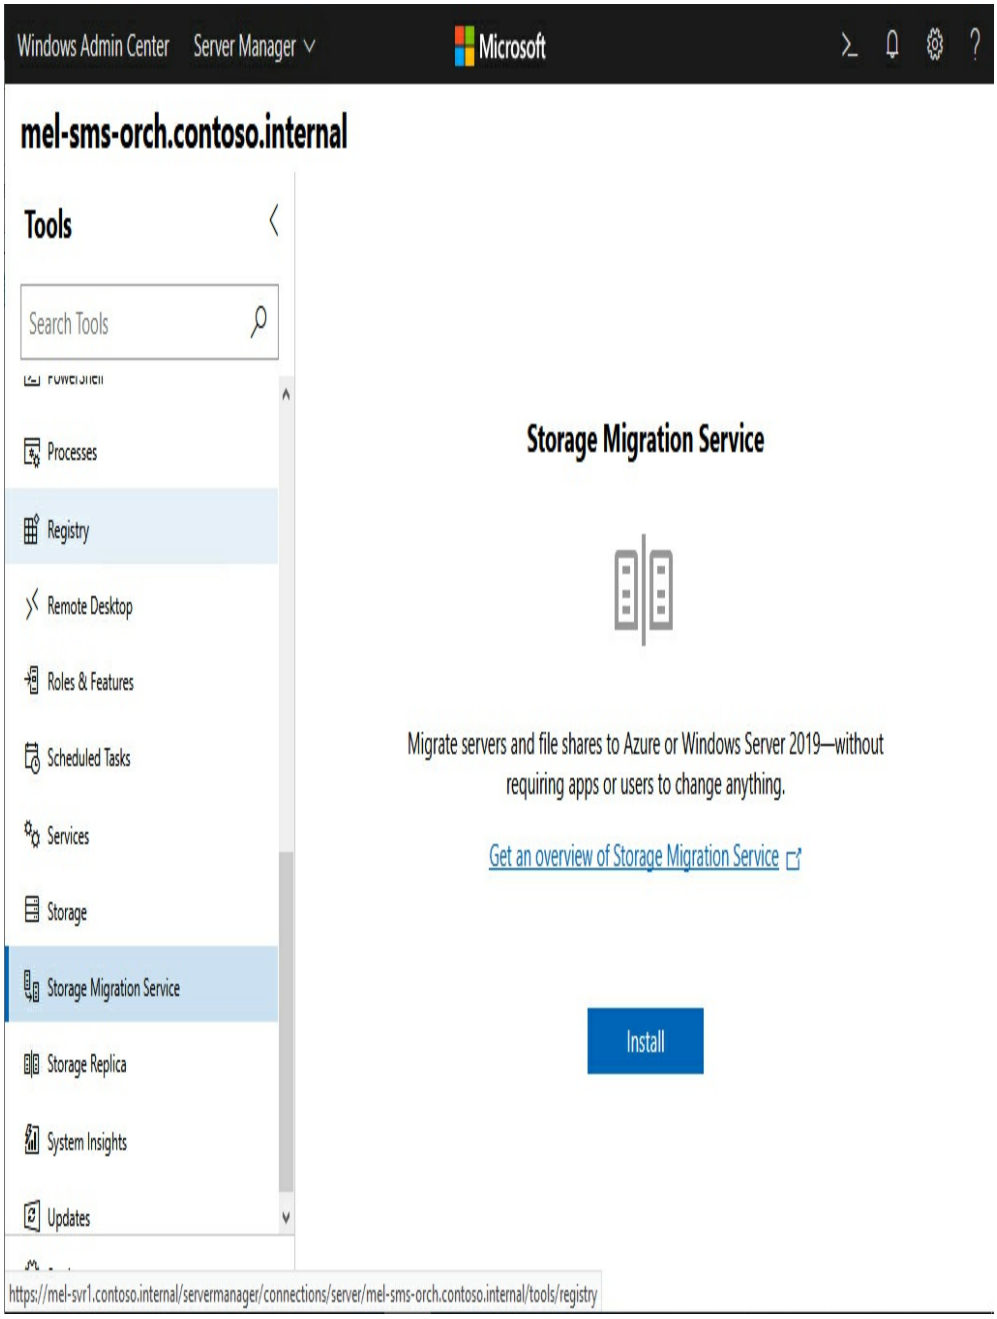


Figure 22-16 Install Storage Migration Service

Storage Migration Service requires that the following rules are enabled within Windows Firewall:

- File and Printer Sharing (SMB-In)
- Netlogon Service (NP-In)
- Windows Management Instrumentation (DCOM-In)
- Windows Management Instrumentation (WMI-In)

If you are using a third-party firewall product, ensure that you open TCP ports 135, 445, and 1025–65535. The service ports are TCP 28940 for the Orchestrator and TCP 28941 for the Proxy.

To perform a migration, you need access to the credentials of accounts with the following permissions:

- Administrator permission on the source and destination computer
- Administrator permissions on the computer that hosts the Storage Migration Service Orchestrator role

Inside OUT

Local users, local groups, Samba, and failover clusters

The version of Storage Migration Services that ships with Windows Server 2019 can only be used to migrate file shares from source servers that run Windows Server 2003 and later operating systems. If you are running Windows Server Semi Annual Channel 1903 or later, or if you have applied the cumulative update that was released in August 2019, you can also use Storage Migration Services to migrate file shares from source servers that use Samba versions 3.6.x and 4.x on RedHat 7.6, CentOS 7, Debian 8, Ubuntu 16.04, and SUSE Linux Enterprise Server (SLES) 11 SP4. Similarly, if you are running Windows Server Semi Annual Channel 1903 or later or have installed the cumulative update available in August 2019, the migration process will also allow you to migrate local users and groups from a source server to a destination server.

Inventory source servers

The first step in migrating file shares from one server to

another is to inventory the source servers to discover what shares are present, what files and folders are present, and all the identity information related to the source server. To do this, perform the following steps:

1. With Storage Migration Services selected in Windows Admin Center, click New Job. Provide a name for the job as shown in Figure 22-17 and then click OK.

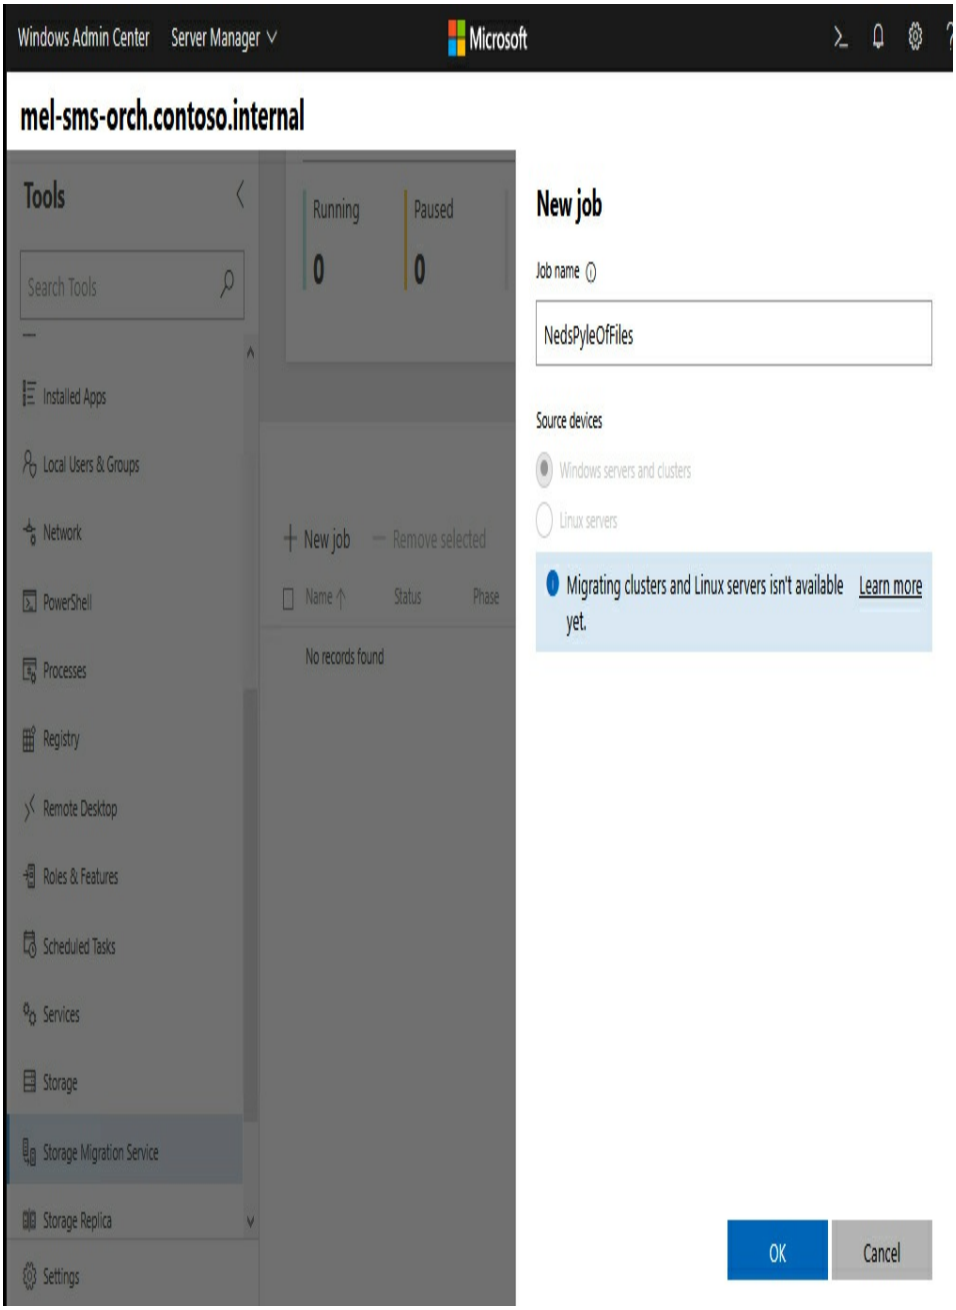


Figure 22-17 Create inventory job

2. On the next page of the Inventory wizard, provide the credentials of an account with administrative privileges. This page also allows you to choose whether to migrate the contents of administrative shares, as shown in Figure 22-18.

Technet24

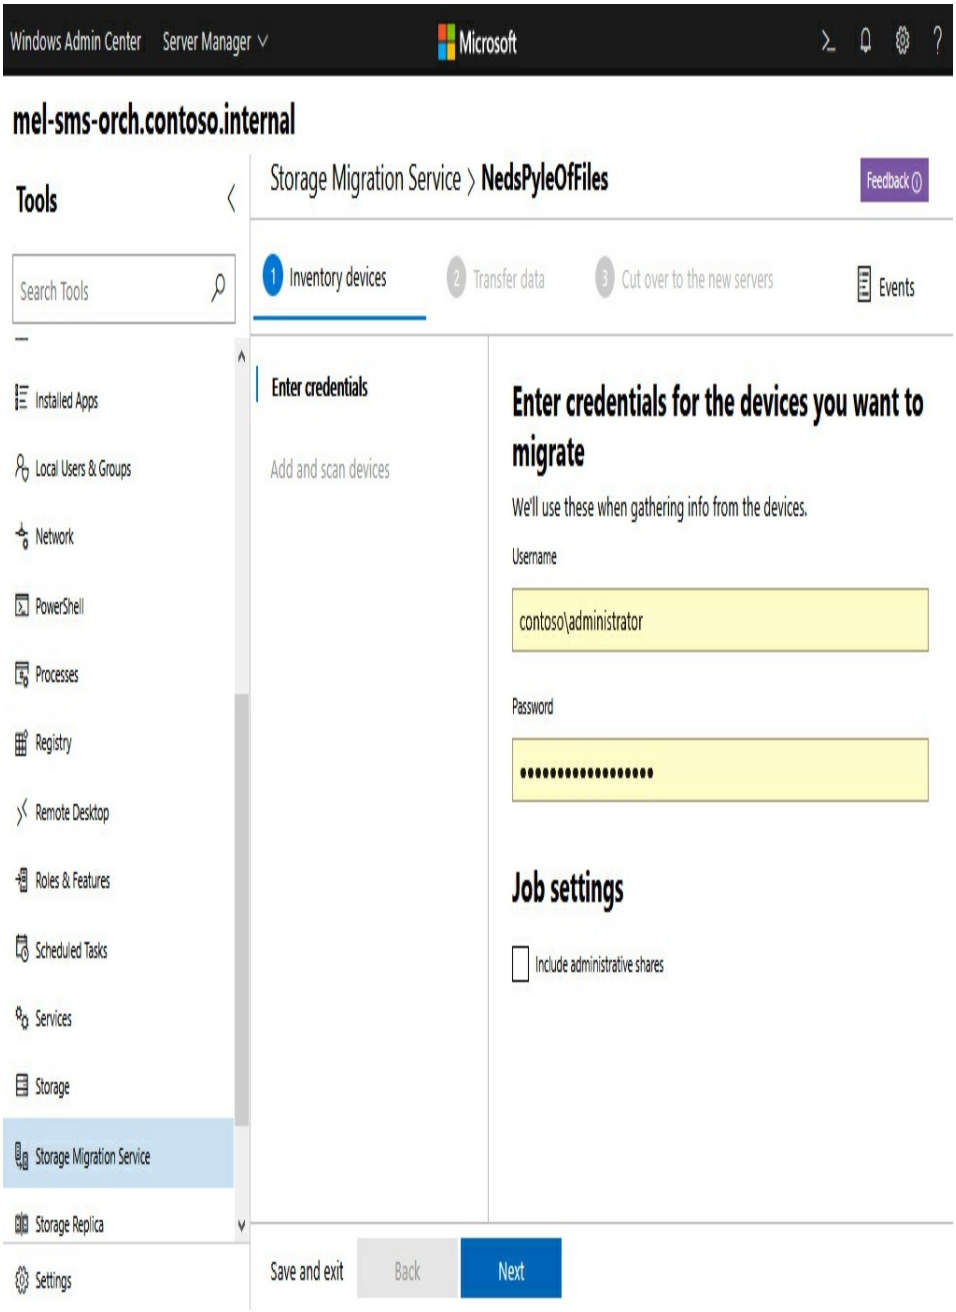


Figure 22-18 Inventory job credentials

3. Click Add to start the Add A Device wizard. In this wizard, provide the hostname for the server that hosts the files that you wish to migrate. You can use the Add button to add multiple source servers on this page.

4. When you are finished adding source servers, click Start Scan. Storage

Migration Services will generate an inventory of the listed servers.

5. After the inventory completes, you can click the server and view the server shares, configuration, network adapters, and volumes, as shown in Figure 22-19.

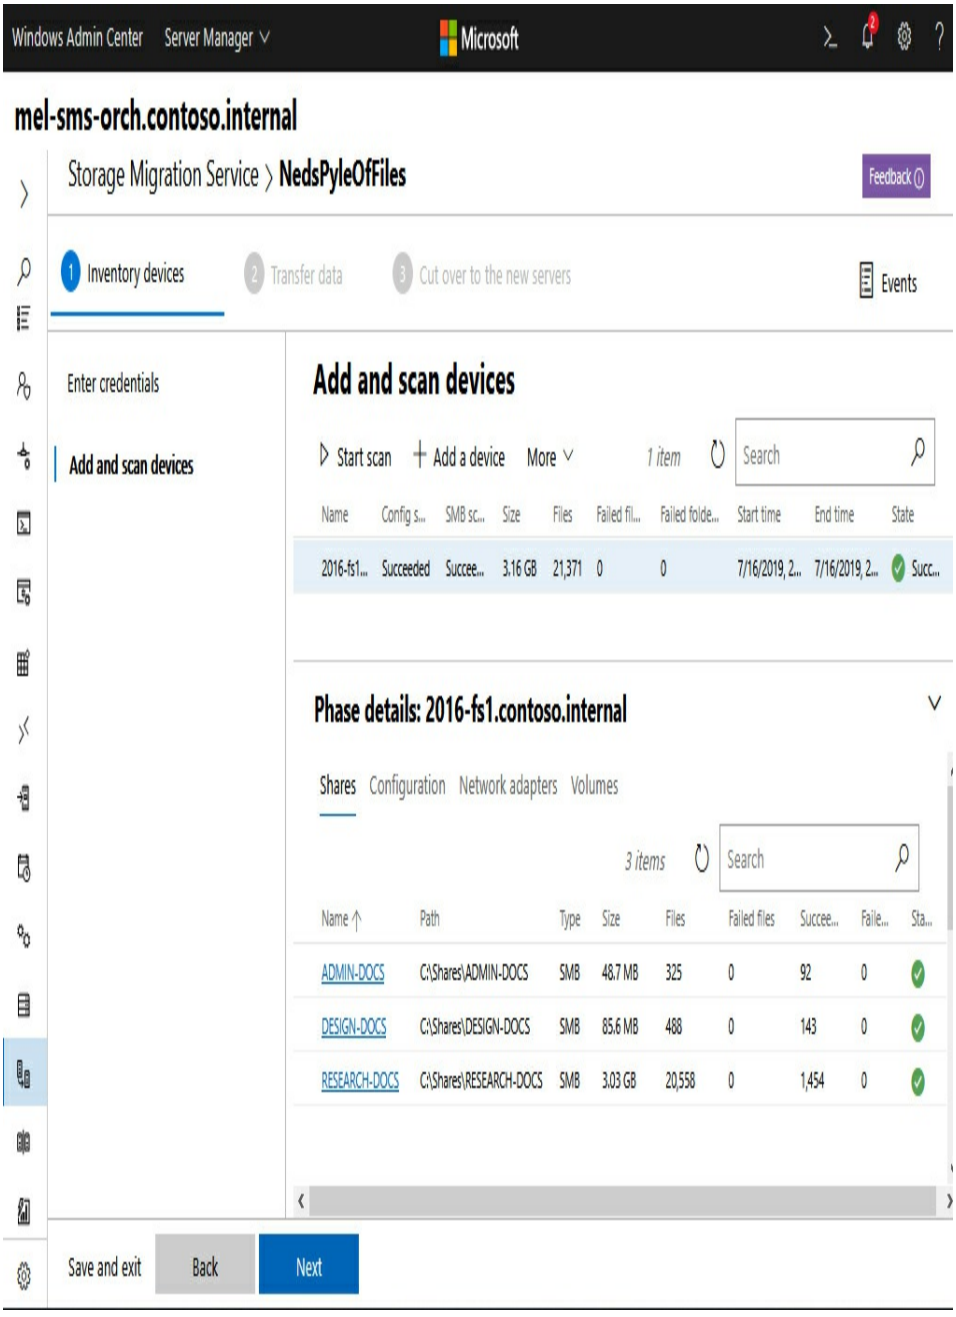


Figure 22-19 Inventory job results

Technet24

6. Once you have reviewed the information generated by the inventory scans, click Next to begin the data copy process.

Copy data

The data copy process copies data from the source to the destination server. Although this process is sometimes termed the *Transfer* process, it's important to note that the data is only copied, and it is not actually moved from one server to another. This means that if something goes wrong, you still have the original source server in its pre-migration condition.

To begin the copy process, perform the following steps:

1. Ensure that the firewall rules listed earlier are configured on the destination server.
2. Provide the credentials of an account that has administrative rights on the destination server.
3. Provide the address of the destination device, as shown in Figure 22-20. You can do this using the IP address or FQDN. Click Scan Device and the destination server will be inventoried. If the scan doesn't work, verify that the firewall rules are configured correctly.

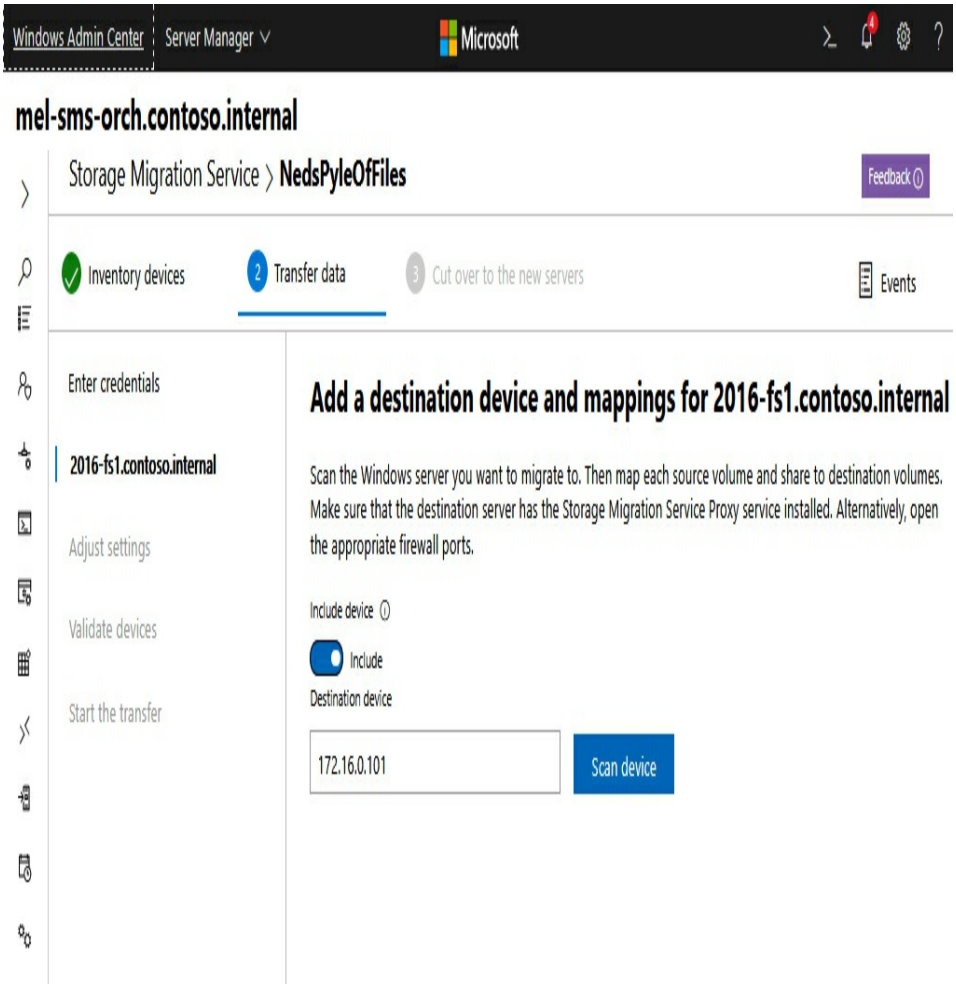


Figure 22-20 Add destination device

4. On the Adjust Transfer Settings page, select checksum settings for transmitted files, the maximum duration of the transfer job, how many retries to attempt for files that fail to transfer properly, and the delay between transfers.
5. On the Validate Source And Destination Devices page, click Validate and review any warnings.
6. On the Start The Transfer page, click Start The Transfer. You can monitor the progress of the transfer, as shown in Figure 22-21.

Technet24

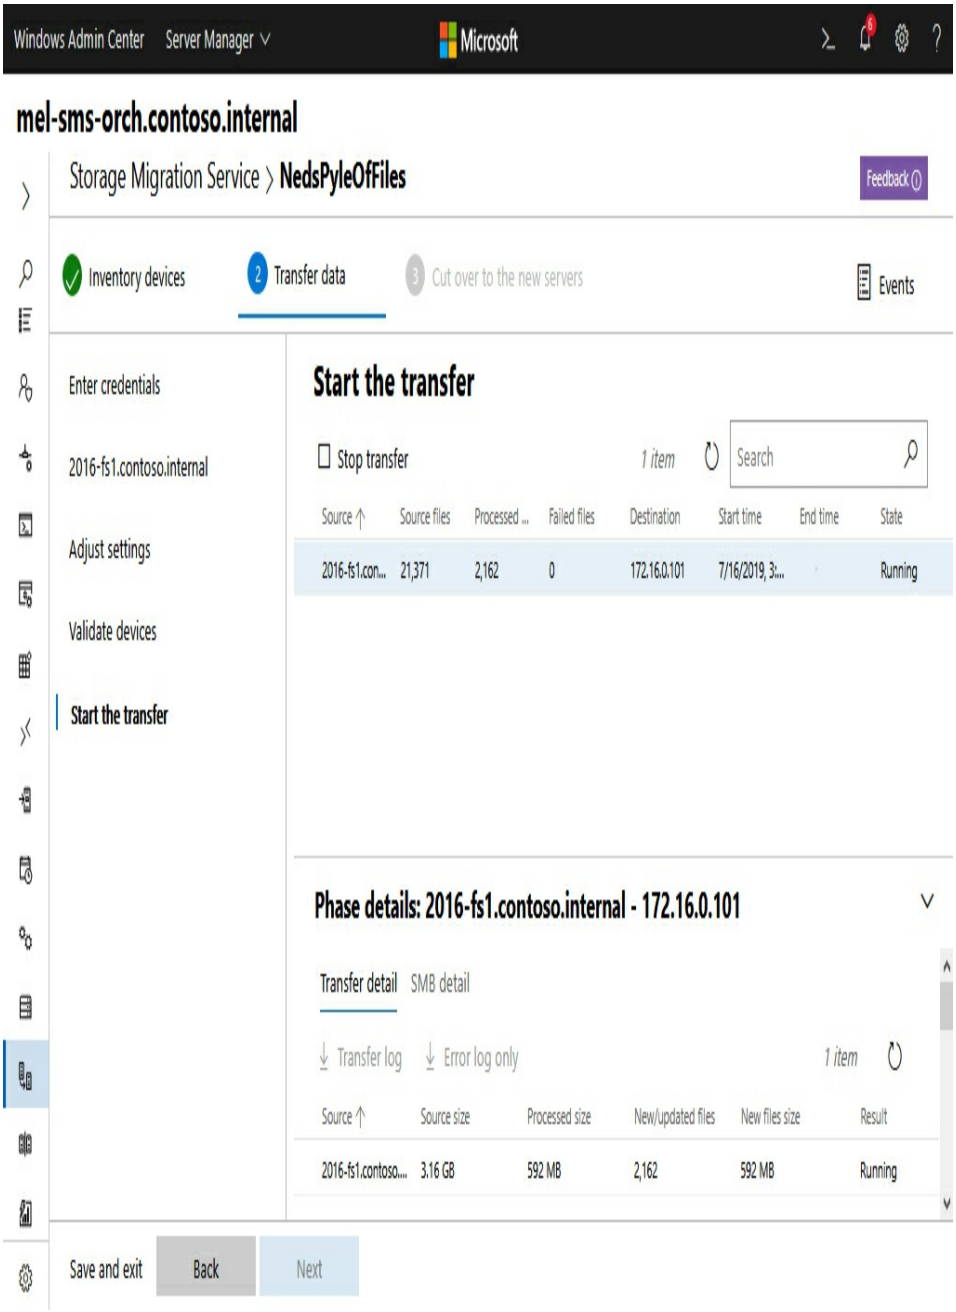


Figure 22-21 Monitor transfer

7. When the transfer completes, you can check the error log to determine which, if any, files and settings failed to transfer.

Cut over to new servers

The cut over phase is an optional phase that allows you to configure the destination server to assume the name and identity of the source server. This is one of the coolest parts of this technology because it means that you can have Storage Migration Service create a Windows Server 2019 clone of a Windows Server 2003 server, with the Windows Server 2019 server assuming the network identity of the server it is replacing; this means users and applications aren't aware that any change has occurred. To perform the cutover, perform the following steps:

1. On the Enter Credentials page, provide administrative credentials for the source and destination devices.
2. On the page with the source server name, specify how you want network identity information transferred from the source to the destination server. Figure 22-22 shows the address information assigned to the Ethernet adapter on the source computer being applied to the Ethernet adapter on the destination computer. You can also configure the source computer to be assigned a random name during the cutover process, or you can have it choose a new name to ensure that there are no conflicts once the destination server assumes the source server's identity.

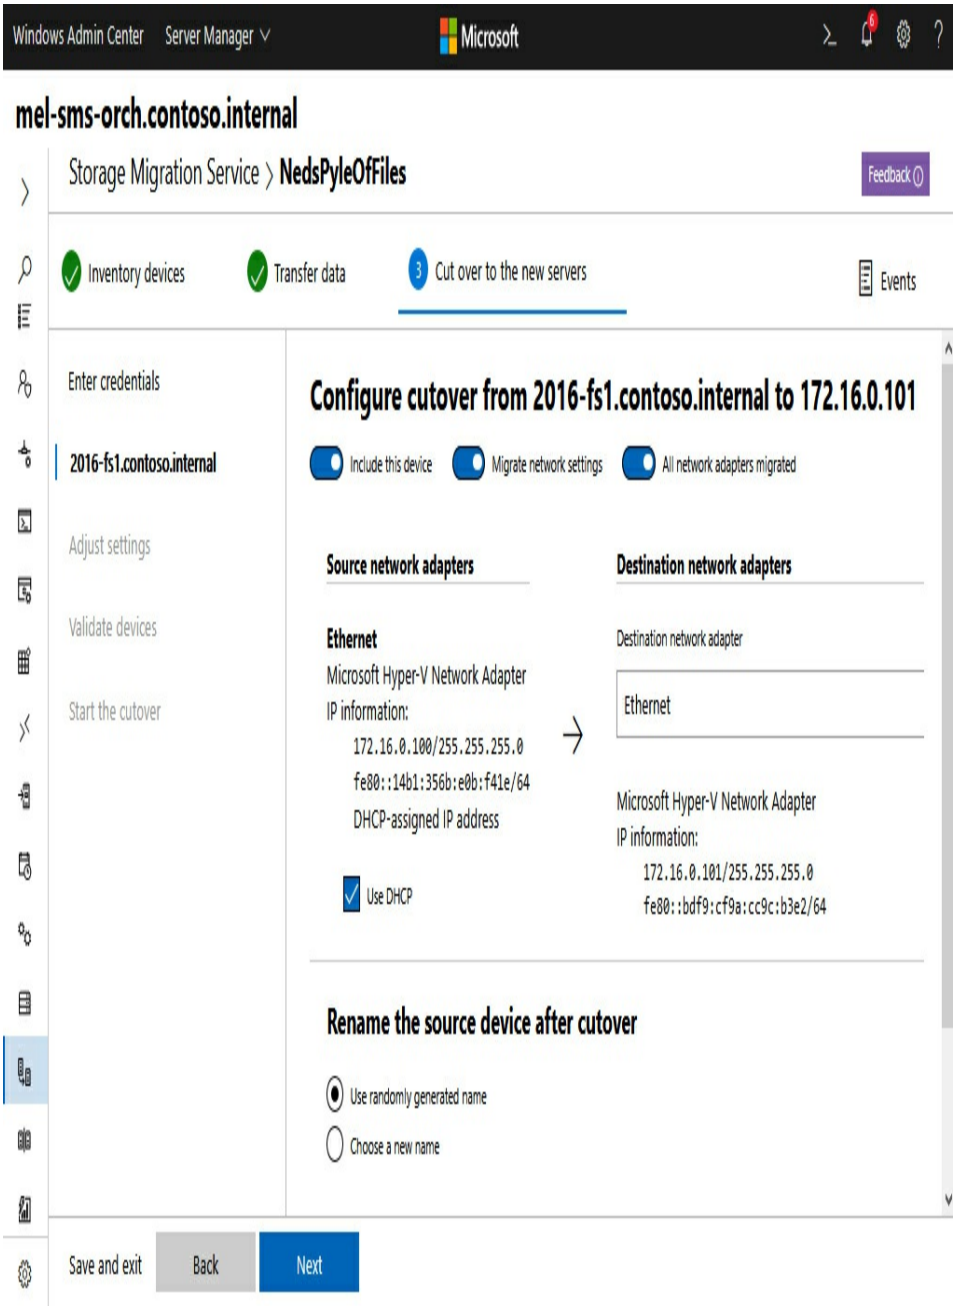


Figure 22-22 Configure cutover identity information

3. After validating settings and providing any credentials required to update the Active Directory information for the computers that are being managed, you can perform one final validation, as shown in Figure 22-23.

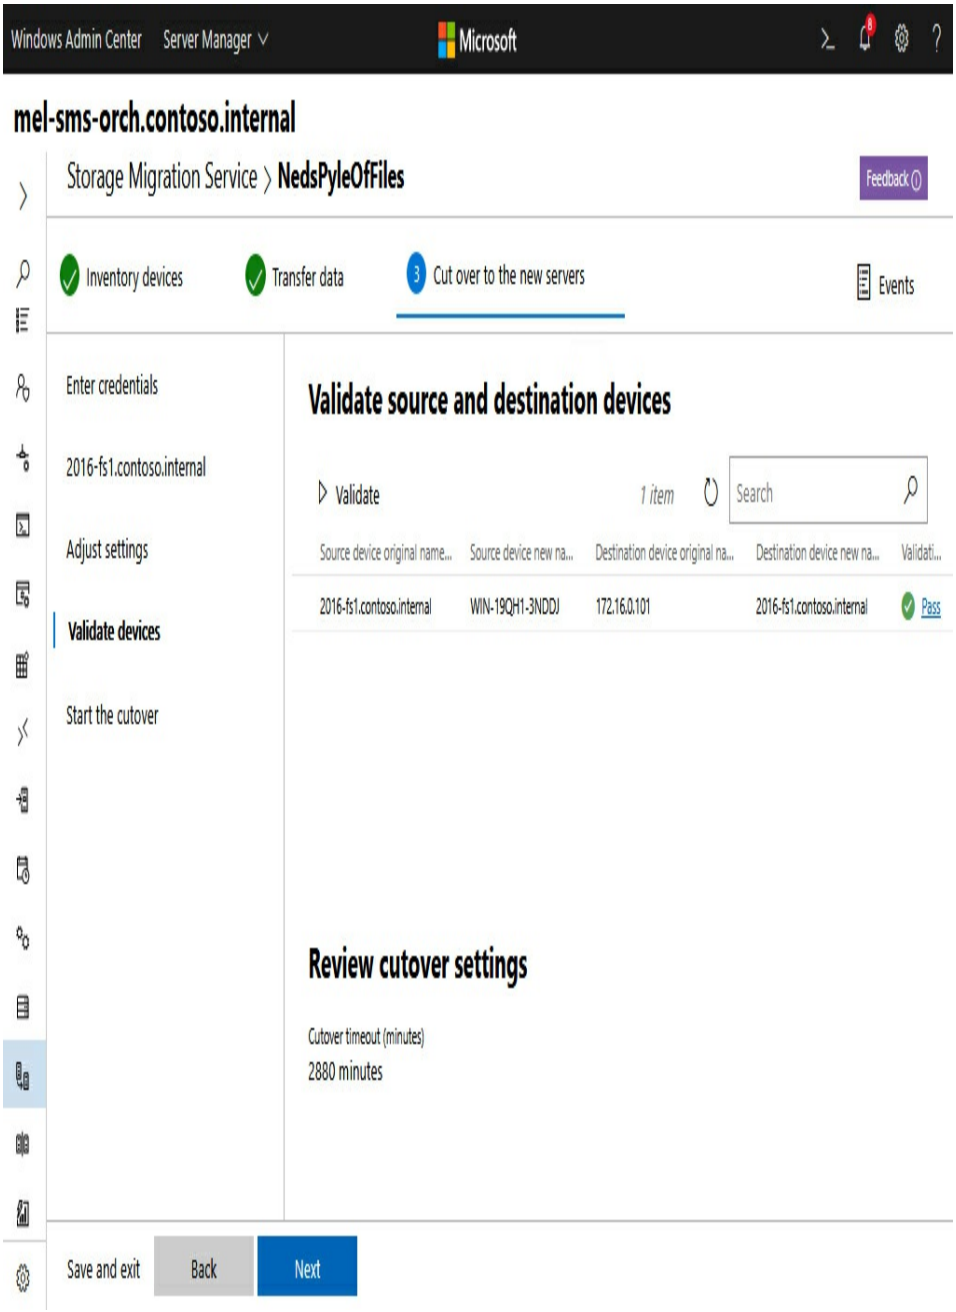


Figure 22-23 Validate cutover configuration

4. You then begin the cutover by clicking Start Cutover. When the cutover completes, click Finish.

More Info

Storage Migration Service

You can learn more about Storage Migration Service at https://docs.microsoft.com/en-us/windows-server/storage/storage-migration-service/overview.

Migrate file and storage servers using WSMT

While you should use Storage Migration Service to migrate file servers to Windows Server 2019, should you really want to, you can use the WSMT to migrate the following file and storage services information from a computer running Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019 to Windows Server 2019:

- Server identity information
- Local users and groups
- Data and shared folders
- Shadow copies of Shared Folders
- Data Deduplication
- DFS Namespaces
- DFS Replication
- File Server Resource Manager (FSRM)
- Group Policy settings related to SMB
- Group Policy settings for Offline Files
- iSCSI Software Target settings
- Network File System (NFS) shares

- Remote Volume Shadow Copy Service

The following migration scenarios are supported for file and storage services:

- The file server is a member of a domain.
- The file server is a member of a workgroup.
- File server data and shares are located on an external storage location, such as a SAN or on direct attached storage, as long as either storage medium preserves data permissions and file-share permissions.

It is important to note that file migration using the WSMT is aimed at moving data sets smaller than 100 GB, because only one file at a time is copied over a HTTPS connection. When moving larger data sets, Microsoft recommends using the *robocopy.exe* utility included with Windows Server 2019.

Migration permissions

Depending on the type of migration that you want to perform, a variety of permissions are necessary. If you are going to migrate data and shared folders, you must use an account that has local administrator permissions on the destination server.

Migration of DFS can be performed if an account is a member of a Domain Admins group. Alternatively, if there is more than one namespace server, the account used doesn't have to be a member of the Domain Admins group as long as it has permission to administer

all namespaces on the source server and is a member of the local Administrators group on the source and destination server.

Users who trigger migration of DFS replicas must either be members of the Domain Admins group or must have been delegated permission to the appropriate replication groups and replication members. If the user account is used to perform this action, it must be a member of the local Administrators group on the source and destination servers.

Preparing to migrate

To prepare to migrate the file and storage services role and data from a source server to a destination server, perform the following steps:

- Ensure that the source and destination server have their time correctly set. If the clocks on the source and destination server are out of synchronization, the migration will fail.
- Ensure that the same File Services role services are present on both the source server and the destination server.
- Ensure that Windows Server Migration Tools are deployed on both the source server and the destination server.
- Configure a firewall rule to allow traffic on TCP and UDP port 7000 between the source server and the destination server on both the source and destination server. TCP and UDP port 7000 is used by the *Send-SmigServerData* and *Receive-SmigServerData* cmdlets when establishing data migration connections. TCP and UDP ports 7001 and 7002 are also used if the migration occurs across subnets.
- Ensure that the destination path has sufficient disk space to migrate

data. If you have a configured NTFS quota or folder management, ensure that the quota limit will allow the migration of data.

- If migrating File Server Resource Manager data and settings, ensure that the destination server has the same volume configuration as the source server.
- While file and folder permissions are preserved during migration, remember that migrated files and folders inherit permission from parent folders. You should ensure that parent folders on the destination server have permissions that match those on the source server.
- If you are performing a DFS migration, ensure that the DFS Namespaces role service is installed and configured on the destination server. The DFS Namespace service must be running prior to commencing migration.
- If you are migrating DFS namespaces, back up the source server using a system state or full server backup. If the DFS namespaces are integrated into AD DS, it will be necessary to back up the system state on a domain controller to ensure that the AD DS configuration information for DFS namespaces is backed up.
- Ensure that the Server service on the destination computer is set to use the Automatic startup type as shown in Figure 22-24.

Technet24

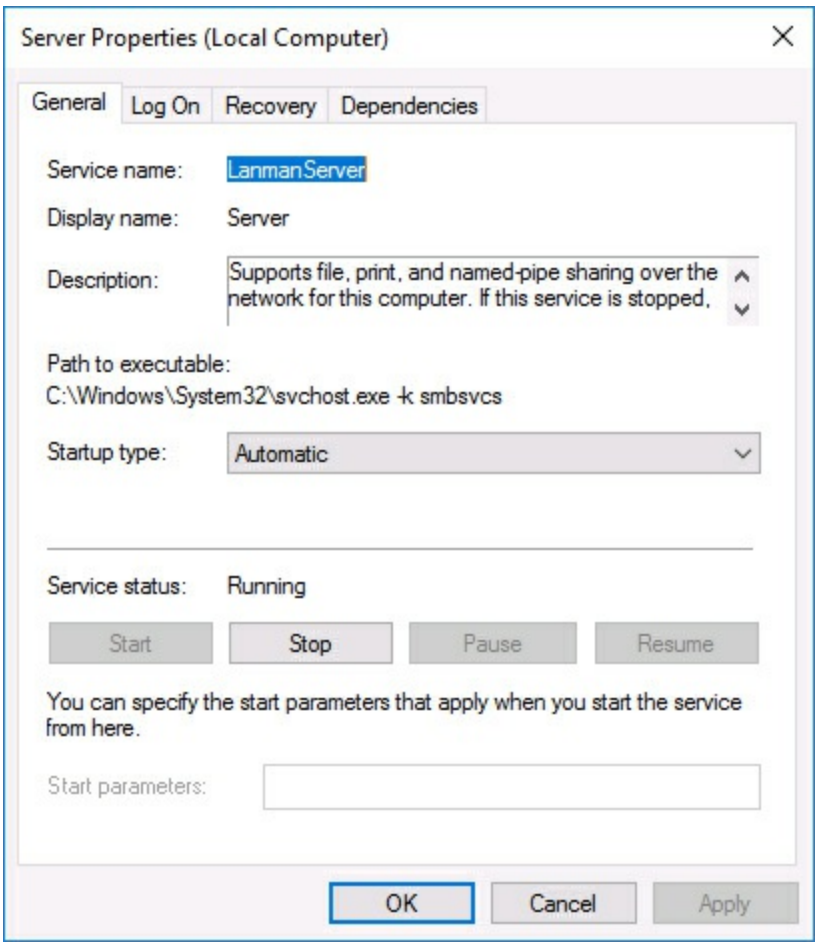


Figure 22-24 Server service settings

- You can back up the configuration information of a domain-based DFS namespace using the following command:

  Click here to view code image

```
dfsutil.exe root export <\\
<DomainName>\NameSpace> <BackupFileName>
```

Migrating File and Storage Services

Once you have completed the migration preparation steps, use the *Export-SmigServerSetting* cmdlet to

capture the following settings on the source server so that you can import them on the destination server using the *Import-SmigServerSetting* cmdlet:

- SMB (Server Message Block) settings
- Offline Files settings
- DFS Namespaces configuration
- File Server Resource Manager settings
- Shadow Copies of Shared Folders

Running the cmdlet will generate the migration file *svrmig.mig*, which you'll need to then transfer to and import on the destination server. If you want to migrate local users and groups from the source server to the destination server, use the User All and Group parameters with the *Export-SmigServerSetting* cmdlet. Remember to specify the same parameters with the *Import-SmigServerSetting* cmdlet on the destination server.

Once settings have been transferred from source server to destination server, use the *Send-SmigServerData* and the *Receive-SmigServerData* cmdlets to transfer data between the source and destination servers.

For example, to send data to the computer named MEL-FS-2K16, of files that are located under the *e:\fileshare* folder on the source computer to the *f:\fileshare* folder on the destination computer, you would run the command:

Click here to view code image

```
Send-SmigServerData -ComputerName MEL-FS-2K16 -SourcePath
e:\fileshare -DestinationPath
f:\fileshare -Recurse -Include All -Force
```

You must start the *Receive-SmigServerData* cmdlet in an elevated session on the destination server within five minutes of it being executed on the source server. Data transfer between the source and the destination server is encrypted.

More Info

File and storage services migration

Though the guide mentions Server 2012 R2, you can learn about migrating file and storage services to Windows Server 2016 at https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/dn479292(v=ws.11).

Chapter 23

Troubleshooting

Troubleshooting methodology

Command-line tools

Sysinternals tools

In this chapter, you'll learn about troubleshooting methodologies for complex problems. This chapter provides general guidelines for solving problems above the simple "Is it plugged into the network?" or "Did you switch it on and off again?" In this chapter, you'll also learn about some of the built-in command line as well as Sysinternals tools, which you can use to diagnose problems.

TROUBLESHOOTING METHODOLOGY

There are a couple of tricks to troubleshooting problems. The first is to realize that you'll be far more

Technet24

effective at troubleshooting if you approach the problem systematically, rather than just trying random solutions.

The next is that, jokes aside, there is far more to IT troubleshooting than typing some problem characteristics into (this being a Microsoft Press book) Bing and diligently applying any solution that turns up.

Some problems with the search engine approach include:

- Posts often provide an incomplete description of problem symptoms. This is often because the person asking for help in the post hasn't developed a full understanding of the nature of the problem that they are trying to solve by posting a request for help.
- Unless you understand all the symptoms, it's possible that the post and the proposed solution aren't relevant to the problem that you are troubleshooting.
- Even when other posters in the same thread respond that the solution worked for them, you can't be certain that they had the same problem as you or that they implemented the solutions proposed in responses to the post correctly.

This doesn't mean that you can't find the solution to a complex problem by consulting online forums; it's just that you should treat solutions offered to other people's problems with an appropriate degree of skepticism. Solutions posted on the Internet might work or they might not but finding out through experimentation on your production environment probably isn't the best way to display your technical aptitude.

Inside OUT

Thinking, fast and slow

Nobel prize winner Daniel Kahneman discusses how we arrive at conclusions in his bestselling book “Thinking, Fast and Slow.” This book is illuminating for people troubleshooting problems in information technology with its ideas around intuitive and reasoned decision-making. Kahneman suggests that our first response when tasked with solving a problem is to rely upon intuitive thinking. The problem with intuitive thinking is that it almost always involves jumping to conclusions with only minimal justification or evidence. As Kahneman explains, intuitive thinking is often no better than a random guess. Sometimes we justify these this guesswork by assuming that our own long prior experience in IT troubleshooting other problems must have given us superhuman intuition about the causes of technological problems. Agood troubleshooting strategy involves reducing the amount of guesswork and intuitive thinking and prefers deductive reasoning. Good troubleshooters know to distrust their first assumptions about a complex problem and to spend some time investigating the problem further before becoming certain as to the cause and the resolution.

Redeployment

When I started in IT, I’d often spend hours, if not days, troubleshooting a problem with a server or desktop computer because each computer had a unique configuration, which meant deploying from scratch would take substantially longer. Not only would it be necessary to reinstall the operating system, but it would be necessary to locate the correct drivers, something that was a lot more challenging in the early days of the web when many hardware vendors weren’t online, let alone providing websites with easy-to-locate driver downloads. With early versions of Windows, it was often necessary to keep boxes of disks around containing drivers to support each hardware configuration. The process of installing server

applications was often similarly fraught. Vendor documentation was scarce, and other than the occasional helpful newsgroup post, there wasn't much help available online when it came to troubleshooting. Getting a server up and running in the 90s could be such an involved process that you would take the time to troubleshoot whatever problem you had as much as possible and only redeploy the workload when all other avenues were exhausted.

With Windows Server 2019, most drivers utilized by the operating system are included or are easily accessible on vendor websites. It's relatively straightforward to back up and reinstall applications, and the web provides a cornucopia of documentation and advice on troubleshooting existing workloads.

If a workload with a known configuration has worked in the past and is still experiencing errors now, one of the options that you have at your disposal is to redeploy rather than spend time troubleshooting. With tools such as desired-state configuration, Puppet, or Chef, it's quite likely that you can redeploy in less time than it takes you to diagnose, research, and solve the problem you face.

The trick with all this is that you need to have workload configuration and deployment highly automated. Building out that configuration automation will take a

substantial amount of time, and you will need to ensure that the configuration automation is updated each time you update the configuration of the workload in question. The argument for this approach is easy to make if you have a large number of servers with similar configurations. In scenarios where you are only dealing with a small number of servers and where each workload is unique, you may end up spending more time automating configuration than you ever would deploying servers.

You'll also need to ensure that data is decoupled properly from the workload. A long-term aspirational goal of any server administrator should be to automate the configuration of any workload that they manage, so that they can redeploy it quickly in another VM or on another physical chassis should something catastrophic happen with the original workload. It's likely that you are far from that goal at the moment in your own environment, but as you get closer to it, you'll find that redeployment as a troubleshooting step becomes a viable first response strategy to complex problems. Admittedly, if it works, redeployment doesn't answer the tricky question of what might have caused the problem in the first place. However, if it does work, you can ponder the question of what actually happened as an existential problem, rather than as a practical one.

In some cases, redeploying the workload from what is

believed to be a good configuration will not solve the issue. In that scenario, you'll need to perform a deeper investigation as to what is causing the problem.

Symptoms and diagnosis

Before you can solve a problem, you need to understand the nature of the problem. The first step in understanding the nature of a complex problem is to have an understanding of the symptoms of that problem.

The easiest symptoms to document are the ones that caused you to notice that there was a problem in the first place. The thing that is important to realize is that there are likely to be other symptoms present that you might not have noticed. Understanding all these symptoms and their relationship with one another will allow you to diagnose the cause and resolution of the problem. It might be that the first thing you notice about a problem is only tangential to that problem. By investigating further, you discover that far more problematic things are occurring than the thing that drew your attention to the problem in the first place.

Inside OUT

More than one cause

Anyone who has diagnosed a complex problem knows that a single set of symptoms may have separate root causes. That's one reason why we might need to try multiple resolutions when troubleshooting. Generally, though, the more evidence you can collect about problem symptoms, the better idea you can have

about diagnosing the problem.

A diagnostic hypothesis is a hypothesis as to what might be causing the problem. It's what you suspect may be causing the problem before you have verification that what you suspect is causing the problem is actually causing the problem.

Rather than latching on to the first diagnostic hypothesis that occurs to you, try to list other things that might also be causing the problem. When thinking up other things, you're likely to gain more of an understanding about how you should be investigating the problem. You'll be looking at the problem more broadly, rather than focusing on the symptoms that first came to your attention.

You want to get as much evidence as possible about what's causing the problem before you try to solve it. As every experienced IT person knows, applying the wrong solution to a problem can often make the problem worse. By having as much information about the problem as possible, you reduce the chance that you'll apply an incorrect solution when attempting to solve the problem.

Dependencies

There's an old IT Pro joke that "it's always DNS." The foundation of this joke is that almost everything in IT is

dependent on DNS, so if an apparently unrelated system, workload, or service isn't functioning, the cause of that malfunction may actually be that DNS isn't working. So, even though it's a bit of a joke, before troubleshooting any complex problem, it probably doesn't hurt to fire up a console and run *nslookup*.

On a more serious note, systems, workloads, and services rarely stand alone, and they almost always depend on other things. As another part of your troubleshooting and diagnostic process, determine the dependency chain of the issue that you are troubleshooting. Working out the dependency chain will give you more items against which to test your hypothesis. It may also allow you to discover more symptoms of the problem than you were initially aware of.

For example, if a Windows Service has failed and won't start, determine what the service is dependent on. You can determine what Windows Service depends on, and what is dependent on that service, by looking on the Dependencies tab of the service's Properties, as shown in Figure 23-1.

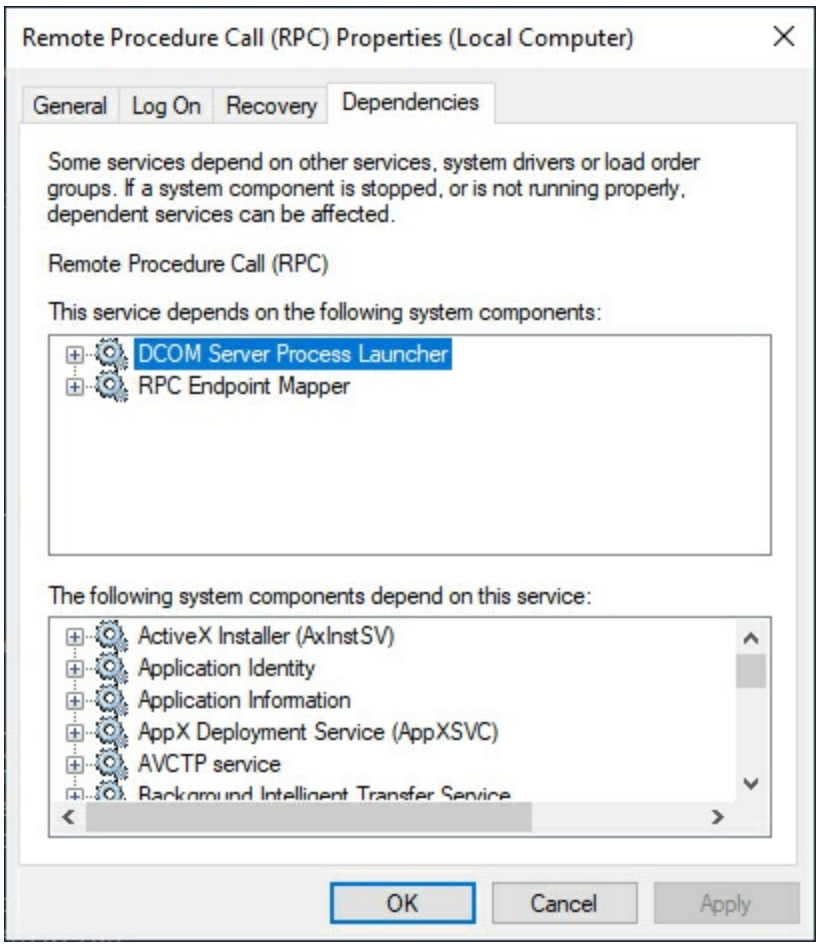


Figure 23-1 Service Dependencies

Ranking hypothetical solutions

A solution is hypothetical until it works. While simple problems often have obvious solutions, complex problems are likely to have a variety of apparent solutions, and you generally won't know which is the appropriate solution until you've applied it and found that it works. Before you apply a hypothetical solution in an attempt to solve a problem, you should rank the

Technet24

hypothetical solutions. Consider the following factors when ranking the order in which you should apply hypothetical solutions:

- Rank hypothetical solutions that are simple to implement and involve minimally changing the existing configuration higher than solutions that involve substantial changes to the existing configuration. This is because if the hypothetical solution turns out to be incorrect, you'll need to roll back the configuration changes before trying the next hypothetical solution. This is easier to do if the changes are less complicated than it is to do with a complex set of configuration changes. If a solution doesn't work, you need to go back to where you started rather than continuing to make further configuration changes when a solution you tried didn't work.
- Rank hypothetical solutions for which you have more evidence higher than solutions where there isn't as much evidence. As mentioned earlier, try to figure out what other evidence you should see if your hypothesis about what is causing the problem is correct rather than just focusing on the symptoms that brought the problem to your attention in the first place. Try to think of ways to disprove a diagnostic hypothesis without taking invasive action. It's better to discard an incorrect solution before implementing it because you find evidence that indicates that it is impossible than to try the incorrect solution in your production environment and find out that it doesn't work.
- Rank hypothetical low-impact solutions above solutions that are high impact. Always keep in the back of your mind that applying a solution could make things worse, so give preference to the solutions that will cause fewer problems if things go wrong than solutions that might require you to stay at work longer because a solution goes wrong.
- Rank hypothetical solutions that require a shorter amount of time to implement above solutions that take much longer to implement. When you don't know which solution will solve your problem, you

are better off trying the one that takes a couple of minutes over the one that takes a couple of hours.

Applying solutions

As mentioned earlier, it's important that when you apply a hypothetical solution and it doesn't resolve the problem, you find a way to revert to the configuration that existed prior to attempting the hypothetical solution. If you are diagnosing problems with a workload in a virtual machine, you should use snapshots or checkpoints to roll back to the "last known-bad" configuration, which is much better than building a "completely undocumented and worse" configuration.

Avoid trying a number of different hypothetical solutions that don't pan out and involve you making a large number of undocumented configuration changes to your production servers. When you continually modify the configuration of one server in an attempt to resolve one problem, you might end up causing a host of other problems.

So, before applying any solution, ensure that you have a rollback strategy that will allow you to return to the current configuration. Put another way, always ensure that you have an escape route for when your meticulously crafted plans go wrong.

The other thing you should do before applying any

solution is to have developed a set of criteria to determine what constitutes a resolution of the problem that you are trying to troubleshoot. You need to know what the problem resolution will look like. You should also test whether any of the characteristics of what you have decided constitutes the resolution are present before you apply the resolution. If you don't do this, it's entirely possible that part of your benchmark for declaring the problem solved may have been met before you attempted the solution. This can lead you to believe that you've solved a problem when the problem is still present. Alternatively, finding that one of your benchmarks for declaring the problem is solved prior to implementing a solution might provide further evidence that you don't fully understand all the characteristics of the problem.

COMMAND-LINE TOOLS

Windows Server includes a large number of command-line tools. While PowerShell is usually the first thing younger IT Pros think of when someone suggests doing work at the Windows Server command prompt, some of the older tools that are still included with the operating system can be useful when diagnosing and resolving problems. These commands can be a bit idiosyncratic, each having its own way of functioning; however, they often allow you to accomplish something with a single

command that might take a short script in PowerShell.

Rather than provide a list of each command (which you can find on the *docs.microsoft.com* website and is linked below), here are some select commands that you're likely to find useful when troubleshooting. Once you find one that sounds like it might be useful, you can dig in and see how it works, either by running it from the command prompt or by checking the online reference.

- **Arp.exe.** Manage the address resolution protocol cache.
- **At.exe.** Used to schedule tasks. Can also use *schtasks* for this purpose.
- **Attrib.exe.** Used to manage attributes assigned to files or directories.
- **Auditpol.exe.** Use this command to view information related to audit policies.
- **Bcdboot.exe.** Repair and manage the boot environment on the system partition.
- **Bcdedit.exe.** Manage the boot configuration data (BCD) stores.
- **Certreq.exe.** Manage certificate requests and interact with a certification authority.
- **Certutil.exe.** Can be used to configure, backup, and restore a certificate authority. Can also be used to manage local certificates.
- **Chkdsk.exe.** Troubleshoot file system for logical and physical errors.
- **Chkntfs.exe.** Manage automatic disk checking when system starts.

Technet24

- **Cipher.exe.** Manage encryption of files and directories on NTFS volumes.
- **Clip.exe.** Redirects command-line output to the Windows clipboard.
- **Comp.exe.** Compare two files.
- **Compact.exe.** Manage the compression of files and directories on NTFS volumes.
- **Cscript.exe.** Starts a script that executes in the command-line environment.
- **Date.exe.** Sets the system date.
- **Dcgpofix.exe.** Regenerates the default group policy objects for a domain.
- **Dfsrmig.exe.** Migrates SYSVOL replication from file replication services to distributed file system replication (DFSR).
- **Diskshadow.exe.** Manage volume shadow copy services.
- **Dnscmd.exe.** Manage DNS servers.
- **Driverquery.exe.** Lists all device drivers installed on the computer.
- **Eventcreate.exe.** Create custom events in event logs.
- **Find.exe.** Search text files or text output for a specific string.
- **Findstr.exe.** Search text files and text output using regular expressions.
- **Format.exe.** Format a volume.
- **Fsutil.exe.** Manage volumes and file systems. *Fsutil.exe* with the *file* parameter allows you to find files by properties, such as locating all files owned by a specific user. *Fsutil.exe* with the *repair* parameter allows you to manage NTFS self-healing operations without having to run *chkdsk.exe*. *Fsutil.exe* with the *tiering* parameter allows you to manage storage tier functionality.

- **Getmac.exe.** Returns MAC address and protocols associated with specific network adapters.
- **Gpfixup.exe.** Repair domain name dependencies in group policy objects and group policy links after a domain has been renamed.
- **Gpresult.exe.** Display cumulative group policy information for a specific user or computer.
- **Gpupdate.exe.** Refreshes group policy settings.
- **Hostname.exe.** View the computer's name.
- **Icacls.exe.** Manage discretionary access control lists on files.
- **Ipconfig.exe.** View TCP/IP configuration settings. Can also be used to refresh DHCP settings.
- **Klist.exe.** View a list of cached Kerberos tickets.
- **Ksetup.exe.** Manage the Kerberos protocol and key distribution center for Kerberos realms.
- **Ktmutil.exe.** Kernel transaction manager utility.
- **Ktpass.exe.** Configure the Service Principal Name for a host or service in Active Directory.
- **Lodctr.exe.** Configure performance counter registry settings.
- **Logman.exe.** Manage Event Trace Sessions and Performance logs.
- **Logoff.exe.** Kick a user off a Remote Desktop Session Host session and terminate their session.
- **Manage-bde.exe.** Manage BitLocker.
- **Mklink.exe.** Manage symbolic links.
- **Mountvol.exe.** Manage volume mount points.
- **Msiexec.exe.** Perform operations on Windows Installer files.
- **Nbtstat.exe.** Displays protocol information about NetBIOS over TCP/IP.

- **Netcfg.exe.** Manage the Windows Preinstallation Environment (WinPE).
- **Netsh.exe.** A command-line environment for managing network configuration.
- **Netstat.exe.** View information about TCP connections, networking statistics, and IP routing information.
- **Net print.** Manage print jobs and queues
- **Nlbmgr.exe.** Manage network load-balancing clusters.
- **Nslookup.exe.** View DNS information.
- **Openfiles.exe.** Display a list of open files and directories. Also can be used to disconnect a user from a remotely open file.
- **Pathping.exe.** Functions similarly to a combination of *traceroute.exe* and *ping.exe*.
- **Ping.exe.** Network diagnostic utility that uses ICMP to verify connectivity to another host.
- **Pnpunattend.exe.** Manage drivers installed on the system.
- **Pnputil.exe.** Manage the system's driver store.
- **Prndrvr.exe.** Manage printer drivers.
- **Prnjobs.exe.** Manage print jobs.
- **Prnmngr.exe.** Manage printers.
- **Prnqctl.exe.** Manage printer queues.
- **Pubprn.exe.** Manage publication of printers to Active Directory.
- **Recover.exe.** Recover readable information from a defective volume.
- **Reg.exe.** Manage registry entries.
- **Regsvr32.exe.** Register *dll* files with the registry.
- **Repair-bde.exe.** Recover data from a damaged Bitlocker

encrypted volume.

- **Schtasks.exe.** Manage scheduled tasks.
- **Secedit.exe.** Analyze security configuration against a security template.
- **Servermanagercmd.exe.** Menu driven command-line utility for server configuration.
- **Sfc.exe.** Verify the integrity of protected system files. Will replace problematic versions with updated versions.
- **Showmount.exe.** View mounted directories.
- **Subst.exe.** Allows you to associate a path with a volume letter.
- **Systeminfo.exe.** View system information at the command line.
- **Takeown.exe.** Recover access to a file by making an administrative account owner of the file.
- **Taskkill.exe.** Terminate a specific process from the command line.
- **Tasklist.exe.** View a list of running processes.
- **Time.exe.** Manage system time.
- **Tracerpt.exe.** Manage Event Trace Logs as well as log files created by Performance Monitor.
- **Tracert.exe.** Determine the path between the host and a remote system.
- **Tree.exe.** View a graphical representation within the command-line environment of the directory structure.
- **Tzutil.exe.** Manage the time zone settings.
- **Verifier.exe.** Manage driver verification.
- **Wbadmin.exe.** Manage backup.
- **Wdsutil.exe.** Manage Windows Deployment Services.
- **Wecutil.exe.** Manage event subscriptions.

- **Wevutil.exe.** Manage information about event logs and publishers.
- **Where.exe.** Locate files based on search keywords.
- **Whoami.exe.** Provides information about the currently signed-on user identity.
- **Winrs.exe.** Can be used to remotely manage and execute programs.
- **Wmic.** Manage WMI information from the command prompt.
- **Wscript.exe.** Windows Script Host is an environment through which scripts can be run.

More Info

Command-line tools

You can learn more about the Windows Server command-line tools as a whole at https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/windows-commands.

SYSINTERNALS TOOLS

The Sysinternals tools are a collection of utilities that you can use to diagnose and troubleshoot Windows Server 2019 as well as other Windows Server and client operating systems. The tools are hosted on Microsoft's website and can be used to interact not only with the GUI version of Windows Server 2019, but many of the tools can also be used with the Server Core version, either directly or remotely. In the following pages, you'll learn about several of the Sysinternals tools that you might find useful in troubleshooting and diagnosing

your Windows Server 2019 deployment.

> **More Info**
>
> **Sysinternals tools**
>
> You can learn more about the Sysinternals tools as a whole at https://docs.microsoft.com/en-us/sysinternals/.

Process Explorer

Process Explorer, shown in Figure 23-2, allows you to view real-time information about how processes are interacting with resources on a particular computer. Process Explorer can be thought of as a highly advanced version of the built in Task Manager tool. You can use Process Explorer to determine which processes have a specific file open. Process Explorer also allows you to determine which handles a specific process has open, as well as allowing you to see which DLLs and memory-mapped files have been loaded by a specific process.

Technet24

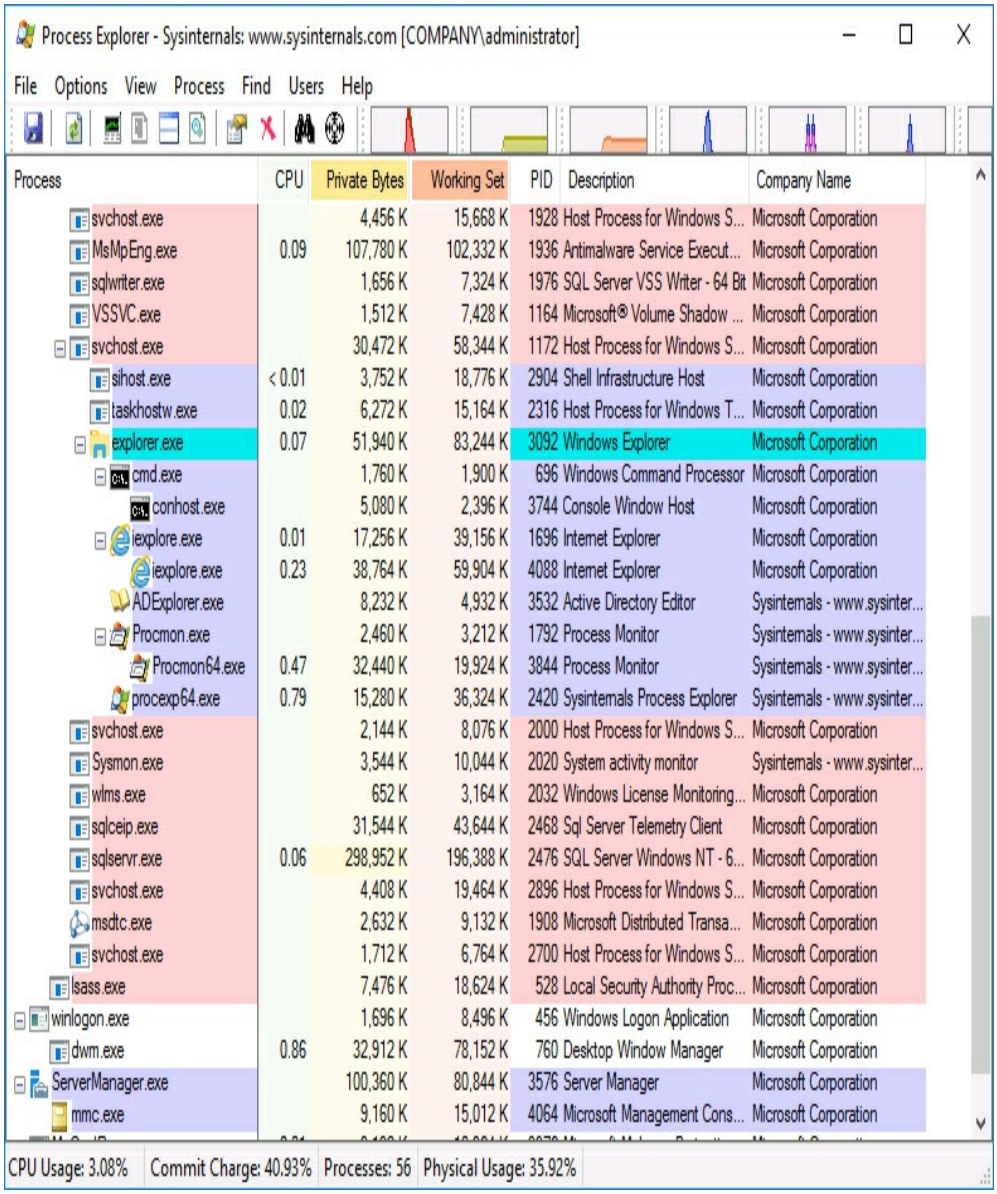


Figure 23-2 Process Explorer

Process Explorer includes the following functionality:

- **Tree view.** This allows you to see the parent–child relationships between processes.
- **Color coding.** Allows you to visually determine the process type and state, with services, .NET processes, suspended processes, and

processes running using the same credentials as the user running Process Explorer each being assigned a separate color code.

- **Fractional CPU reporting.** Shows where processes are consuming minimal amounts of CPU resources, rather than being rounded down to zero, which might imply that the processes were completely inactive.
- **Suspicious process identification.** Process Explorer will highlight processes that are flagged as suspicious by *VirusTotal.com*, an anti-malware service.
- **DLL and mapped file identification.** Identification of all DLLs and mapped files loaded by a process.
- **Kernel object handle identification.** Identification of all handles to kernel objects opened by a process.
- **Identifying open handles.** Ability to determine which processes have open handles to kernel objects, including files and folders.
- **Identifying processes loading specific DLLs.** Ability to determine which processes have loaded a specific DLL.
- **Viewing process threads.** View process threads, including start addresses and stacks.
- **Managing processes.** Modify a process's priority, suspend the process, and terminate a process or a process tree.
- **Process dumps.** Create a process dump.

More Info

Process Explorer

You can learn more about Process Explorer at https://docs.microsoft.com/en-us/sysinternals/downloads/process-explorer.

Process Monitor

Process Monitor is a real-time monitoring tool that you can use to view file system, registry, and process/thread activity. You can configure Process Monitor to log data to a file for later analysis, and Process Monitor logging supports scaling to tens of millions of captured events and gigabytes of log data. Process Monitor also includes the ability to capture thread stacks for specific operation, allowing you to determine the root cause of an operation. You can also capture the details of processes including image path, command line, related user or service account, and session ID. Process Monitor is shown in Figure 23-3.

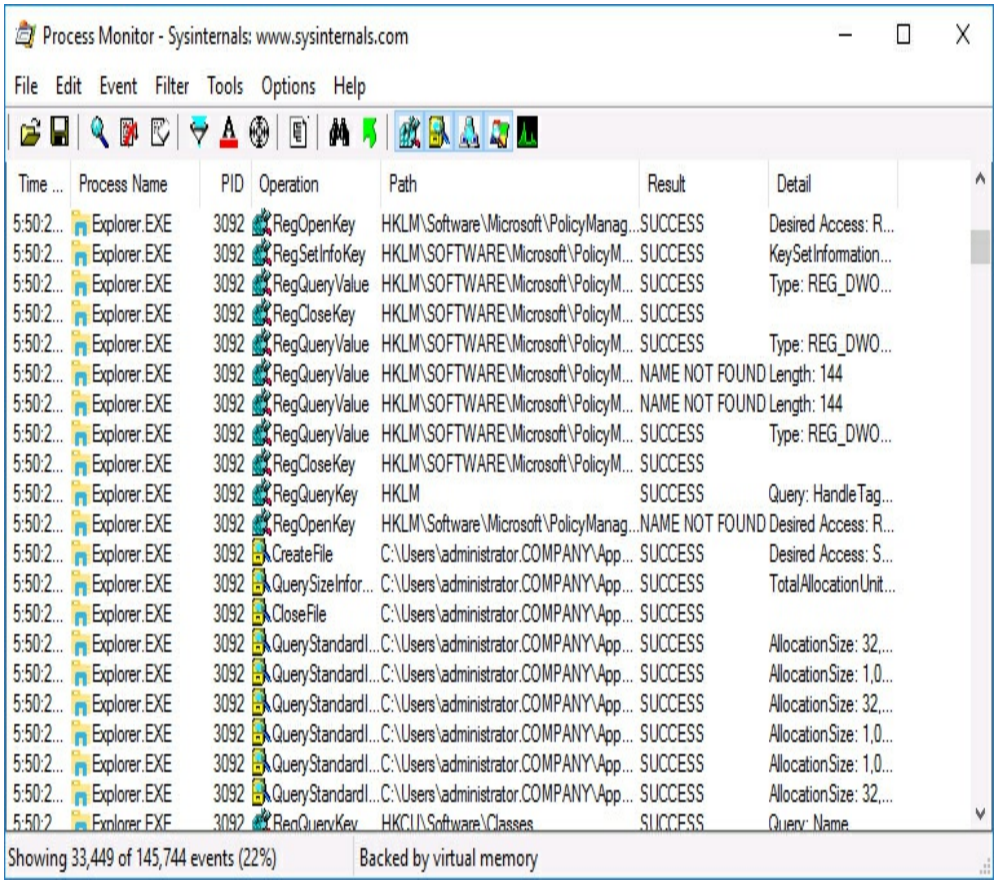


Figure 23-3 Process Monitor

More Info

Process Monitor

You can learn more about Process Monitor at https://docs.microsoft.com/en-us/sysinternals/downloads/procmon.

ProcDump

ProcDump is a command line utility that you can use to monitor a specific application for CPU utilization spikes. You can then generate diagnostic dumps during the spike that allow a developer to determine what causes the spike. ProcDump includes hung-window monitoring and unhandled exception monitoring, and it can be configured to trigger diagnostic dumps based on the values of specific system performance counters.

More Info

ProcDump

You can learn more about ProcDump at https://docs.microsoft.com/en-us/sysinternals/downloads/procdump.

PsTools

The PsTools Suite is a collection of tools named after the UNIX operating system PS utility rather than Microsoft's more recent PowerShell administrative scripting language. You can download the PSTools Suite in its entirety from Microsoft's website or download

Technet24

individual tools as you need them. Some tools in the PSTools Suite have functionality that is mirrored by Windows Server 2019 command-line utilities or PowerShell cmdlets. The PSTools Suite has been around since the Sysinternals Tools were first released in the 1990s. The command-line utilities and PowerShell cmdlets only became available more recently. Tools included in the PSTools Suite are as follows:

- **PsExec.** Enact processes remotely
- **PsFile.** View files that have been opened remotely
- **PsGetSid.** View the SID of a user or computer
- **PsInfo.** View system information
- **PsPing.** Covered in detail later in this chapter, this allows you to measure network performance
- **PsKill.** Terminate processes on the basis of name or process ID
- **PsList.** Determine detailed information about active processes on a system
- **PsLoggedOn.** Determine which accounts are signed on locally and through resource sharing
- **PsLogList.** Extract event log records
- **PsPasswd.** Change account passwords
- **PsService.** View and manage services
- **PsShutdown.** Shut down or reboot a computer
- **PsSuspend.** Suspends a running process

Note that some anti-malware tools will flag utilities in the PSTools Suite as malware. This is because some

malware toolchains include and leverage utilities in the PSTools Suite. If you are concerned that one of the utilities in the PSTools Suite is malware, use the SigCheck utility, covered later in this chapter, to verify the integrity of the tool in question.

> **More Info**
>
> **PsTools**
>
> You can learn more about the PsTools collection at https://docs.microsoft.com/en-us/sysinternals/downloads/pstools.

VMMap

VMMap allows you to analyze physical and virtual memory utilization. You can use VMMap to view a specific process's committed virtual memory types and the amount of physical memory (working set) that the operating system has allocated to the process. VMMap is useful when you need to troubleshoot application memory utilization. Figure 23-4 shows VMMap used to analyze the application ADExplorer.exe.

Technet24

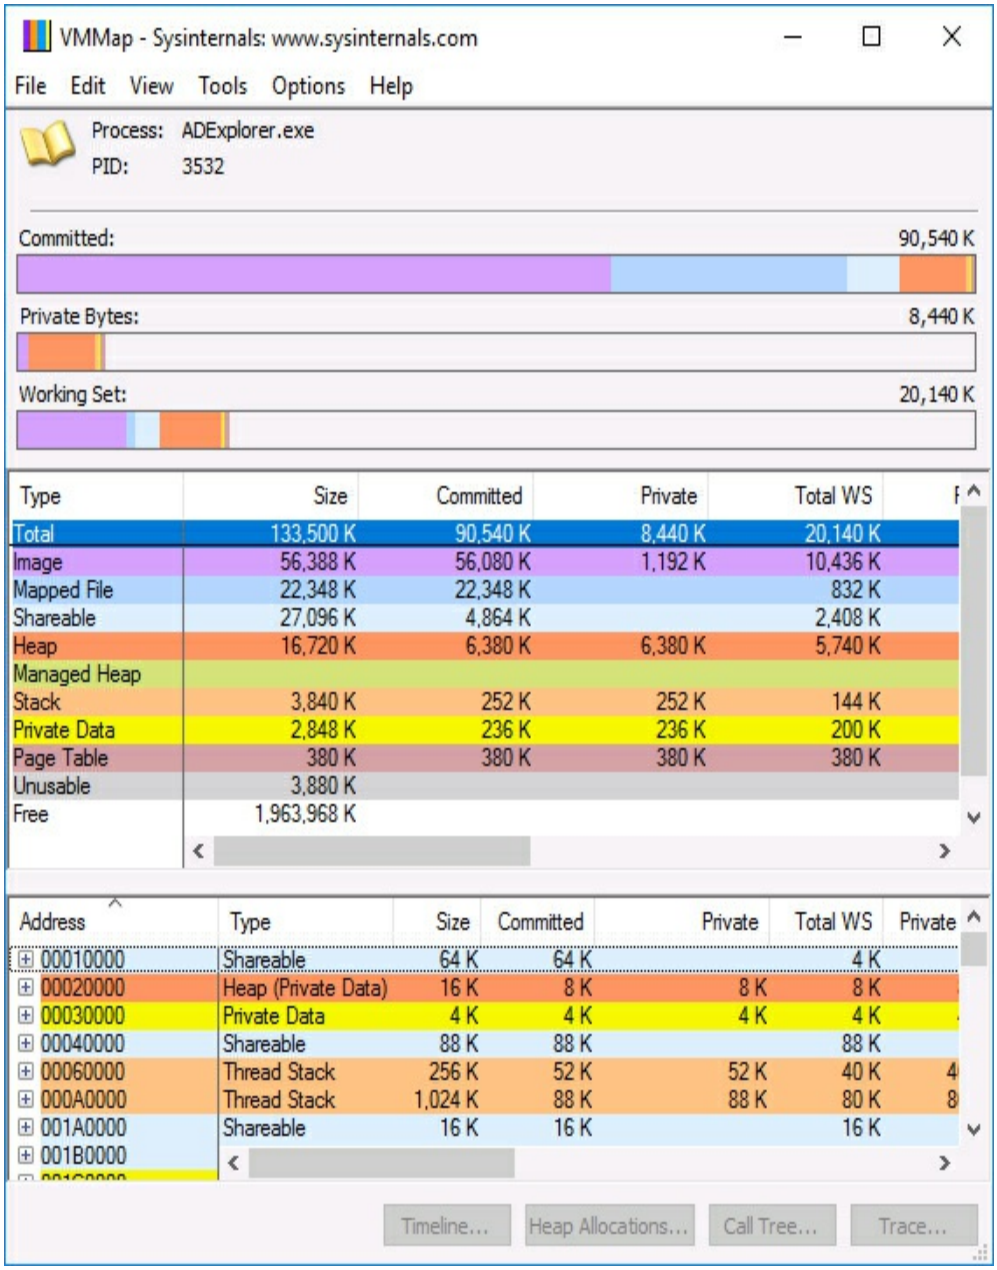


Figure 23-4 VMMap

More Info

VMMap

You can learn more about VMMap at https://docs.microsoft.com/en-us/sysinternals/downloads/vmmap.

SigCheck

Sigcheck is a command-line utility that allows you to view the following information:

- File version number
- Time stamp
- Digital signature details, including certificate chains

Sigcheck also allows you to check a file's status in VirusTotal, a database that scans against 40 separate antivirus engines. You can also use SigCheck to upload the file to be scanned by VirusTotal. Use SigCheck when you suspect a file may have been tampered with or may, for some other reason, not be legitimate. Figure 23-5 shows SigCheck used to check the signature of the file *c:\Windows\notepad.exe*.

Technet24

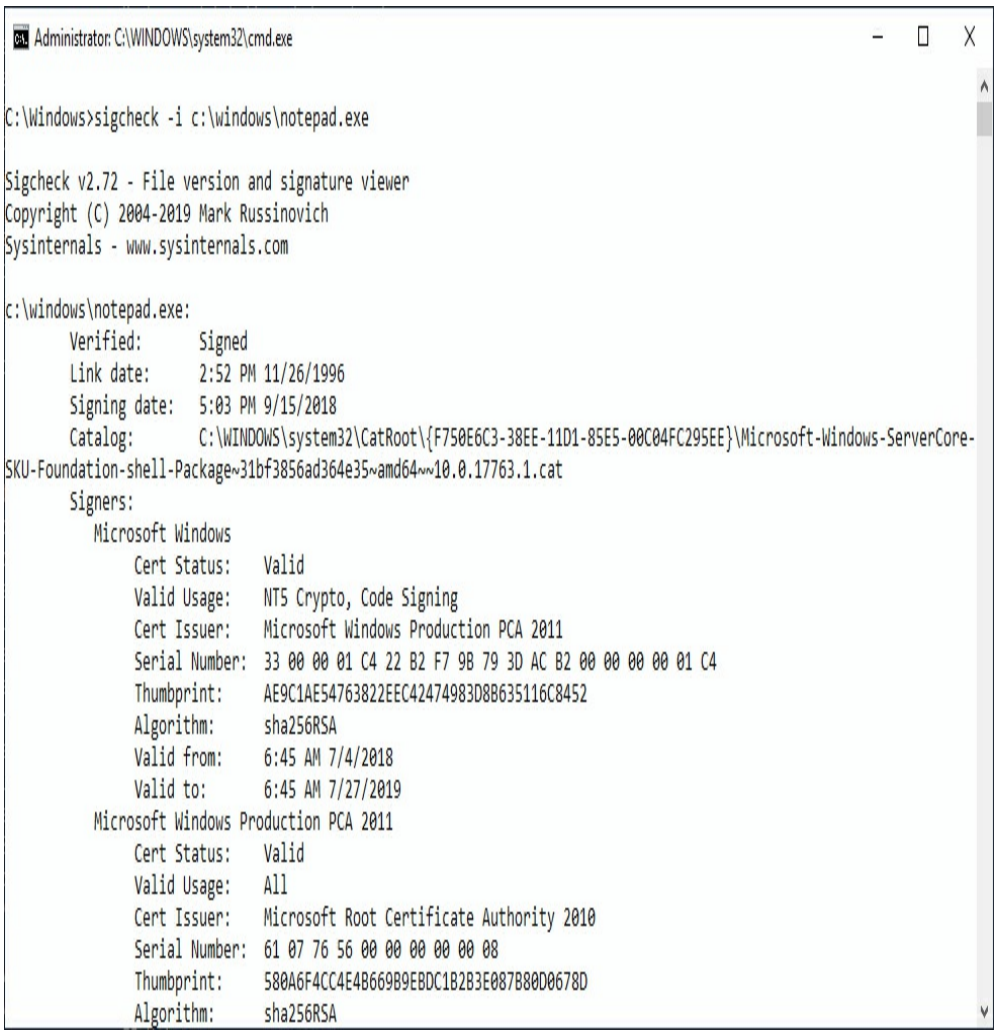


Figure 23-5 SigCheck

More Info

SigCheck

You can learn more about SigCheck at https://docs.microsoft.com/en-us/sysinternals/downloads/sigcheck.

AccessChk

AccessChk is a command line utility that allows you to

determine what access user, service, or group accounts have to specific files, folders, registry keys, or Windows services. Using the *Accesschk -?* command provides a list of all options that can be used, as shown in Figure 23-6.

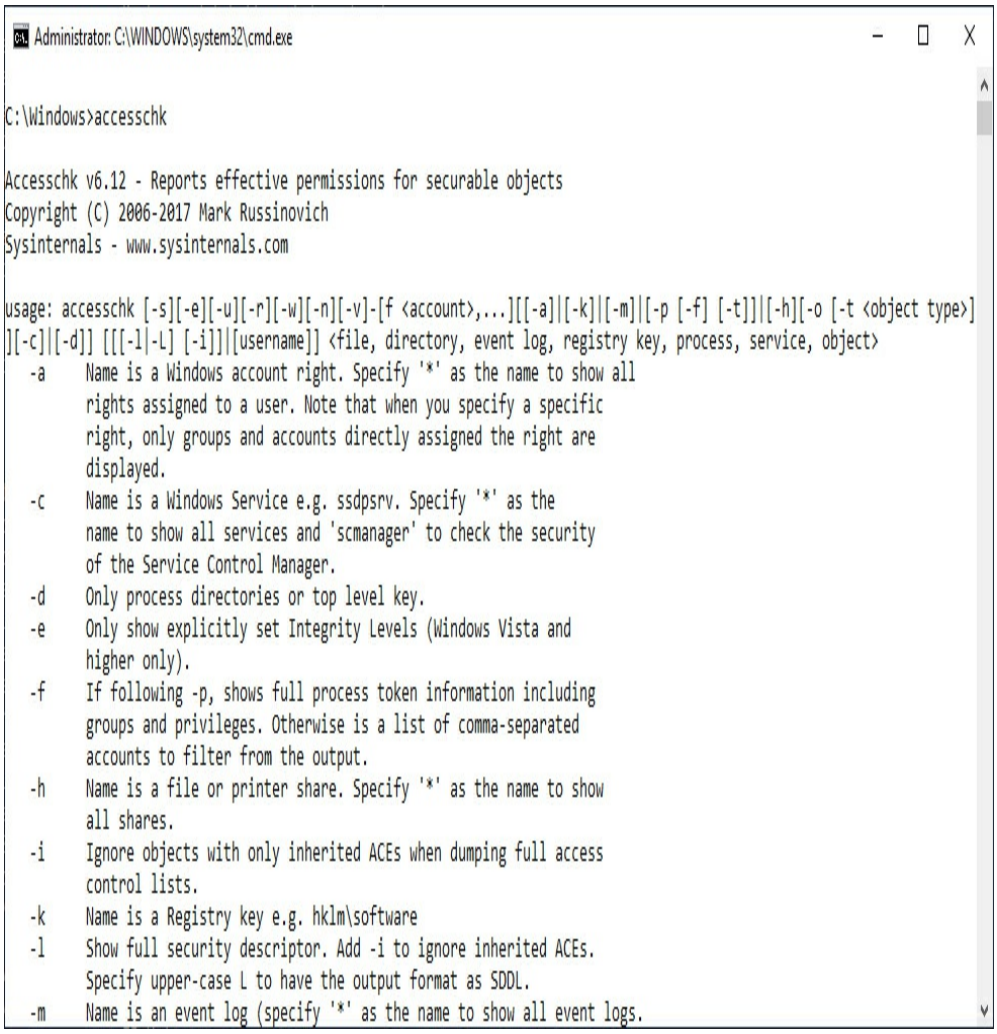


Figure 23-6 AccessChk options

More Info

AccessChk

You can learn more about AccessChk at https://docs.microsoft.com/en-us/sysinternals/downloads/accesschk.

Sysmon

System Monitor (Sysmon) is a service and device driver that you can install on Windows Server 2019 that monitors and logs system activity to the event logs. Sysmon will remain installed and running across reboots. Sysmon allows you to identify anomalous or malicious activity. It is used by security experts to determine how malware and malicious actors operate when compromising a computer. It is important to understand that while Sysmon monitors and logs potentially hostile activity, it provides no method for you to analyze that activity. A competent attacker will be able to determine that Sysmon is present on a computer and will take appropriate countermeasures.

Sysmon includes the following functionality:

- **Logs the creation of processes.** This includes information about the command line for the currently created process as well as any parent process.
- **Creates a hash of the process image file.** The default is SHA1, but you can configure Sysmon to use MD5, SHA256 or IMPHASH. You can have multiple hashes created for the same process image file.
- **Logs the loading of drivers and DLLS.** Records driver signatures and DLL hashes.
- **Includes a GUID for each process.** This allows you to track

activity when Windows Server reuses process IDs for different processes over time.

- **Includes a GUID related to each session.** This allows you to correlate events to specific logon sessions.
- **Log raw disk access.** Logs opens when a raw read access to disks and volumes occurs.
- **Log network connections.** Data written to the log includes connection source processes, IP addresses, port numbers, host names used, and port names.
- **Detects modifications made to file creation time.** Malware and intruders often modify file creation timestamps when cleaning up evidence of a breach.
- **Reloads configuration.** Reloads configuration if the configuration is modified in the registry.
- **Supports rule filtering.** This allows you to include or exclude specific events from being monitored and logged.
- **Allows you to generate events during the system startup process.** This allows you to capture activity caused by kernel-mode malware.

Sysmon events are stored in the *Applications and Services Logs\Microsoft\Windows\Sysmon\Operational* log folder, as shown in Figure 23-7.

Technet24

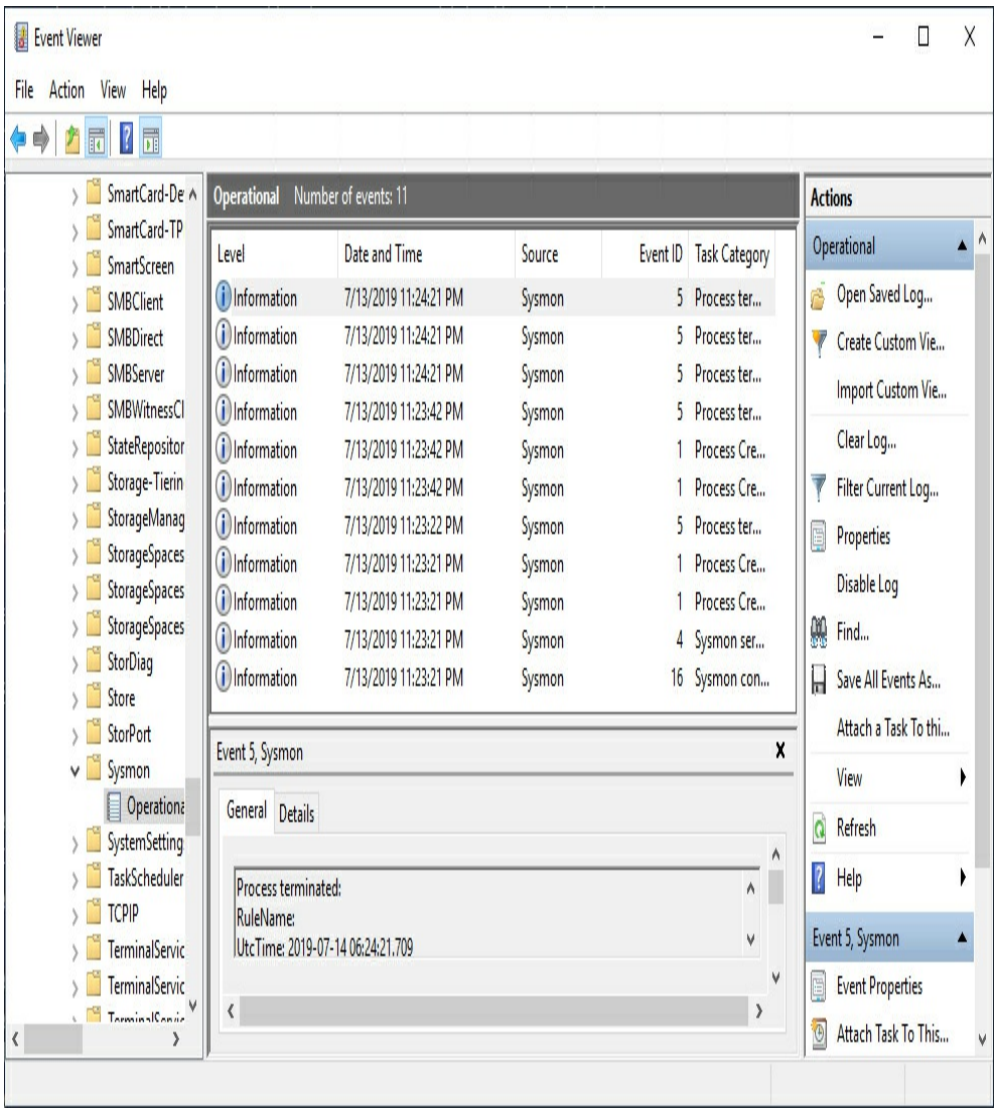


Figure 23-7 Sysmon events

Table 23-1 lists important event IDs generated by Sysmon.

Table 23-1 Important Sysmon Event IDs

Event ID	Description

1: Process creation	Includes detailed information about newly created processes. Includes the *ProcessGUID* field that uniquely identifies the process. Hash generated is a full hash of the associated file using the configured hashing algorithms.
2: A process changed a file creation time	An event generated when file creation time is modified by a process. You can even use this to track the accurate creation time of a file. Intruders often alter file creation times in an attempt to mask malware to make it look as though it was part of the original operating system installation. It is important to note that many standard processes alter file creation times, and it will be necessary to investigate further to determine if the change was part of standard operating system behavior or malicious activity.
3: Network co	Logs TCP and UDP connections to the server. This option is disabled by default when you install Sysmon. When enabled, each connection is associated with a *ProcessID* and *ProcessGUID*. Event will record source, destination host names, IP addresses, and port numbers.

Technet24

nnection	
4: Sysmon service state changed	Records when the Sysmon service starts or stops.
5: Processterminated	Records when a process terminates.
6: Driverlo	Records when a driver has loaded. Provides hash and signature information for the driver. Signatures will be created asynchronously to minimize performance impact.

added	
7: ImageLoaded	Records when a process loads a specific module. It is disabled by default when Sysmon is installed and must be manually enabled. It generates hash and signature information asynchronously to minimize performance impact. Only enable with appropriate filtering, because enabling this process will cause a large number of events to be written to the monitoring log.
8: CreateRemoteThread	Records when one process creates a thread in another process. This is one strategy used by malware when it injects code to hide it in the execution of another process. Event records both the source and target process. Provides information on the code that runs in the new thread, including *StartAddress*, *StartModule*, and *StartFunction*.
9: RawAccessRead	Records when a process performs read operations from storage using the \\.\ denotation. This technique is tracked because it is often used by malware when exfiltrating files locked for reading. It also avoids most file access auditing technologies. The event records the source process and target device.
10	Records when one process opens another process, which may

Technet24

: ProcessAccess	involve queries reading or writing data in the address space of the target process. This event ID often records hacking tools that copy the contents of memory used by processes such as the Local Security Authority (Lsass.exe) as part of a Pass-the-Hash attack. Enabling this type of auditing will generate a substantial number of events if diagnostic utilities that query the state of processes are being used on the server. In that case, create filters to remove this expected access.
11: FileCreate	Records when a file is created or overwritten. This is useful when investigating the status of autostart locations as well as temporary and download locations. These locations are often used by attackers when storing malware during initial server compromise events.
12: RegistryEvent (Object create and delete)	Registry modification events, where keys and values are created or deleted, are mapped to this event. Use this event to determine if malware has made modifications to the registry.

13: RegistryEvent (ValueSet)	Registry modification events, where values are changed. The changes are written as event data.
14: RegistryEvent (KeyandValuerename)	Registry modification events, where keys and values are renamed. Name changes are written as event data.
15	Records when a named file stream is created. Creates a hash of

: FileCreateStreamHash	the contents of the file. Useful to track malware variants that use browser downloads to drop executables or configuration settings. This problem is less likely to occur if you deploy Server Core, as the web browser is not available.
255: Error	Generated when Sysmon experiences an error. May occur if Sysmon is under extreme load.

More Info

Sysmon

You can learn more about Sysmon at https://docs.microsoft.com/en-us/sysinternals/downloads/sysmon.

AccessEnum

AccessEnum allows you to determine where a file, directory, or registry key has separate permissions from its parent object. This allows you to view where permissions have been configured at the object level rather than where the object has inherited the permission from the parent level. This sort of

configuration is almost always done manually, although it won't be obvious unless surfaced by a tool such as AccessEnum. Figure 23-8 shows where one executable in a directory has different permissions applied to that of the parent directory.

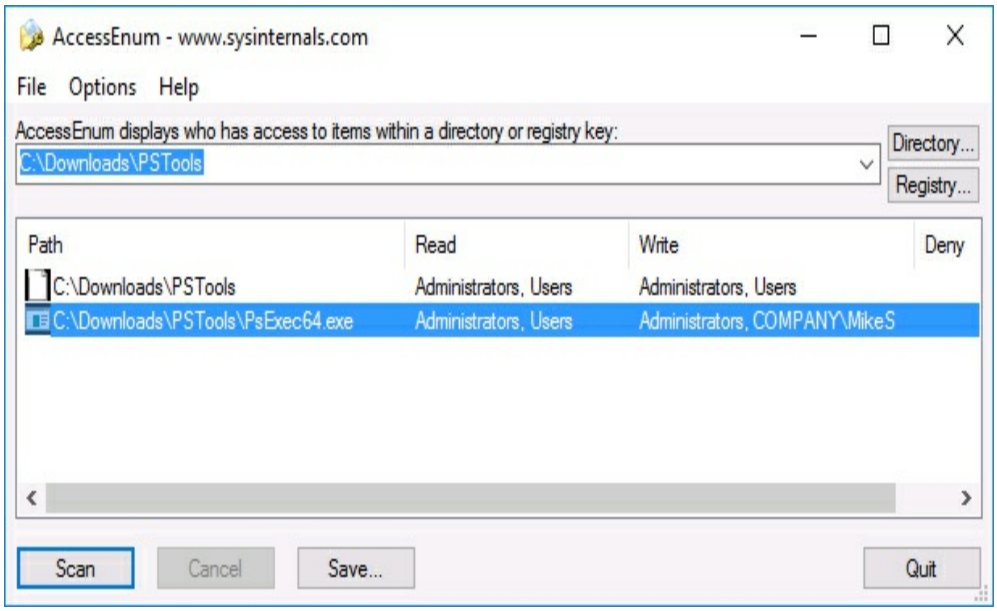


Figure 23-8 AccessEnum

More Info

AccessEnum

You can learn more about AccessEnum at https://docs.microsoft.com/en-us/sysinternals/downloads/accessenum.

ShellRunAs

ShellRunas provides you with the ability to quickly launch applications, including cmd.exe and PowerShell sessions, using the credentials of other users by right-

Technet24

clicking and selecting Run As Different User from the context menu. You must run *shellrunas /reg* from an elevated command prompt to register ShellRunas with the computer. Figure 23-9 shows *ADExplorer.exe* being configured to run with alternate credentials.

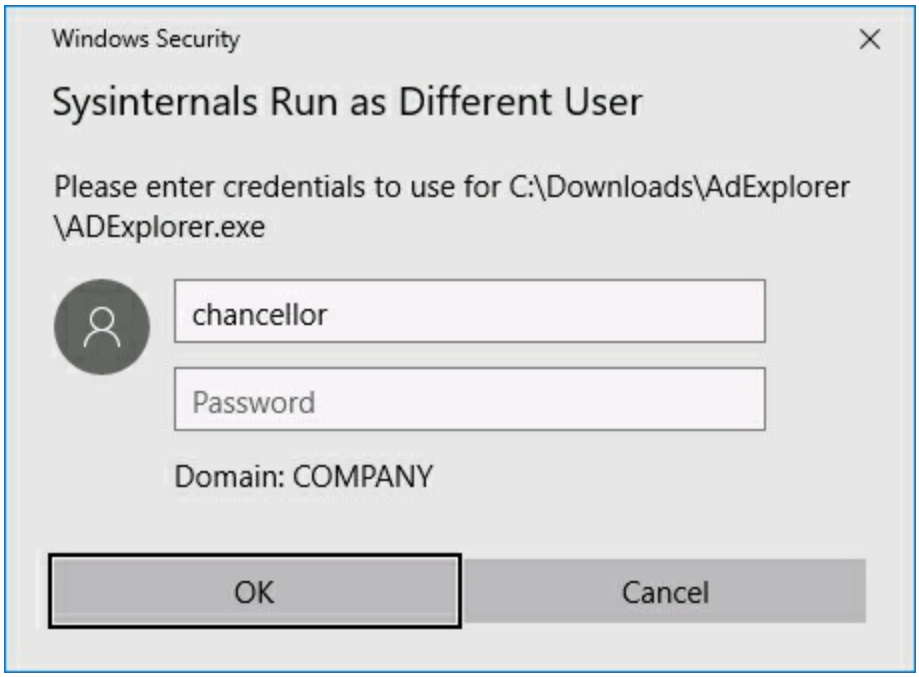


Figure 23-9 ShellRunAs

More Info

ShellRunAs

You can learn more about ShellRunAs at https://docs.microsoft.com/en-us/sysinternals/downloads/shellrunas.

LogonSessions

LogonSessions is a command-line utility that allows you to view the activity of all active sessions on a Windows

Server 2019 system. When used with the *-p* option, as shown in Figure 23-10, you're also able to see which processes are running in each session. This tool can also be used when connected to a remote PowerShell session to determine which active sessions are present on a server.

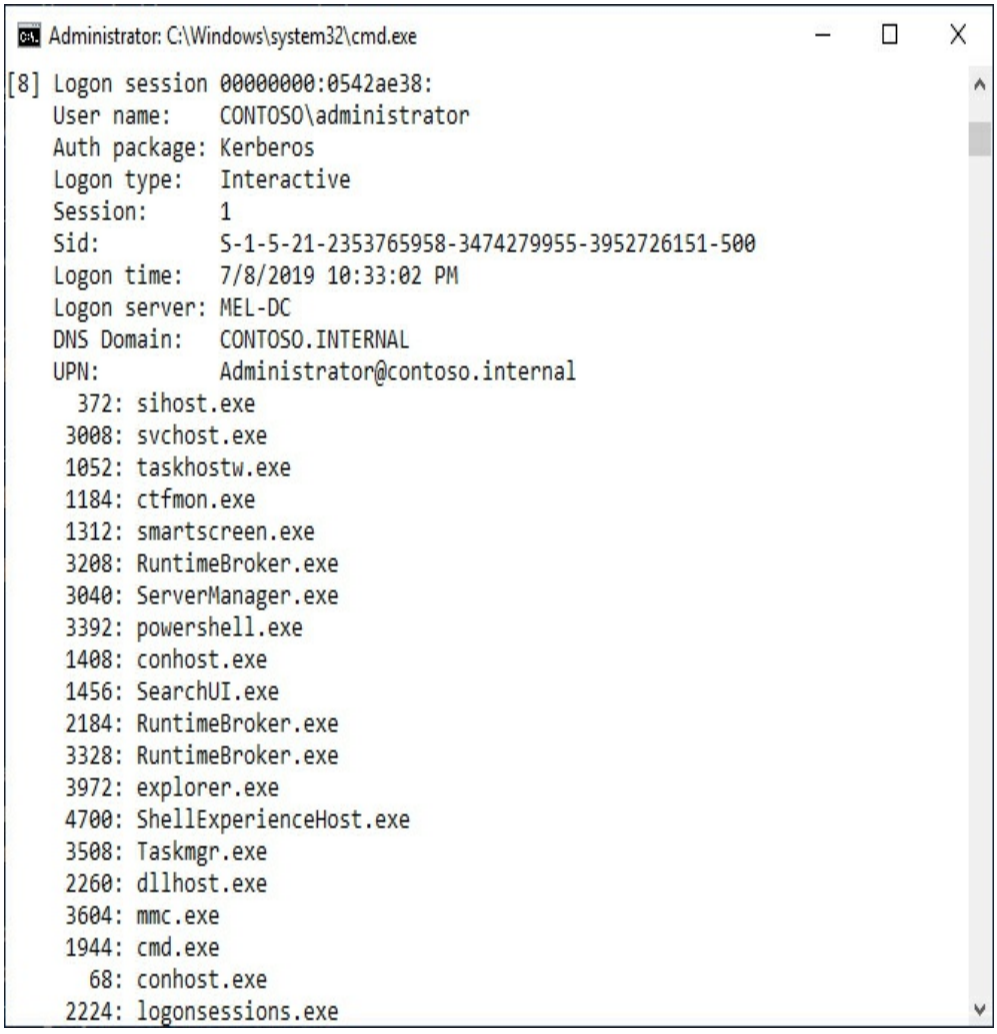


Figure 23-10 Logon Sessions

More Info

LogonSessions

You can learn more about LogonSessions at https://docs.microsoft.com/en-us/sysinternals/downloads/logonsessions.

Active Directory Explorer

Active Directory Explorer (AD Explorer) provides you with a detailed view of and ability to edit Active Directory objects. It's in some ways a more advanced and substantially more user-friendly version of the ADSIEdit utility that is built in to Windows Server 2019. You can use AD Explorer to view Active Directory object properties and attributes, as shown in Figure 23-11.

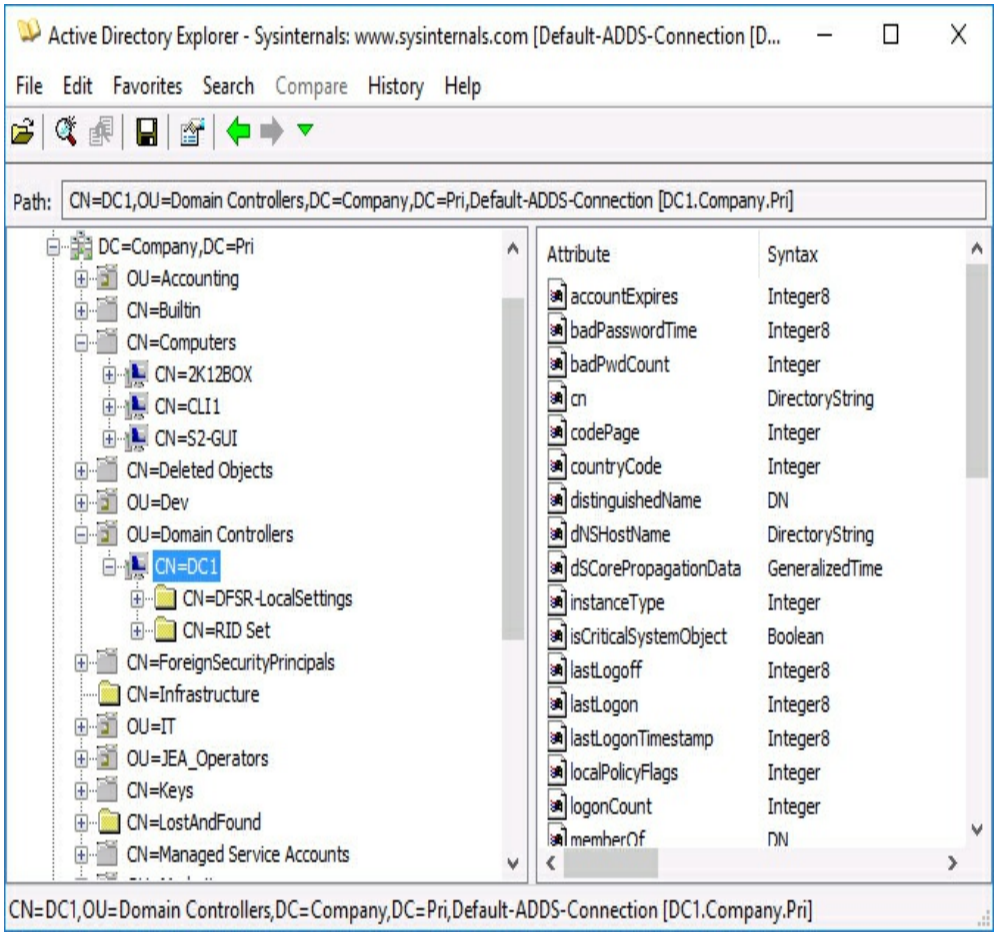


Figure 23-11 Active Directory Explorer

You can also use AD Explorer to save snapshots of the AD database for offline viewing. You can also use the tool to directly compare separate snapshots of the AD database, which can be very useful in recovery scenarios where you have several copies of the AD database and are unsure which version to restore to your production environment.

When connecting to a live AD instance, specify the domain name, user name, and password as shown in

Figure 23-12. You can also save connection properties, allowing you to reconnect to different domains as necessary.

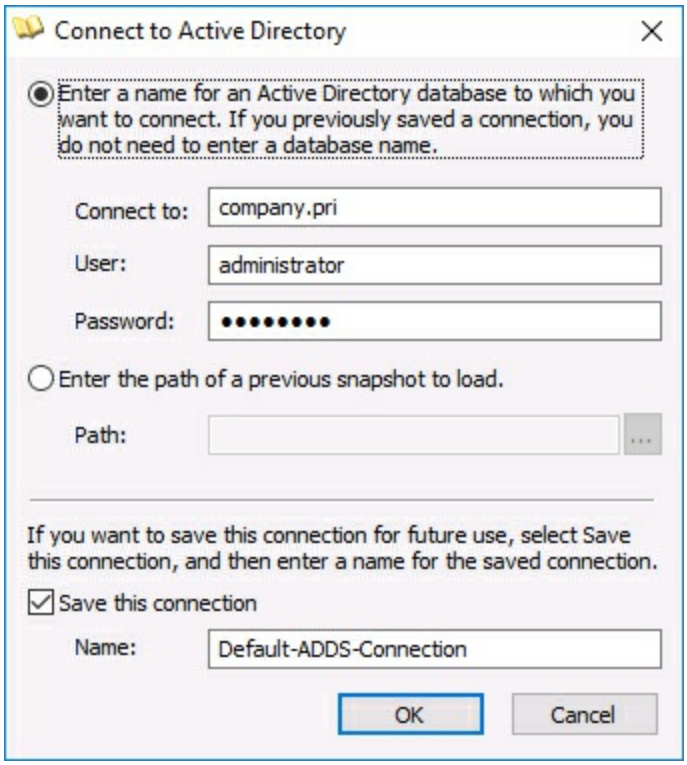


Figure 23-12 Connect to AD Database

When you specify the domain name, AD Explorer will connect to a domain controller in the same site. If you want to connect to a specific domain controller, you can specify the address of the domain controller using FQDN or the IP address. In addition to mounting existing snapshots of the Active Directory database, you can also use AD Explorer to connect to AD LDS instances.

You can use AD Explorer to rename or delete objects. You can also view and modify individual attributes of an object on the Attributes tab of the object's properties, as shown in Figure 23-13.

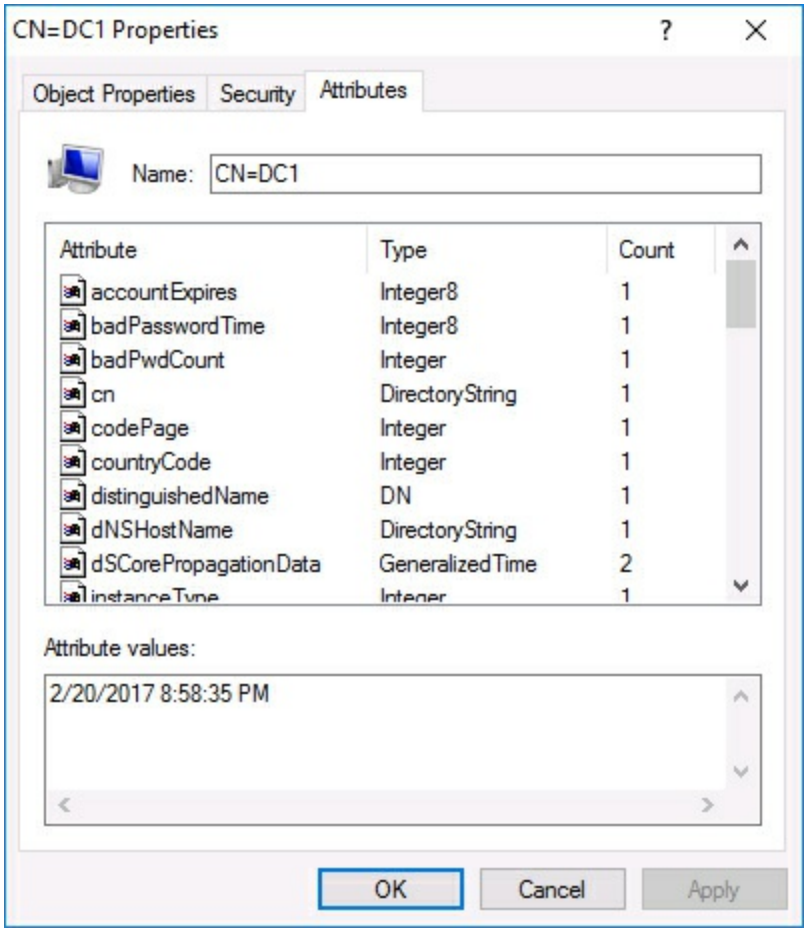


Figure 23-13 Object properties

You can also use AD Explorer to create new objects. When you create a new object, you must specify which class it belongs to, as shown in Figure 23-14. Object names must begin with the designator CN= and must be

unique within the container in which you create the object. The object's attribute list will automatically be created and set based on the attributes that are mandatory for the class of object that you have selected. Depending on the class you select, you'll need to configure several of these objects before you can successfully create the object.

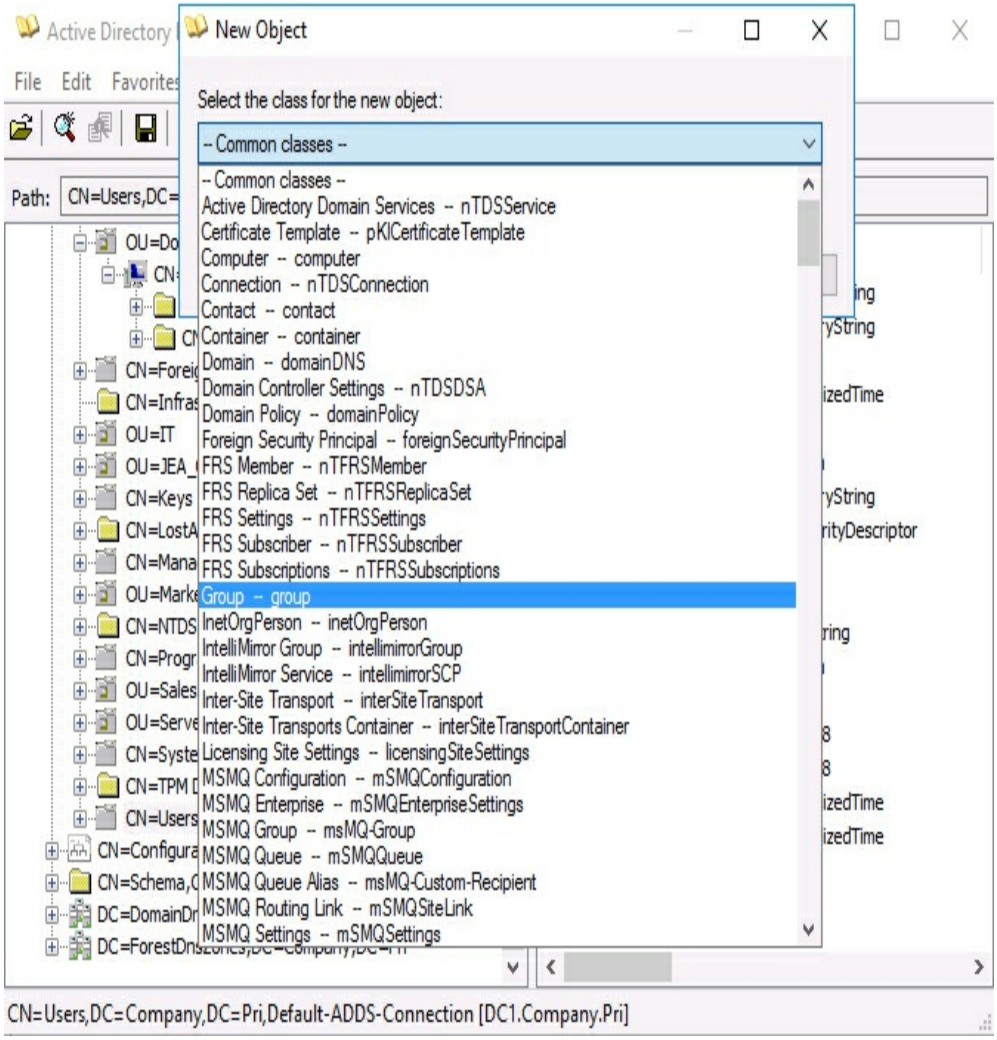


Figure 23-14 Create object

More Info

Active Directory Explorer

You can learn more about ADExplorer at https://docs.microsoft.com/en-us/sysinternals/downloads/adexplorer.

Insight for Active Directory

ADInsight (Insight for Active Directory) allows you to monitor LDAP traffic. You can use ADInsight to resolve problems with Windows authentication, DNS, and other applications and services that interact with Active Directory.

ADInsight allows you to troubleshoot LDAP traffic by intercepting calls that applications make to the *wldap32.dll* library. This library is the foundation used for Active Directory APIs including ADSI and ldap. ADInsight differs from traditional network monitoring tools in that it will intercept and interpret all Active Directory related client-side APIs. This includes those Active Directory related client-side APIs that do not directly transmit data to a domain controller. When used with administrator privileges, ADInsight allows you to monitor Windows services and other system processes.

More Info

Insight for Active Directory

You can learn more about Insight for Active Directory at https://docs.microsoft.com/en-us/sysinternals/downloads/adinsight.

PsPing

You can use PsPing as an advanced network diagnostic utility that has a superset of the ping.exe utility's functionality. PsPing allows you to measure TCP ping, network latency, and network bandwidth. PsPing can report information to within 0.01 milliseconds, which is 100 times more accurate than the standard Ping utility. Use the following commands to determine information about each mode:

- **Psping -? i.** Provides information on ICMP ping functionality.
- **Psping -? t.** Provides information on TCP ping functionality.
- **Psping -? l.** Provides information on TCP/UDP latency tests.
- **Psping -? b.** Shown in Figure 23-15, provides information on TCP/UDP bandwidth tests.

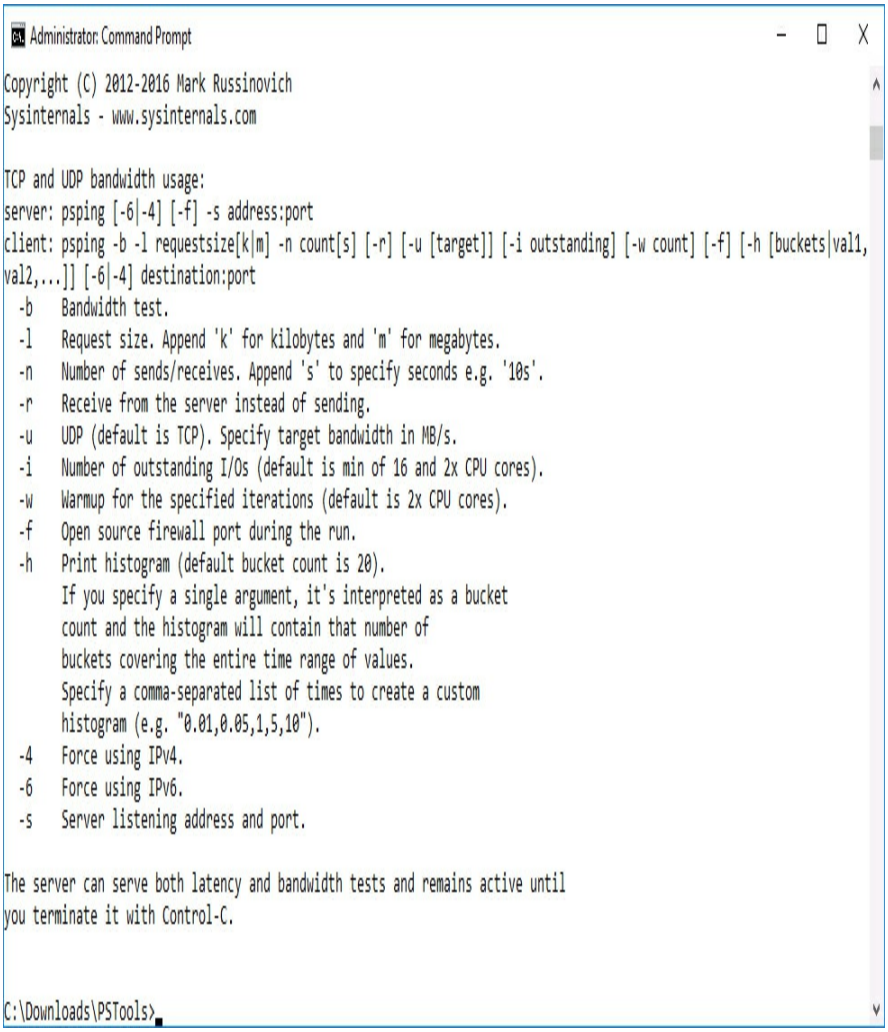


Figure 23-15 PsPing bandwidth check

More Info

PsPing

You can learn more about PsPing at https://docs.microsoft.com/en-us/sysinternals/downloads/psping.

RAMMap

RAMMap provides advanced information about how physical memory is used on computers running

Windows operating systems. The separate tabs of RAMMap provide the following information:

- **Use Counts.** Shown in Figure 23-16, this tab provides summary information by type. Categories include Process Private, Mapped File, Sharable, Page Table, Paged Pool, Nonpaged Pool, System PTE, Session Private, Metafile, Address Windowing Extension (AWE), Driver Locked, Kernel Stack, Large Page, and Unused.

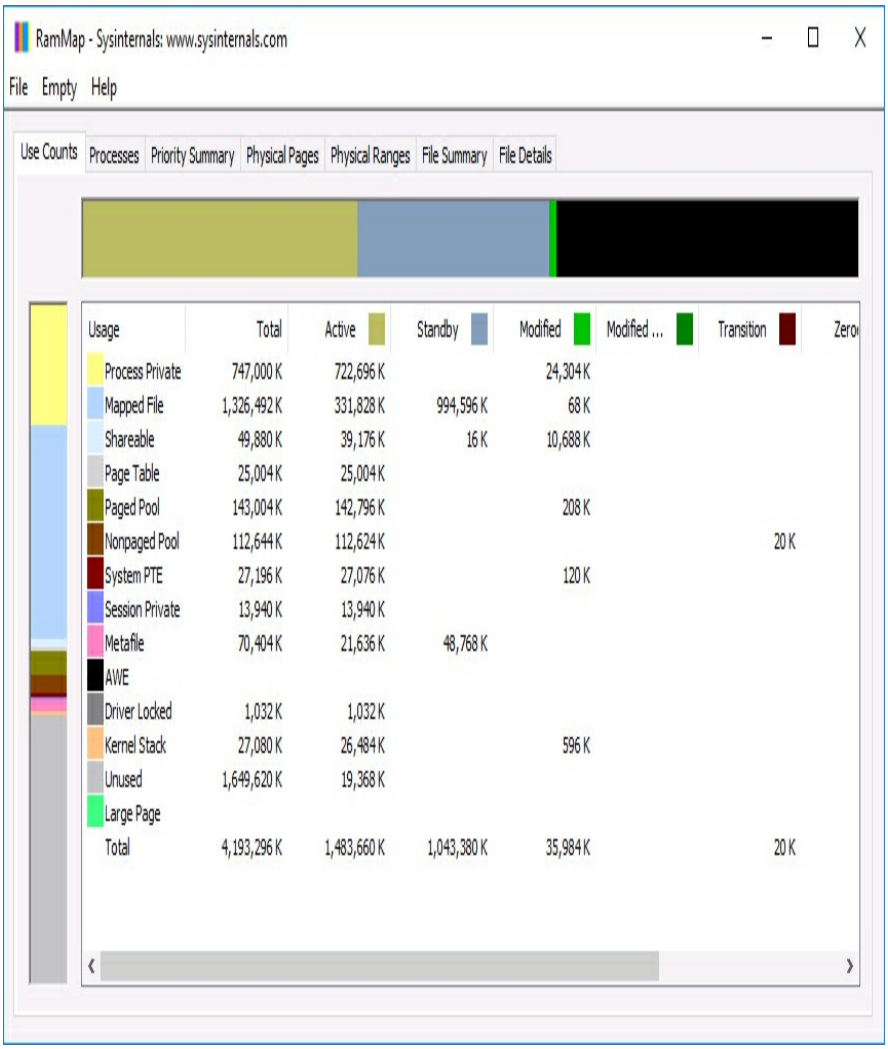


Figure 23-16 RAMMap Use Counts

- **Processes.** Provides information about the size of process working

sets. Includes information about PID, Private memory, Standby, Modified, Page Table, and Total physical memory utilization.

- **Priority Summary.** Shows the size of prioritized standby lists. If items with priority 5 and higher show high repurpose counts, it's likely that the server is under memory pressure and would likely benefit from the addition of more physical memory.
- **Physical Pages.** Shows per page information of all physical memory in the computer. Provides physical page information including physical address, list membership, allocation type, priority, whether the page stores part of an image, offset within a page table, file name of mapped file, originating process, virtual address, and any paged and nonpaged memory pool tag.
- **Physical Ranges.** Lists the physical ranges used in terms of start address, end address, and overall size.
- **File Summary.** Provides details of files that are stored in physical memory.
- **File Details.** Shown in Figure 23-17, displays memory pages in physical RAM on a per file basis.

Technet24

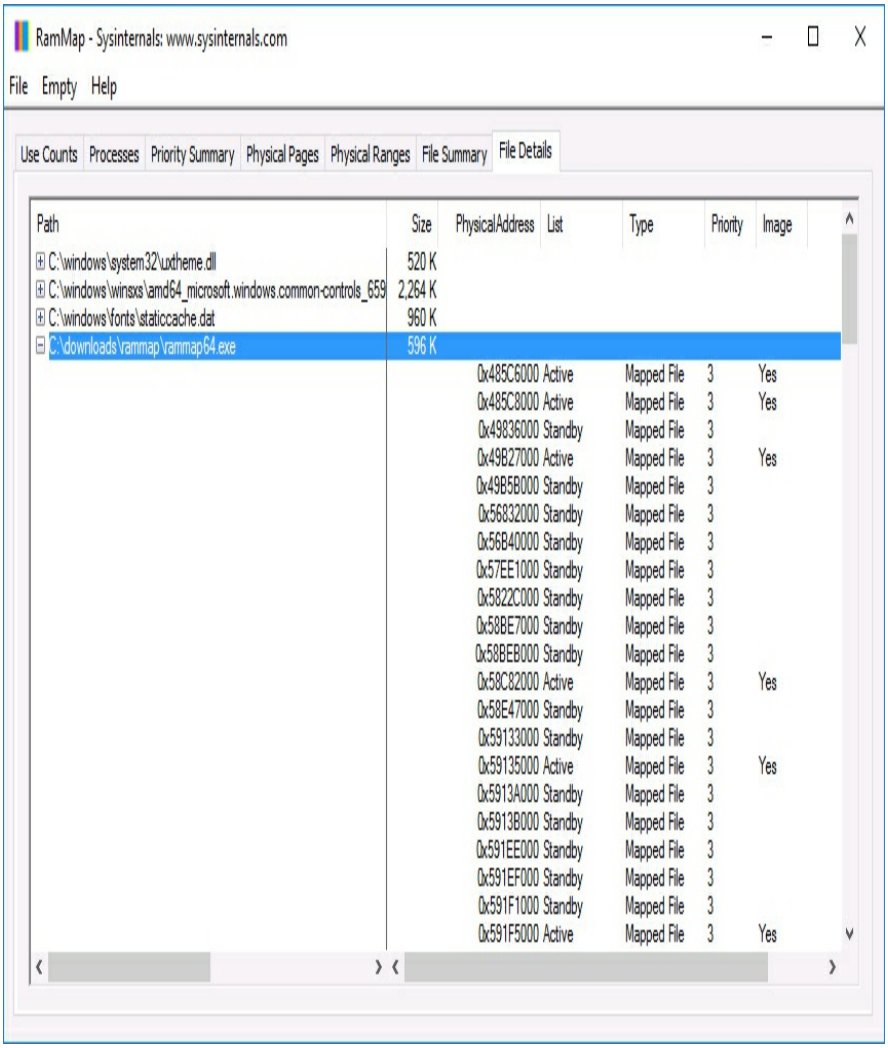


Figure 23-17 RAMMap File Details

You use RAMMap when you need to understand how Windows Server 2019 manages memory. For example, you can use it to determine how physical memory is allocated to a specific process or file. You can search memory for a specific file or process by using Ctrl+F. You can also RAMMap to take snapshots of how physical memory is being allocated at a specific point in time for later analysis. You can also use RAMMap to

purge portions of RAM when testing memory management scenarios.

More Info

RAMMap

You can learn more about RAMMap at https://docs.microsoft.com/en-us/sysinternals/downloads/rammap.

Inside OUT

Support calls

It's an old truism in IT that when all else fails, read the manual, and when that fails, call support. The trick to calling support is to have documented enough about the problem as possible so that support doesn't need to spend time going back over pathways that you've already fruitlessly explored. The more information you can provide to support, the more quickly they can get down to investigating your issue and hopefully providing you with a resolution. You'll get a lot more out of your support experience if you go into it prepared rather than if you approach support without having done your homework and hoping that they can magically resolve everything for you.

Index

A

access rights, 597. *See also* file access
access rules, DAC, 459–460
AccessChk tool, 740
Access-Denied Assistance, DAC, 292–293, 463–464
AccessEnum tool, 744–745
account security options, 586–587
accounts
    AD environment, 130–138
    creating in containers, 587
    rights assignments, 598
    types, 130–138
ACLs (access control lists), command-line tool, 730
AD (Active Directory). *See also* Azure Active Directory
    accounts, 130–138
    authoritative restore, 149–151
    consoles, 12
    Container Service Accounts, 326
    Domains And Trusts console, 117
    editing objects, 746–748

Enterprise CA, 393
Enterprise Subordinate CA, 402–403
FRS to DFSR migration, 695–696
migration to forest, 696–698
Migration Tool, 697–698
module cmdlets, 152–155
overview, 109
Recycle Bin, 147–149
remote administration, 110
Sites And Services Console, 114–117
snapshots, 151–152
Standalone CAs, 403–406
upgrade and migration paths, 693–695
upgrading domain controllers, 693–695
Users And Computers console, 113–114
AD CS (Active Directory Certificate Services). *See also* CAs (Certificate Authorities)
Certutil.exe and *Certreq.exe*, 435–436
migration, 699–704
overview, 391
PowerShell cmdlets, 433–434
role services, 412
roles, 412
AD DS structure, 125–130, 152–157

Technet24

AD Explorer, 746–749
AD FS (Federation Services)
attribute stores, 438, 442
certificate relationship, 441–442
claim rules, 443–444
claims and claim rules, 438
claims provider, 439
claims provider trust, 440–441
components, 437
multifactor authentication, 449
overview, 437
PowerShell cmdlets, 450–453
relying party, 439–440
web application proxy servers, 445–447
WebApplicationProxy module, 453
Workplace Join, 447–449
AD RMS. *See also* DAC (Dynamic Access Control)
administrators and super users, 469–470
certificates and licenses, 465–466
exclusion policies, 471–472
installing, 464–465
PowerShell cmdlets, 473–474
publishing domains, 469–470
templates, 466–469, 471–472

TPDs (Trusted Publishing Domains), 471
TUDs (Trusted User Domains), 470–471
ADAC (Active Directory Administrative Center), 110–113
ADCSAdministration PowerShell module cmdlets, 433–434
address resolution protocol cache, managing, 729
ADDSDeployment module cmdlets, 157
ADInsight tool, 749
administrative permissions, IIS, 331. *See also* permissions
admin-trusted attestation, shielded VMs, 631–633
ADModify.Net tool, 553
ADMX file format, 145
ADSI Edit console, 12
advanced auditing, 615–617
AGPM (Advanced Group Policy Management), 139–140
alerts, performance counter, 655
alias records, 166
AliasDefinitions role-capability file, 606
answer files, 57–59
antimalware solution, 650. *See also* malware toolchains
appcmd.exe utility, 312–313
application pools, IIS, 326–329
applications, launching with ShellRunAs, 745

Technet24

application-specific backups, 665
archived certificate, recovering, 431
arp.exe command-line tool, 729
ASP.NET Core Runtime, 346
AssembliesToLoad role-capability file, 607
assume breach, 635
at.exe command-line tool, 729
Attack Surface Analyzer, 638–639
attestation types, shielded VMs, 632
attrib.exe command-line tool, 729
attribute stores, AD FS, 438, 442
attributes, managing, 729
audit policies, command-line tool, 729
auditing and security, rights assignment, 599
auditing policies, GPOs, 615–617
auditpol.exe command-line tool, 617, 729
authentication
    Azure AD Connect, 540
    IPsec connections, 622
    Kerberos V5, 622
    NTLMv2, 622
    policy silos, 589–590
    rights assignment, 598
    VPNs (virtual private networks), 482

websites, 318–319
authoritative restore, AD (Active Directory), 149–151
authorization rules, URLs, 323–324
automating tasks, 2
Autounattend.xml file, 57–58
Azure
Chef Infra Server, 85
and Remote Desktop, 25
Azure Active Directory. *See also* AD (Active Directory)
overview, 538–539
password protection, 562–563
self-service password reset, 560–561
Azure AD Connect
connectivity requirements, 542
deployment accounts, 543–544
Directory Federation, 541
Health tool, 554
installing, 544–551
non-routable domains, 551–553
object filters, 557–559
overview, 540–541
pass-through authentication, 540
passwords, 540
server requirements, 541–544

Technet24

SQL Server requirements, 543
synchronization, 555–559
UPN suffixes, 551–553
User sign-in, 546
Azure AD DS (Active Directory Domain Services), 563–564
Azure Arc, 572
Azure Backup, 666–670. *See also* backup; Windows Server Backup
Azure Bastion, 533–534
Azure Cli, 346
Azure File Sync, 569–571
Azure Hybrid services, 6, 564–567, 681
Azure Iaas VMs
creating, 520–524
WAC (Windows Admin Center), 567–569
Azure Monitor, 661–662
Azure Network Adapter, 573
Azure Site Recovery, 572–573
Azure Update Management, 679–682. *See also* updates
Azure VPN gateways, 531

B

backup. *See also* Azure Backup; Windows Server Backup

command-line tool, 733
and restore, 140–141
rights, 597
security, 663
using with WAC (Windows Admin Center), 5
bandwidth management, 207
bare metal chassis, 47, 68, 75–77
base images, containers, 345–347
batch job, rights assignment, 599
bcdboot.exe command-line tool, 729
bcdedit.exe command-line tool, 729
BitLocker, 2, 731–732
blocking writing files, 282
boot environment, command-line tool, 729
boot.wim file, 48
brace matching, PowerShell, 22
BranchCache, 299–302, 305–306

C

CA management
backup and recovery, 425–427
certificate requests, 424–425
key archiving and recovery, 427–431
overview, 422–424

Technet24

CAL licenses, RDS, 514
CAPolicy.inf file, 432
CAs (Certificate Authorities). *See also* AD CS (Active Directory Certificate Services)
autoenrollment and renewal, 420–422
backup and recovery, 425–427
certificates, 391–393
command-line tools, 729
CRLs (certificate revocation lists), 406–412
customizing settings, 432
deployments, 394
Enterprise CAs, 393–403
hash algorithms, 399–400
hierarchies, 392
key archiving and recovery, 427–431
KRA key, 431
management, 422–424
PKIview, 392
private key cryptography, 399–400
renewing, 392
requests, 424–425
Standalone CAs, 403–406
templates, 413–420
types, 391–392

CAU (Cluster Aware Updating), clusters, 373–374
CCI (Configurable Code Integrity) policy, 642
Center for Internet Security, 596
central access policies, DAC, 461–462
central access rules, DAC, 459–460
certificate chains, viewing, 739–740
certificate relationship, AD FS, 441–442
Certificate Services, managing, 435–436
Certificates, using with WAC (Windows Admin Center), 5
Certification Authority console, 12
certreq.exe command-line tool, 436, 729
certutil.exe command-line tool, 435, 729
CFG (Control Flow Guard), 647
checkpoints, Hyper-V virtualization platform, 202–203
checkpoints, managing, 202–203
Chef Infra Server
    agents, 94–95
    cookbooks and recipes, 95
    desired state, 46
    vs. DSC, and Puppet, 103
    servers, 85–88
ChefDK (Development Kit), 89–94
chkdsk.exe command-line tool, 729

Technet24

chkntfs.exe command-line tool, 729
Chocolatey package manager, 11, 105–107
cipher.exe command-line tool, 729
claims and claim rules, AD FS, 438, 443–444. *See also* user and device claims
claims provider, AD FS, 439
claims provider trust, AD FS, 440–441, 444
Client Naming Policy, WDS, 64
clip.exe command-line tool, 729
cloud witness and clusters, 376–377
cluster networking, Hyper-V, 214–215
cluster quorum, Hyper-V, 213–214, 370–371
cluster storage, Hyper-V, 212
Cluster-Aware Updating console, 12
clusterng, future of, 370
clusters
    CAU (Cluster Aware Updating), 373–374
    host management, 387–388
    MPIO (multipath I/O), 373
    networks, 372–373
    NLB operation modes, 386–387
    Rolling Upgrades, 379–380
    sets, 382
    storage and shared volumes, 372

workgroups, 381
cmd.exe, using with containers, 347
CNAME records, 166
color coding, PowerShell, 21
command prompt, 36
command-line tools, 729–733. *See also* troubleshooting methodology
command-line output, redirecting, 729
compact.exe command-line tool, 729
comp.exe command-line tool, 729
Component Services console, 12
compression of files, command-line tool, 729
computer access rights, 597
computer accounts, 132–133, 587
Computer Management console, 12
computer name, viewing, 730
Connection Manager Administration Kit console, 12
connection security rules, WFAS, 623–627. *See also* security
consoles, 11–17
container images, 353–354. *See also* images; OS images
containers. *See also* Docker engine; LCOW (Linux Containers on Windows)
applying updates, 356–357

Technet24

and Chocolatey, 351
concepts, 337–339
creating accounts in, 587
creating images from, 351
and data, 339
Dockerfiles, 351–352
images, 353–354
isolation modes, 339–340
listing, 348–349
modifying running, 350
networking, 357–363
orchestration, 364–368
registries and images, 345–347
retrieving OS images, 344–345
service accounts, 355–356
starting, 347–350
using with WAC (Windows Admin Center), 5
context-sensitive help, PowerShell, 22
control panels, 37
Controlled Folder Access, 645–647
CPU utilization spikes, monitoring, 736
Credential Guard, 2, 640–642
CRLs (certificate revocation lists), 406–412
cscript.exe command-line tool, 729

CSP (Cryptographic Service Provider), 465, 700–701
CSVs (Cluster Shared Volumes), Hyper-V, 215–216

D

DAC (Dynamic Access Control). *See also* AD RMS
  Access-Denied Assistance, 463–464
  central access policies, 461–462
  central access rules, 459–460
  group policy, 456
  overview, 455–456
  PowerShell cmdlets, 474
  resource properties, 457–459
  staging, 463
  user and device claims, 456–457
DANE (DNS-based Authentication of Named Entities), 172–173
data, backing up, 664
data collector sets, 653–654
Datacenter edition, 31–32, 687
date panel, 37
date.exe command-line tool, 729
dcgpofix.exe command-line tool, 730
DCs (domain controllers)
  and Credential Guard, 642

Technet24

deployment, 118–120
evaluation, 689
global catalog servers, 121
hardening, 582–583
installation options, 120
Server Core, 120–121
as targets, 109
upgrading, 693–695
virtual domain controller cloning, 124
virtualization platforms, 121
DDA (Discrete Device Assignment), 195–196, 512
debug programs, rights assignment, 598
debugging, PowerShell, 21
deduplication
enabling, 273
PowerShell cmdlets, 266
process, 260–263
support, 204
default groups, 134–136
Defragment and Optimize Drives console, 12
deleted items, restoring, 119–120, 146–152
Demon.json file, Docker, 342–344
deny access, rights assignment, 598
DEP (Data Execution Prevention), 647

dependencies, 338, 726–727
Desired State Configuration (DSC), 46, 57, 81–85, 103
Desktop Experience installation options, 42
detached clusters, 216
device and user claims, configuring for DAC, 456–457
device drivers
  command-line tool, 730
  rights assignment, 599
Device Guard policy, 642
Devices, using with WAC (Windows Admin Center), 5
DFS (Distributed File System)
  migration from FRS (File Replication Service), 695–696
  namespace, 293–295
  PowerShell cmdlets, 306–308
  replication, 296–299
  witnesses, 370
DFS Management console, 12
DFSR (distributed file system replication), 730
dfsrmig.exe command-line tool, 730
DHCP (Dynamic Host Configuration Protocol)
  administration, 182–184
  console, 12
  failover, 181

Technet24

filtering, 179
migration, 705–709
multicast scopes, 180
name protection, 180–181
reservations, 178–179
scopes, 177
server and scope options, 178
split scopes, 180
superscopes, 179–180
using with WAC (Windows Admin Center), 5
diagnostic hypothesis, 726
digital signature details, viewing, 739–740
DirectAccess
configuring, 493–496
overview, 489–490
server, 490–492
topologies, 490
directories, rights assignment, 600. *See also* mounted directories
directory browsing, website management, 321–322
directory service data, rights assignment, 600
directory structure, command-line tool, 732
DISA (Defense Information Systems Agency), 596
disk checking, command-line tool, 729

Disk Cleanup console, 12
diskshadow.exe command-line tool, 730
DISM (Deployment Image Servicing and Management) tool, 49–50, 52
DNS (Domain Name Service)
cache locking, 170–171
conditional forwarders, 163–164
console, 12
delegated administration, 174
event logs, 169–170
forwarders, 163
GlobalNames zones, 164–165
managing with PowerShell, 174–176
migration, 704–705
netmask ordering, 171
PNRP (Peer Name Resolution Protocol), 165–166
policies, 173
recursion, 171
resource records, 166–167
response rate limiting, 171–172
reverse lookup zones, 162
scavenging, 167–168
socket pool, 170
stub zones, 164

Technet24

using with WAC (Windows Admin Center), 5
zone aging, 167–168
zone delegation, 162–163
zone types, 159
DNS information, command-line tool, 731
DNS servers, command-line tool, 730
dnscmd command-line tool, 170
dnscmd.exe command-line tool, 730
DNSSEC (Domain Name System Security Extensions), 168–169
Docker engine. *See also* containers
Demon.json file, 342–344
Experimental mode, 363
installing, 341–342
Linux containers on Windows, 363–364
overview, 340
registries and images, 345–347
retrieving container OS image, 344–345
docker ps command, 360
Docker Swarm orchestration, 365–368
Dockerfiles, using, 351–352
docking station, rights assignment, 599
domain controller cloning, Hyper-V virtualization platform, 223

domain dominance, 109
domain local groups, 133
domain name filtering and IP addresses, 322–323
domain names, repairing dependencies, 730
domain naming master, 129
domains, AD DS structure, 125–126
downloading files, 44
driver verification, command-line tool, 732
driverquery.exe command-line tool, 730
drivers and updates
  adding to images, 53–54
  command-line tool, 731
DSC (Desired State Configuration), 46, 57, 81–85, 103
DSRM (Directory Services Restore Mode), 119–120
dynamic memory, Hyper-V virtualization platform, 191–192
dynamic virtual machine queue, 207

E

editions of Windows Server, upgrading, 31–33, 687, 689
EMET (Enhanced Mitigation Experience Toolkit), 647
encryption
  command-line tool, 729
  IPsec, 622

Technet24

encryption-supported VMs, 633. *See also* VMs (Virtual Machines)
Enhanced Session Mode, Hyper-V virtualization platform, 195
Enterprise CA, 393
Enterprise Root CA, 394–402
Enterprise Subordinate CA, 402–403
EnvironmentVariables role-capability file, 607
error response, modifying in IIS, 319
ESAE (Enhanced Security Administrative Environment), 127, 594–596. *See also* security
Essentials edition, 31, 687
evaluation version, converting to licensed version, 688–689
event logs
  DNS, 169–170
  event customization, 730
  PowerShell cmdlets, 683–684
  and publishers, 733
  system activity, 741–744
event subscriptions, command-line tool, 733
Event Trace Logs, command-line tool, 732
Event Trace Sessions, command-line tool, 731
Event Viewer, 12, 655–659

eventcreate.exe command-line tool, 730
events, using with WAC (Windows Admin Center), 5
exclusion policies, AD RMS, 471–472
Exploit Protection, 647–649

F

Failover Cluster Manager console, 12
failover clustering
- CAU (Cluster Aware Updating), 373–374
- cloud witness, 376–377
- cluster quorum modes, 370–371
- cluster sets, 382
- Hyper-V, 212–217
- MPIO (multipath I/O), 373
- multisite clusters, 375–376
- networks, 372–373
- NLB (Network Load Balancing), 385–389
- overview, 369–370
- PowerShell cmdlets, 383–385
- preference settings, 374–375
- rolling upgrades, 379–380
- shared volumes, 372
- storage spaces, 372
- VMs (Virtual Machines), 377–379

Technet24

workgroup clusters, 381
Fax Service Manager console, 12
features and roles, adding, 45–46, 54–56
Federation Server, 437
file access
recovering, 732
speeding up, 299–302
file and storage servers. *See also* file and storage servers
Storage Migration Service, 710–718
WSMT (Windows Server Migration Tools), 718–722
File Explorer, shared folder permissions, 275–276
file management, command-line tool, 729
file screens, FSRM, 282–286
File Server Resource Manager console, 13
file servers
BranchCache, 299–302
BranchCache cmdlets, 305–306
cleaning up, 292
DFS (Distributed File System), 293–299
DFS cmdlets, 306–308
FSRM cmdlets, 303–305
overview, 273
protecting from ransomware, 645
Shared Folder cmdlets, 302–303

shared folder permissions, 274–279
Storage Migration Service, 710–718
workgroup clusters, 381
WSMT (Windows Server Migration Tools), 718–722
file systems command-line tool, 729
file version number, viewing, 739–740
files. *See also* *Openfiles.exe* command-line tool; protected system files; system files
auditing, 616–617
comparing, 729
compression, 729
downloading, 44
locating, 733
rights assignment, 600
updating, 673–674
using with WAC (Windows Admin Center), 5
find.exe command-line tool, 730
findstr.exe command-line tool, 730
firewall
profiles, 619–620
using with WAC (Windows Admin Center), 683–683
firmware environment values, rights assignment, 599
folder level quotas, FSRM, 280–282
folders

Technet24

access control, 645–647
auditing, 616–617
sharing, 258–259, 274–279
force shutdown, rights assignment, 598
forests
AD DS structure, 126–127
migration to, 696–698
policy in, 142
format.exe command-line tool, 730
FormatsToProcess role-capability file, 607
FRS (File Replication Service), to DFS (Distributed File System) migration, 695–696
FSMO (Flexible Single Master Operations), 128–130
FSRM (File Server Resource Manager)
Access-Denied Assistance, 292–293
file classification, 288–290
file management tasks, 290–292
file screens, 282–286
folder level quotas, 280–282
PowerShell cmdlets, 303–305
storage reports, 286–287
fsutil.exe command-line tool, 730
FTP protocol, managing, 332–334
FunctionDefinitions role-capability file, 606

G

garbage collection, 262
Generation 2 VMs, Hyper-V virtualization platform, 194–195
Get-Credential command, PowerShell, 20
getmac.exe command-line tool, 730
global catalog servers, 121
global groups, 133
global objects, rights assignment, 598
Global Resource Property list, DAC, 458
gMSA (group-managed service account), 137–138, 355–356, 601–603
gpfixup.exe command-line tool, 730
GPOs (Group Policy Objects)
    AD DS structure, 127–128
    audit policies, 615
    command-line tools, 730
    comparing, 636
    DAC (Dynamic Access Control), 456
    management, 139–141
    preferences, 143
    RDS (Remote Desktop Services), 507–509
gpresult.exe command-line tool, 730
gpupdate.exe command-line tool, 730

Technet24

group accounts, 133–136
group policy
  administrative templates, 145–146
  management, 139–141
  policy processing, 141–142
  preferences, 143–145
Group Policy Management console, 13
guarded fabric, shielded VMs, 631–633

H

hardening, process, 581
hardening Active Directory
  account security option, 586–587
  authentication policies silos, 589–590
  block server operations, 592
  disable NTLM, 591
  domain controllers, 582–583
  ESAE (Enhanced Security Administrative Environment), 594–596
  KRBTGT account password, 593–594
  least privilege, 583–584
  LSA (Local Security Authority) protection, 592–593
  password policies, 585–586
  protected accounts, 588–589

RBAC (Role-Based Access Control), 584
scheduling tasks, 592
hardening Windows Server
advanced auditing, 615–617
JEA (Just Enough Administration), 603–609, 613
LAPS (Local Administrator Password Solution), 613–615
PAM (Privileged Access Management), 609–613
service accounts, 600–603
user rights, 596–600
hash algorithms, CAs (Certificate Authorities), 399–400
HGS (Host Guardian Service), 631
host records, 166
hostname.exe command-line tool, 730
HVC for Linux, Hyper-V virtualization platform, 198
Hyper-V Manager console, 13
Hyper-V virtualization platform
checkpoints, 202–203
and containers, 340
copying VMs, 221–222
DDA (Discrete Device Assignment), 195–196
domain controller cloning, 223
dynamic memory, 191–192
Enhanced Session Mode, 195

Technet24

exporting VMs, 221–222
failover clusters, 212–217
Generation 2 VMs, 194–195
guest clusters, 217–219
guest integration services, 193–194
host validation, 631
HVC for Linux, 198
importing VMs, 221–222
live migration, 219–221
nested virtualization, 197–198
network isolation, 209
optimizing network performance, 206–208
PowerShell cmdlets, 224–227
PowerShell Direct, 198
private switches, 206
replica, 209–212
resource metering, 193
shielded virtual machines, 223
smart paging, 192–193
storage migration, 221
storage optimization, 204–205
Storage QoS (Quality of Service), 204
V VHD Sets, 219
Virtual Fibre Channel adapters, 203

virtual hard disks, 199–202
virtual machine network adapters, 206
virtual switches, 205–206
VM drain on shutdown, 222
VM MAC addresses, 208
VM network adapters, 206
workgroup clusters, 381
hypothetical solutions, ranking, 727–728
Iaas (Azure Infrastructure as a Service) VMs. *See also* VMs (Virtual Machines)
administration, 532–538
Azure Bastion, 533–534
backup, 536–537
creating from WAC, 567–569
encryption, 537–538
networking, 524–531
NSGs (Network Security Groups), 530–531
overview, 519–524
ICMP (Internet Control Message Protocol), 621, 731
IDE (Integrated Development Environment), 23
idempotency, 80
IIS (Internet Information Services). *See also* website management
administrative permissions, 331

Technet24

application pools, 326–329
container operating system, 346
FTP management, 332–334
Management console, 332–334
missing items, 310
overview, 309–310
PowerShell cmdlets, 334–336
test certificates, 312
users and delegation, 329–331
IKEv2 (Internet Key Exchange version 2), VPNs, 479–480
images. *See also* container images; OS images
benefits, 48
build and capture, 56–57
committing, 56
configuring WDS, 62–66
creating from containers, 351
drivers and updates, 53–54
listing, 349
managing with WDS, 62
modifying, 49–50
mounting, 50–53
roles and features, 54–55
servicing, 50

infrastructure configuration as code, 79–80
Infrastructure Developer, 80
infrastructure master, 130
in-place upgrade, 686
Insider Preview builds, 34–35
installation options, 35–36
Installed Apps, using with WAC (Windows Admin Center), 5
install.wim file, 44, 48
integrity scrubbing, 262
IntelliSense, PowerShell, 21
Intl.cpl, 37
IoT Core base image, 346
IP addresses. *See also* NAT (Network Address Translation)
  and domain name filtering, 322–323
  Iaas VMs, 527–530
IPAM (IP Address Management), 185–190
ipconfig.exe command-line tool, 730
IPsec, configuring for WFAS, 621–622
IPv6
  DHCP, 178
  PNRP (Peer Name Resolution Protocol), 165–166
iSCSI

Technet24

PowerShell cmdlets, 266–267
and storage devices, 252–255, 266–267
iSCSI Initiator consoles, 13, 37
iSNS (Internet Storage Name Service) server, 256–257
isolation modes, containers, 339–340

J

JEA (Just Enough Administration), 26, 603–609, 613
JIT (Just In Time). *See also* VMs (Virtual Machines)
Administration, 609–613
VM Access, 534–536
jump servers, using with PAW, 3

K

Kerberos protocol, command-line tool, 730
Kerberos V5 authentication, 622
kernel transaction manager utility, 731
klist.exe command-line tool, 730
KPS (Key Protection Service), 631
KRA key, CAs (Certificate Authorities), 431
KRBTGT account password, 593–594
ksetup.exe command-line tool, 730
ktmutil.exe command-line tool, 731
ktpass.exe command-line tool, 731

Kubernetes orchestration, 365

L

L2 (Layer 2) Bridge networks, 362
L2TP/IPsec protocols, VPNs, 481
LAN routing, 487
languages, updating, 672
LAPS (Local Administrator Password Solution), 613–615. *See also* passwords
lcacls.exe command-line tool, 730
LCM (Local Configuration Manager), DSC, 83
LCOW (Linux Containers on Windows), 346, 363, 575, 578. *See also* containers
LDAP traffic, monitoring, 749
least privilege, 583–584. *See also* privileged access
licensed version, converting to, 688–689
licensing
    Iaas VMs, 519
    RDS (Remote Desktop Services), 513–515
Linux containers on Windows, 363–364
Linux distro selection, Chef and Puppet, 96–97
Linux on Windows Server. *See* WSL (Windows Subsystem for Linux)
Local Security Policy console, 13
Local Service account, 136

Technet24

Local System account, 136
Local Users and Groups, using with WAC (Windows Admin Center), 5
local vs. remote administration, 1
locctr.exe command-line tool, 731
lock pages in memory rights assignment, 599
logman.exe command-line tool, 731
logoff.exe command-line tool, 731
logon rights, 586, 597–599
LogonSessions tool, 746
LSA (Local Security Authority) protection, 592–593
LTSC (Long Term Servicing Channel), 6, 33

M

MAC addresses
  command-line tool, 730
  enabling spoofing, 361
  VMs, 208
malware toolchains, 737. *See also* antimalware solution
manage-bde.exe command-line tool, 731
MD hash algorithms, CAs, 399–400
Message Analyzer, 660
Microsoft Update Catalog, 49
migration. *See also* upgrade and migration paths

AD CS (Active Directory Certificate Services), 699–704
DHCP, 705–709
DNS, 704–705
to forest, 696–698
FRS to DFSR, 695–696
vs. upgrade, 685
Migration Tool, Active Directory, 697–698
MIM (Microsoft Identity Manager), 610–611
mirroring and storage spaces, 234
mklink.exe command-line tool, 731
ModulesToImport role-capability file, 606
MOF file, DSC configuration script, 82–84
monitoring and maintenance, PowerShell cmdlets, 683–684
monitoring tool, Process Monitor, 735–736
mounted directories, command-line tool, 732. *See also* directories
mounting images, 50–53
mountvol.exe command-line tool, 731
MPIO (multipath I/O), clusters, 373
msiexec.exe command-line tool, 731
MSInfo32.exe, 37
multicast options, WDS, 65–66

Technet24

multifactor authentication, AD FS, 449–450
MX (mail exchanger) records, 166–167

N

name of computer, viewing, 730
Nano Server, 44, 345
NAT (Network Address Translation), 359–360, 487–488. *See also* IP addresses
nbstat.exe command-line tool, 731
nested virtualization, Hyper-V virtualization platform, 197–198
.NET Framework 3.5, 56
Net print command-line tool, 731
NetBIOS over TCP/IP command-line tool, 731. *See also* TCP/IP (Transmission Control Protocol/Internet Protocol)
netcfg.exe command-line tool, 731
netsh.exe command-line tool, 731
netstat.exe command-line tool, 731
network configuration, command-line tool, 731
network isolation, Hyper-V virtualization platform, 209
Network Load Balancing Manager console, 13
network monitoring, 659–660
network performance, optimizing, 206–208
Network Policy Server console, 13

Network Service account, 136
Network tool, using with WAC (Windows Admin Center), 5
networking Iaas VMs, 524–531
NFS (Network File System), 259–260, 267–268
NIC teaming, VMs, 208
NLB (Network Load Balancing), 385–389
nlbmgr.exe command-line tool, 731
NLS (Network Location Server), 492
Notepad, 37
NSGs (Network Security Groups), IaaS VMs, 530–531
nslookup.exe command-line tool, 726, 731
NTLM (NT Lan Manager), disabling, 591
NTLMv2 authentication, 622

O

object label, rights assignment, 599
ODBC Data Sources consoles, 13
Offline Root CAs, 406
Online Responder Management console, 13
OOBE (Out-of-Box Experience), 57
openfiles.exe command-line tool, 731. *See also* files
OpenSSH on Windows Server, 29
operating system, rights assignment, 597

Technet24

orchestration of containers, 364–368
OS container image, retrieving, 344–345
OS images, keeping, 327. *See also* images
OS Rolling Upgrade, clusters, 380
OU (organizational unit) structure, 127–128, 132–133
overlay networks, creating, 362, 367
Overview tool, using with WAC (Windows Admin Center), 5

P

package-management utilities, 103–107
pagefile, rights assignment, 597
PAM (Privileged Access Management), 609–613
parity data, storage, 234
passwords. *See also* LAPS (Local Administrator Password Solution)
  Azure Active Directory, 562–563
  KRBTGT account, 593–594
  policies, 585–586
  resetting, 585
  RODC replication, 122–123
  security options, 586
pathping.exe command-line tool, 731
paths, associating with volume letters, 732

paths, command-line tools, 732
PAW (Privileged Access Workstations), 2–3
PDC emulator, 129
PDK file, 634
performance counters, 655, 731
Performance logs, command-line tool, 731
Performance Monitor console, 13
permissions. *See also* administrative permissions
  AccessEnum tool, 744–745
  CA autoenrollment, 420
  shared folders, 274
  staging, 463
persistent memory, 229
ping.exe utility, 731, 750
PKIview, CAs (Certificate Authorities), 392
pnputil.exe command-line tool, 731
PNRP (Peer Name Resolution Protocol), 165–166
pnunattend.exe command-line tool, 731
Policy Analyzer tool, 636–638
policy processing, 141–142
port rules, NLB, 388
PowerShell
  DHCP management, 182–184
  direct connection, 24–25

Technet24

and DISM (Deployment Image Servicing and Management) tool, 52
DNS management, 174–176
documentation, 17–18
features, 17
Gallery, 19, 104–105
Get-Credential command, 20
help commands, 18
IPAM management, 188–190
ISE (Integrated Scripting Environment), 21–24
modules, 18
remote tabs, 22–23
remoting, 19–21
snippets, 23–24
storage-related cmdlets, 266–271
using with WAC (Windows Admin Center), 5, 10
wget command, 44
PowerShell cmdlets
AD (Active Directory) module, 152–155
AD Deployment module, 157
AD FS, 450–453
AD RMS, 473–474
ADCSAdministration, 433–434
BranchCache, 305–306

container images, 346
DAC (Dynamic Access Control), 474
deduplication module, 266
DFS (Distributed File System), 306–308
DHCP, 182–184
DNS, 174–176
Event Logs, 683–684
event logs, 683–684
failover clustering, 383–385
FSRM (File Server Resource Manager), 303–305
vs. functions and aliases, 605
Group Policy module, 156
Hyper-V virtualization platform, 224–227
IIS (Internet Information Services), 334–335
IPAM, 188–190
iSCSI module, 266
iSCSI Target module, 267
monitoring and maintenance, 683–684
NFS module, 267
NLB (Network Load Balancing), 389
Remote Access, 224–227, 496–498, 515–517
shared folder, 302–303
Storage module, 268
Storage Replica, 271

Technet24

web administration, 335–336
Web Application Proxy, 453
Windows Server Backup, 683–684
WSUS (Windows Server Update Services), 684
PowerShell Direct, Hyper-V virtualization platform, 198
PowerShell session
entering, 37
initiating remotely, 532–533
PowerShell.exe, 347
PPTP VPN protocol, 481–482
print jobs and queues, command-line tools, 731
Print Management console, 13
printer drivers, command-line tool, 731
printer queues, command-line tool, 732
printers, command-line tool, 732
privileges. *See also* least privilege
delegating, 584–585
securing, 596
prndrvr.exe command-line tool, 731
prnjobs.exe command-line tool, 731
prnmngr.exe command-line tool, 732
prnqctl.exe command-line tool, 732
ProcDump tool, 736
Process Explorer tool, 734–735

process level token, rights assignment, 600
Process Monitor tool, 735–736
process working set, rights assignment, 599
processes
command-line tools, 732
rights assignment, 599
using with WAC (Windows Admin Center), 5
Production Checkpoints feature, 121
products, updating, 672
programs, managing and executing remotely, 733
protected accounts, 588–589
protected system files, command-line tool, 732. *See also* files; system files
provisioning data file, 634
PsPing tool, 750
.psrc extension, 604–605
PsTools Suite, 737–738
PTR (pointer) records, 167
pubprn.exe command-line tool, 732
Puppet environment, 96–103
PXE (Pre-boot Execution Environment), 60, 63

R

RAC (rights account certificate), 466

Technet24

RAMMap tool, 751–753
ransomware, protecting file servers from, 645
RBAC (Role-Based Access Control), 584–585
RD (Remote Desktop) Gateway, 475–479
RD (Remote Desktop) Session Host
    Group Policy configuration, 507–509
    overview, 503–504
    personal session desktops, 506
    RemoteApp, 506–507
    session collections, 504–506
RD (Remote Desktop) Virtualization Host, 509–512
RD (Remote Desktop) Web Access, 513
RDS (Remote Desktop Services)
    CAL licenses, 514
    Connection Broker, 502
    deployment, 499–501
    deployment properties, 502–503
    licensing, 513–515
    PowerShell cmdlets, 515–517
    Session Host, 503–509, 731
    Virtualization Host, 509–512
recover.exe command-line tool, 732
Recovery Vault, Azure Backup, 668. *See also* Windows Server Backup

Recycle Bin, AD (Active Directory), 147–149
redeployment, 724–725
ReFS (Resilient File System), 264–265
Regedit.exe/Regedt32.exe, 37
reg.exe command-line tool, 732
Regional Settings control panel, 37
Registry, using with WAC (Windows Admin Center), 5
registry entries, command-line tool, 732
regsvr32.exe command-line tool, 732
relying party, AD FS, 439–440
relying party trust, AD FS, 439, 443
Remote Access Management console, 13. *See also* routing and remote access
Remote Access PowerShell cmdlets, 224–227
remote administration, 1
Remote Desktop, 1–2, 5, 13, 25–27
Remote FX, 196
RemoteApp, 506–507
RemoteFX, 512
repair-bde.exe command-line tool, 732
replication, DFS, 296–299
request filters, website management, 324–326
reset password right, delegating, 585
resource metering, Hyper-V virtualization platform, 193

Technet24

Resource Monitor, 13, 659
Resource Properties, configuring for DAC, 457–459
resource records, DNS, 166–167
restore files and directories, rights assignment, 600
restoring
  from Azure Backup, 669–670
  from backups, 665–666
  deleted items, 119–120, 146–152
  to locations, 666
revoking certificates, 408–410
RID (relative identifier) master, 130
RODC (Read Only Domain Controllers), 121–124
role-capability files, JEA, 604–607
RoleDefinitions session configuration file, 607
roles and features
  adding, 45–46
  adding to images, 54–56
  upgrading, 687–688
  using with WAC (Windows Admin Center), 5
role-specific backups, 665
Rolling Upgrades, clusters, 379–380. *See also* upgrade and migration paths
root certificates, deploying, 405
routing and remote access, 493–496. *See also* Remote

Access Management console; VPNs (virtual private networks)
DirectAccess, 489–496
LAN routing, 487
NAT (Network Address Translation), 487–488
PowerShell cmdlets, 496–498
RD (Remote Desktop) Gateway, 475–479
Routing and Remote Access console, 13
RSAT (Remote Server Administration Tools) tools, 6, 11–17, 61
RunAsVirtualAccount session configuration file, 608
RunAsVirtualAccountGroups session configuration file, 608

S

S2D, Storage Spaces Direct, 240–246
SAC (Semi Annual Channel), 33–34
sandbox, using with containers, 338
scheduled tasks
blocking, 592
command-line tool, 729, 732
using with WAC (Windows Admin Center), 729–723
scheduling priority, rights assignment, 599
schema master, 129
schtasks.exe command-line tool, 732

Technet24

scripts, command-line tool, 729
ScriptsToProcess role-capability file, 606
SCT (Security Compliance Toolkit)
    Attack Surface Analyzer, 638–639
    Credential Guard, 640–642
    Policy Analyzer tool, 636–638
searching files, 730, 733
secedit.exe command-line tool, 732
secrets, isolating, 640–642
security. *See also* connection security rules; ESAE (Enhanced Security Administrative Environment)
    of backups, 663
    connection security rules, 623
    Controlled Folder Access, 645–647
    Exploit Protection, 647–649
    virtualization-based, 644–645
security audits, rights assignment, 598
security classifications, updating, 672
security configuration, command-line tool, 732
security options, accounts, 586–587
security roles, WSUS, 674–675
security settings, Exploit Protection, 647–649
security vulnerabilities, locating, 638–639
server configuration, command-line tool, 732

Server Core
App Compatibility Features on Demand, 42–43
base image, 346
containers, 338
DCs (domain controllers), 120–121
deployment, 43–44
interface, 36–37
and migration, 691
overview, 34–35
roles, 37–42
Server for NFS (Network File System), 259–260
Server Manager console, 14–17, 277–279
server operators, blocking from scheduling tasks, 592
Server with Desktop Experience, 44–45
servermanagercmd.exe command-line tool, 732
service accounts
containers, 355–356
hardening, 600–603
types, 136
service dependencies, 726–727
Service Principal Name, command-line tool, 731
Services, using with WAC (Windows Admin Center), 5
Services console, 13
Services for NFS (Network File System) console, 13

Technet24

servicing branches, 33–35
session activity, viewing, 746
session-configuration files, JEA, 607–608
SessionType configuration file, 607
sfc.exe command-line tool, 732
SHA hash algorithm, CAs, 399–400
shared folder cmdlets, 302–303
shared objects, rights assignment, 598
sharing folders, 258–259, 274–279
ShellRunAs tool, 745
shielded VMs. *See also* VMs (Virtual Machines)
    vs. encryption-supported VMs, 633
    guarded fabric, 631–634
    Hyper-V virtualization platform, 223
    overview, 628–630
shielding data file, 13, 634
showmount.exe command-line tool, 732
shutdown, rights assignment, 598, 600
SigCheck tool, 739–740
signed-on user identity, command-line tool, 733
sites. *See* website management; websites
SLC (server licensor certificate), 465–466
smart card security option, 587
smart paging, Hyper-V virtualization platform, 192–193

SmartScreen, Windows Defender, 651
SMB (Server Message Block) 3.1.1, 251–252
snapshots, AD (Active Directory), 151–152
snippets, PowerShell, 23–24
SoFS (Scale-Out File Server), 258–259
solutions, applying, 728
SQL Server
  base image, 346
  workgroup clusters, 381
SSD (Storage Spaces Direct), 240–246
SSDs (solid-state disks), 204
SSH client and server, 27–29
SSTP VPN protocol, 481
staging, DAC, 463
Standalone CAs, 403–406
Standard edition, 31–32, 687
storage
  adding to storage pools, 230–233
  using with WAC (Windows Admin Center), 5
storage devices, iSCSI, 252–255
storage migration, Hyper-V virtualization platform, 221
Storage Migration Service, 5, 710–718. *See also* file and storage servers
Storage module, PowerShell cmdlets, 268–270

Technet24

storage optimization, Hyper-V virtualization platform, 204–205
storage pools
  features, 229–233
  thin provisioning and trim, 237–239
Storage QoS (Quality of Service), 204, 263–264
Storage Replica
  configuration replication, 249–251
  configurations, 248–249
  features, 247–248
  PowerShell cmdlets, 271
  using with WAC (Windows Admin Center), 5
Storage Reports, FSRM, 286–287
storage servers, WSMT (Windows Server Migration Tools), 718–722
storage spaces
  resiliency, 234–235
  and storage pools, 229
  tiering, 235–236
  virtual disks, 239–240
Storage Spaces Direct, 240–241
storage tiering, 204–205
storage-related PowerShell cmdlets, 266–271
subnets, adding in Active Directory, 116

subst.exe command-line tool, 732
support calls, 753
Swarm mode, Docker, 365–368
symbolic links
    command-line tool, 731
    rights assignment, 598
symptoms and diagnosis, documenting, 725–726
Sysinternals tools
    AccessChk, 740
    AccessEnum, 744–745
    AD Explorer, 746–749
    ADInsight, 749
    LogonSessions, 746
    overview, 733
    ProcDump, 736
    Process Explorer, 734–735
    Process Monitor, 735–736
    PsPing, 750
    PsTools Suite, 737–738
    RAMMap, 751–753
    ShellRunAs, 745
    SigCheck, 739–740
    Sysmon, 741–744
    VMMap, 738–739

Technet24

Sysmon tool, 741–744
Sysprep.exe utility, 57
system activity, logging to event logs, 741–744
System Configuration console, 13
system date, command-line tool, 729
system files, blocking extensions, 283. *See also* files; protected files
system information
  command-line tool, 732
  console, 14
System Insights, using with WAC (Windows Admin Center), 5
system performance, rights assignment, 599
system time, rights assignment, 597
systeminfo.exe command-line tool, 732
SYSVOL replication
  migrating, 730
  upgrading, 695–696

T

takeown.exe command-line tool, 732
Task Manager, 37
Task Scheduler console, 14
task scheduling, blocking, 592

taskill.exe command-line tool, 732
tasklist.exe command-line tool, 732
tasks, automating, 2
TCP/IP (Transmission Control Protocol/Internet Protocol). *See also* NetBIOS over TCP/IP command-line tool
command-line tool, 730
networks, 252–255
subnets, 115
Template Disk Wizard console, 14
test certificates, IIS (Internet Information Services), 312
text files, searching, 730
thin provisioning and trim storage, 237–238
time stamp, viewing, 739–740
time zone settings, command-line tool, 732
TimeDate.cpl, 37
time.exe command-line tool, 732
TLS certificates
configuring, 315–317
DirectAccess, 491
token object, rights assignment, 597
token-signing certificates, AD FS, 442
TPDs (Trusted Publishing Domains), AD RMS, 471
TPM (Trusted Platform Module), 223, 631–633

Technet24

traceroute.exe, command-line tool, 731
tracerpt.exe command-line tool, 732
tracert.exe command-line tool, 732
traffic, distribution, 385–389
TranscriptDirectory session configuration file, 608
transparent networks, 360–362
traverse checking, rights assignment, 597
tree.exe command-line tool, 732
trim storage, 238
troubleshooting methodology. *See also* command-line tools
  applying solutions, 728
  dependencies, 726–727
  overview, 723–724
  ranking hypothetical solutions, 727–728
  redeployment, 724–725
  symptoms and diagnosis, 725–726
trustlets, 640
TTL (Time to Live) value, DNS, 170–171
TUDs (Trusted User Domains), AD RMS, 470–471
tunnel rules, firewalls, 627
TypesToProcess role-capability file, 607
tzutil.exe command-line tool, 732

U

universal groups, 133
"unknown records," 167
unoptimization, 262
updates. *See also* Azure Update Management; WSUS (Windows Server Update Services)
  applying, 356–357
  and drivers added to images, 53–54
  using with WAC (Windows Admin Center), 5
upgrade and migration paths. *See also* migration; Rolling Upgrades
  AD CS (Active Directory Certificate Services), 699–704
  DHCP, 705–709
  DNS, 704–705
  editions, 689
  evaluation to licensed version, 688–689
  overview, 685–688
  support, 685–693
upgrading
  DCs (domain controllers), 693–695
  editions, 689
  roles and features, 687–688
URL authorization rules, 323–324

Technet24

user accounts, 130–132
user and device claims, configuring for DAC, 456–457. *See also* claims and claim rules
user identity, getting information about, 733
user rights, assignment policies, 596–600

V

V VHD Sets, Hyper-V virtualization platform, 219
VariableDefinitions role-capability file, 606
verifier.exe command-line tool, 732
virtual account, 602
virtual directories, IIS, 313
virtual disks, creating, 239–240
Virtual Fibre Channel adapters, Hyper-V virtualization platform, 203
virtual hard disks, Hyper-V virtualization platform, 199–202
virtual switches, Hyper-V virtualization platform, 205–206
Virtual Switches, using with WAC (Windows Admin Center), 5
Virtualization Host, RDS, 509–512
virtualization-based security, 640–642, 644–645
virtualized domain controllers, 121, 124
virtualized workloads, 47–48

VisibleAliases role-capability file, 606
VisibleCmdlets role-capability file, 606
VisibleExternalCommands role-capability file, 606
VisibleFunctions role-capability file, 606
VisibleProviders role-capability file, 606
visual debugging, PowerShell, 21
VM MAC addresses, Hyper-V virtualization platform, 208
VM network adapters, Hyper-V virtualization platform, 206
VM Network Health Detection, 222
VMM (Virtual Machine Manager)
  adding WDS, 75–79
  host groups, 77
  multitier VM service, 80
  networking, 71–75
  storage, 69–71
  templates, 68–69
VMMap tool, 738–739
VMs (Virtual Machines). *See also* encryption-supported VMs; Iaas (Azure Infrastructure as a Service) VMs; JIT (Just In Time) VM Access; shielded VMs
  copying with Hyper-V, 221–222
  drain on shutdown, 222

Technet24

exporting with Hyper-V, 221–222
failover clustering, 377–379
Generation 2, 194–195
importing with Hyper-V, 221–222
load balancing, 378–379
MAC addresses, 208
Network adapters, 206
NIC teaming, 208
resizing, 536
WAC (Windows Admin Center), 5
Volume Activation Tools console, 14
volumes
command-line tools, 730–732
formatting and managing, 730
maintenance task rights, 599
recovering readable information, 732
VPNs (virtual private networks). *See also* routing and remote access
authentication, 482
certificates, 479
default conditions and settings, 485
disabling protocols, 483
IKEv2 (Internet Key Exchange version 2) VPN protocol, 479–480

L2TP/IPsec protocols, 481
PPTP VPN protocol, 481–482
server access, 483–486
server deployment, 482–483
SSTP VPN protocol, 481
VSCode (Visual Studio Code), 23
VSS (Volume Shadow Copy Services), 670–671
Vssadmin, 670–671

W

WAC (Windows Admin Center)
AD DS feature, 120–121
Azure hybrid cloud services, 565–567
Azure Iaas VMs, 567–569
extensions, 9–10
features, 4–6
installing, 6–8
shared folder permissions, 276–277
show script, 10
tools and functionality, 5
WDAC (Windows Defender Application Control), 642–644
WDS (Windows Deployment Services)
adding to VMM, 75–76

Technet24

boot options for images, 64–65
configuring, 62–66
configuring for images, 62–66
driver groups and packages, 68
managing images, 62
requirements, 60–61
transmissions, 67–68
using, 59–60
wdsutil.exe command-line tool, 733
Web Access, RDS (Remote Desktop Services), 513
Web Application Proxy, AD FS, 437, 445–447
web applications, adding, 314–315
webadmin.exe command-line tool, 733
website management. *See also* IIS (Internet Information Services)
adding sites, 310–313
default document, 320–321
directory browsing, 321–322
domain name filtering, 322–323
error response, 319–320
IP address restrictions, 322–323
modifying settings, 314
overview, 309–310
request filters, 324–326

site authentication, 318–319
test certificates, 312
TLS certificates, 315–317
URL authorization rules, 323–324
virtual directories, 313
websites, creating, 115
wecutil.exe command-line tool, 733
wevutil.exe command-line tool, 733
WFAS (Windows Firewall with Advanced Security)
connection security rules, 623–628
firewall profiles, 618–619
inbound rules, 619–620
IPsec configuration, 621–622
outbound rules, 620–621
overview, 618
wget command, PowerShell, 44
where.exe command-line tool, 733
whoami.exe command-line tool, 733
WIM (Windows Imaging) file format, 48
.wim files, 48
Windows 10 for Server 2019, 11
Windows clipboard, redirecting output to, 729
Windows Defender, 2, 650–651
Windows Deployment Services, 14, 733

Technet24

Windows Firewall with Advanced Security console, 14
Windows Installer files, command-line tool, 731
Windows Memory Diagnostics console, 14
Windows Script Host, command-line tool, 733
Windows Server
  editions, 31–33
  servicing branches, 33–35
Windows Server 2003, end of support, 686
Windows Server 2008 mode, 295
Windows Server Backup, 14, 662–666, 683–684. *See also* Azure Backup; backup; Recovery Vault
Windows Server Update Services console, 14
Windows SIM, 58
WinPE (Windows Preinstallation Environment), command-line tool, 731
winrs.exe command-line tool, 733
WINS console, 14
wmic.exe command-line tool, 733
workgroup clusters, 381
Workplace Join, 447–449
writing files, blocking, 282
wscript.exe command-line tool, 733
WSL (Windows Subsystem for Linux)
  installing, 576–579

overview, 575–576
updating, 579
version 2.0, 579–580
WSLab project, 79–80
WSMT (Windows Server Migration Tools)
DHCP server, 705–709
file servers, 718–722
storage servers, 718–722
using, 689–693
WSUS (Windows Server Update Services). *See also* updates
automatic approval rules, 679
autonomous and replica modes, 673
groups, 675
languages, 672
policies, 675–676
PowerShell cmdlets, 684
products, 672
security classifications, 672
security roles, 674–675
update deployment, 677–679
update files, 673–674

Z

Technet24

zone aging and scavenging, 167–168

Plug into learning at

MicrosoftPressStore.com

The Microsoft Press Store by Pearson offers:

- Free U.S. shipping
- Buy an eBook, get three formats – Includes PDF, EPUB, and MOBI to use with your computer, tablet, and mobile devices
- Print & eBook Best Value Packs
- eBook Deal of the Week – Save up to 50% on featured title
- Newsletter – Be the first to hear about new releases, announcements, special offers, and more
- Register your book – Find companion files, errata, and product updates, plus receive a special coupon* to save on your next purchase

Discounts are applied to the list price of a product. Some products are not eligible to receive additional discounts, so your discount code may not be applied to all items in your cart. Discount codes cannot be applied to products that are already discounted, such as eBook Deal of the Week, eBooks that are part of a book + eBook pack, and products with special discounts applied as part of a promotional offering. Only one coupon can be used per order.

Pearson

Technet24

Hear about it first.

Since 1984, Microsoft Press has helped IT professionals, developers, and home office users advance their technical skills and knowledge with books and learning resources.

Sign up today to deliver exclusive offers directly to your inbox.

- New products and announcements
- Free sample chapters
- Special promotions and discounts
- ... and more!

MicrosoftPressStore.com/newsletters

Pearson

Code Snippets

Many titles include programming code or configuration examples. To optimize the presentation of these elements, view the eBook in single-column, landscape mode and adjust the font size to the smallest setting. In addition to presenting code and configurations in the reflowable text format, we have included images of the code that mimic the presentation found in the print book; therefore, where the reflowable format may compromise the presentation of the code listing, you will see a “Click here to view code image” link. Click the link to view the print-fidelity code image. To return to the previous page viewed, click the Back button on your device or app.

```
msiexec /i <WindowsAdminCenterInstallerName>.msi /qn /L*v log.txt SME_PORT=<port>
SME_THUMBPRINT=<thumbprint> SSL_CERTIFICATE_OPTION=installed
```

```
Install-WindowsFeature -IncludeAllSubFeature RSAT
```

```
$cred = Get-Credential
Enter-PSSession -computername <computername> -Credential $cred
```

```
Set-Item wsman:\localhost\Client\TrustedHosts -Value 192.168.3.200 -Concatenate
```

```
Invoke-Command -ScriptBlock { get-service wuauserv } -computername $Computers
```

```
$Computers = Get-Content c:\Computers.txt
Invoke-Command -FilePath c:\FixStuff.ps1
```

```
Set-ItemProperty -Path "HKLM:\System\CurrentControlSet\Control\Terminal Server" -Name
"fDenyTSConnections" -Value 0
```

```
Enable-NetFirewallRule -DisplayGroup "Remote Desktop"

Add-WindowsCapability -Online -Name OpenSSH.Server~~~~0.0.1.0
Add-WindowsCapability -Online -Name OpenSSH.Client~~~~0.0.1.0

Install-Module -Force OpenSSHUtils -Scope AllUsers

Set-Service -Name ssh-agent -StartupType 'Automatic'
Set-Service -Name sshd -StartupType 'Automatic'

Set-Service ssh-agent -StartupType 'Automatic'
Start-service ssh-agent
Ssh-add ~\.ssh\id_rsa

Ssh chancellor@172.16.0.15 mkdir c:\users\chancellor\.ssh\
Scp c:\users\chancellor\.ssh\id_rsa.pub chancellor@172.16.0.15:C:\Users\Administrator\.
ssh\authorized_keys

Repair-AuthorizedKeyPermission C:\users\Chancellor\.ssh\authorized_keys

# Match Group administrators
#       AuthorizedKeysFile __PROGRAMDATA__/ssh/administrators_authorized_keys
```

```
Add-WindowsCapability -Online -Name ServerCore.AppCompatibility~~~~0.0.1.0

Mount-DiskImage -ImagePath X:\ISO_Folder\ISO_filename.iso

Add-WindowsCapability -Online -Name ServerCore.AppCompatibility~~~~0.0.1.0 -Source
<Mounted_Server_FOD_Drive> -LimitAccess

$Computers = Get-Content c:\virt-servers.txt
$Invoke-Command -ScriptBlock { Add-WindowsFeature Hyper-V -Restart } -Computername $Computers
```

```
Dism.exe /get-wiminfo /wimfile:c:\images\install.wim

Get-WindowsImage -ImagePath c:\images\install.wim

Dism.exe /mount-image /imagefile:c:\images\install.wim /index:2 /mountdir:c:\mount

Mount-WindowsImage -ImagePath c:\images\install.wim -index 2 -path c:\mount

Dism.exe /image:c:\mount /Add-Driver /driver:c:\drivers\ /recurse

Add-WindowsDriver -Path c:\mount -Driver c:\drivers -Recurse

Dism.exe /image:c:\mount /Add-Package /PackagePath:"c:\updates\windows10.0-kb4505056-x64
_09c6216e684ce667fc0c07cc30e84ef21f04c6f1.msu"

Add-WindowsPackage -path c:\mount -packagepath "c:\updates\AMD64-all
-windows10.0-kb4505056-x64_09c6216e684ce667fc0c07cc30e84ef21f04c6f1.msu"

Dism.exe /image:c:\mount /Enable-Feature /all /FeatureName:NetFx3ServerFeatures

Dism.exe /image:c:\mount /Enable-Feature /all /FeatureName:NetFx3ServerFeatures

Dism.exe /image:c:\mount /Enable-Feature /all /FeatureName:NetFx3 /Source:d:\sources\sxs

Dism.exe /Unmount-Wim /MountDir:c:\mount /commit

Dismount-WimImage -Path c:\mount -Save

Choco install -y windows-adk-all

Configuration TestDscConfiguration {
Node “TEST-SVR” {
WindowsFeature InstallWINS {
Ensure = "Present"
Name = "WINS"
}
WindowsFeature InstallBackup {
Ensure = "Present"
Name = "Windows-Server-Backup"
}
}
}
TestDscConfiguration

Sudo dpkg -i ./chef-server-core_12.19.31-1_amd64.deb
```

```
Sudo chef-server-ctl reconfigure

Sudo chef-server-ctl user-create sydneymelbourne Sydney Melbourne sydney@company.
pri P@ssw0rd --filename ~/certs/sydneymelbourne.pem

Sudo chef-server-ctl org-create contoso "Contoso-Chef" --association_user
sydneymelbourne --filename ~/certs/contoso.pem

Sudo chef-server-ctl install chef-manage

Sudo check-server-ctl reconfigure
Sudo chef-manage-ctl reconfigure --accept-license

Sudo chef-server-ctl install opscode-reporting
Sudo chef-server-ctl reconfigure
Sudo opscode-reporting-ctl reconfigure --accept-license

Knife cookbook site download windows
Knife cookbook site download ohai

Chef generate cookbook wins-windows

windows_feature "WINSRuntime" do
action :install
end

name "wins-windows"
description "test WINS install"
run_list "recipe[wins-windows]"

Knife cookbook upload --all

Gem update knife-windows

Tar -xvf puppet-enterprise*

Puppet parser validate c:\test_manifests\services.pp

Puppet apply c:\test_manifests\services.pp --noop
```

Puppet module install puppetlabs/windows

```
Choco install -y windows-adk-all

Set-ExecutionPolicy Unrestricted

iwr https://chocolatey.org/install.ps1 -UseBasicParsing | iex
```

```
Install-WindowsFeature AD-Domain-Services -IncludeManagementTools
Install-ADDSDomainController -DomainName contoso.internal -InstallDNS:$True
-credential (Get-Credential)
```

```
Install-ADDSForest -DomainName contoso.internal -InstallDNS
```

```
Set-ADObject "CN=NTDS Settings,CN=MEL-DC1,CN=Servers,CN=Melbourne-Site,CN=Sites,
CN=Configuration,DC=Contoso,DC=Internal" -Replace @{options='1'}
```

```
New-ADDCCloneConfigFile -IPv4Address 10.10.10.42 -IPv4DefaultGateway 10.10.10.1
-IPv4SubnetMask 255.255.255.0 -IPv4DNSResolver 10.10.10.10 -Static -SiteName MEL-SITE
```

```
Get-ADForest contoso.internal | FT SchemaMaster
```

```
Get-ADForest contoso.internal | FT DomainNamingMaster
```

```
Get-ADDomain contoso.internal | FT PDCEmulator
```

```
Get-ADDomain contoso.internal | FT InfrastructureMaster
```

```
Get-ADDomain contoso.internal | FT RidMaster
```

```
Move-ADDirectoryServerOperationMasterRole -Identity MEL-DC2 -OperationMasterRole
RIDMaster,InfrastructureMaster,DomainNamingMaster -Force
```

```
Test-ComputerSecureChannel -Credential Domain\<AdminAccount> -Repair
```

```
Add-KdsRootKey -EffectiveTime ((get-date).addhours(-10))
```

```
New-ADServiceAccount MEL-SQL-GMSA -DNSHostname MEL-SQL-GMSA.contso.internal
-PrincipalsAllowedToRetrieveManagedPassword MEL-SQL-Servers
```

```
Install-ADServiceAccount MEL-SQL-GMSA
```

```
Install-WindowsFeature RSAT-ADDS-Tools
```

```
Dsquery * "cn=Directory Service,cn=Windows NT,cn=Services,cn=Configuration,dc=Contoso,
dc=Internal" -scope base -attr tombstonelifetime
```

```
Ntdsutil "authoritative restore" "restore object cn=Mercury,ou=Planets,dc=contoso,
dc=com" q q
```

```
Ntdsutil "authoritative restore" "restore subtree OU=Planets,
DC=contoso,DC=com" q q
```

```
Dsamain.exe /dbpath C:\restore\ntds.dit /ldapport 51389
```

Technet24

```
Ntdsutil “Activate Instance NTDS” snapshot create quit quit
Ntdsutil snapshot “list all” quit quit
Ntdsutil snapshot “list all” “mount 1” quit quit
```

```
Dsamain.exe /dbpath C:\$SNAP_201606100306_VOLUMEC$\WINDOWS\NTDS\ntds.dit /ldapport 51389
```

```
Add-DNSServerDirectoryPartition -Name Tasmania

Add-DnsServerPrimaryZone -Name cpandl.com -ReplicationScope Forest

Add-DnsServerSecondaryZone -Name "australia.contoso.com" -ZoneFile "australia.contoso.
com.dns" -MasterServers 192.168.15.100

Add-DnsServerPrimaryZone -NetworkID "192.168.15.0/24" -ReplicationScope "Forest"

Add-DnsServerZoneDelegation -Name "contoso.com" -ChildZoneName "australia" -NameServer
"ausdns.australia.contoso.com" -IPAddress 192.168.15.100 -PassThru
-Verbose

Add-DnsServerForwarder 10.10.10.111

Add-DnsServerConditionalForwarderZone -MasterServers 10.10.10.102 -Name tailspintoys.com
-ReplicationScope Forest

Add-DnsServerStubZone -MasterServers 10.10.10.222 -Name fabrikam.com -ReplicationScope
Forest -LoadExisting

Add-DnsServerPrimaryZone -Name GlobalNames -ReplicationScope Fores

Set-DnsServerGlobalNameZone -ComputerName DNSServerName -Enable $True

Set-DnsServerScavenging -ApplyOnAllZones -RefreshInterval 10.0:0:0 -ScavengingInterval
10.0:0:0 -ScavengingState $True

Dnscmd /config /socketpoolsize 4000

Set-DNSServerCache -LockingPercent 80

Set-DNSServerRecursion -RetryInterval 3 -PassThru

Add-DHCPServerv4Scope -Name "MelbourneNet" -StartRange 192.168.10.1 -Endrange
192.168.10.254 -SubnetMask 255.255.255.0

Set-DhcpServerv4OptionValue -Computername dhcp.contoso.com -ScopeID 192.168.10.0
-DNSServer 192.168.0.100

Add-DhcpServerv4Reservation -ScopeId 192.168.10.0 -IPAddress 192.168.10.8 -ClientId
A0-FE-C1-7A-00-E5 -Description "A4 Printer"

Add-DhcpServerv4Filter -List Allow -MacAddress D1-F2-7C-00-5E-F0 Description
"Rooslan's Laptop"
```

```
Add-DhcpServerv4Superscope -SuperscopeName "MelbourneCBD" -ScopeId 192.168.10.0,
192.168.11.0
Set-DhcpServerv4DnsSetting -Computer MEL-DC -NameProtection $true
```

```
Set-VM -Name VMName -AutomaticStopAction TurnOff

Set-VMProcessor -VMName NameOfVM -ExposeVirtualiationExtensions $true

Get-VMNetworkAdapter -VMName NameOfVM | Set-VMNEtworkAdapter -MacAddressSpoofing On

New-VMSwitch -name VMNAT -SwitchTypeInternal

New-NetNAT -Name LocalNAT -InternalIPInterfaceAddressPrefix “192.168.15.0/24”

Get-NetAdapter “vEthernet (VmNat)” | New-NetIPAddress -IPAddress 192.168.15.1
-AddressFamily IPv4 -PrefixLength 24

Enter-PSSession -vmname NameOfVM

Get-Disk 3 | Add-VMHardDiskDrive -VMName Alpha-Test
```

```
Get-StoragePool -FriendlyName ExamplePool | Get-PhysicalDisk | Select UniqueID, MediaType, Usage

Install-WindowsFeature -Name "Failover-Clustering", "Hyper-V", "Data-Center-Bridging", "RSAT-Clustering-PowerShell", "Hyper-V-PowerShell", "FS-FileServer"

New-Cluster -Name TestS2DCluster -Node S2DN1,S2DN2 -NoStorage

Enable-ClusterStorageSpacesDirect -CimSession TestS2DCluster

Test-Cluster -Name TestS2DCluster -node S2DN1,S2DN2,S2DN3 -Include "Storage Spaces Direct", "Inventory", "Network", "System Configuration"

Test-SRTopology -SourceComputerName  S3-DCenter -SourceVolumeName e: -SourceLogVolumeName f: -DestinationComputerName S4-DCenter -DestinationVolumeName e: -DestinationLogVolumeName f: -DurationInMinutes 15 -ResultPath c:\temp

New-SRPartnerShip -SourceRGName SR71 -DestinationRGName SR72 -SourceComputerName S3-DCenter -SourceVolumeName e: -SourceLogVolumeName f: -DestinationComputerName S4-DCenter -DestinationVolumeName e: -DestinationLogVolumeName f:

Set-SRPartnerShip -NewSourceComputerName S4-DCenter -SourceRGName SR72 -DestinationComputerName S3-DCenter -DestinationRGName SR71

Set-SmbServerConfiguration -AuditSmb1Access $true

Remove-WindowsFeature -Name FS-SMB1

New-IscsiVirtualDisk -Path "e:\disks\disk1.vhd" -Size 10GB

New-IscsiServerTarget -Targetname "Syd-A-Syd-B-Target" -InitiatorIDs DNSName:Syd-a.contoso.com,DNSName:syd-b.contoso.com

Set-WmiInstance -Namespace root\wmi -Class WT_iSNSServer -Arguments @{ServerName="ISNSservername"}

Add-ClusterScaleOutFileServerRole -Cluster ClusterName -Infrastructure -Name InfrastructureSoFSName

New-StorageQosPolicy -Name Alpha -PolicyType Dedicated -MinimumIops 150 -MaximumIops 250

New-Volume -FriendlyName "ExampleVolume" -filesystem CSVFS_ReFS -StoragePoolFriendlyName "ExamplePool" -StorageTierFriendlyNames Performance, Capacity -StorateTierSizes 200GB, 800GB
```

```
Install-WindowsFeature FS-DFS-Namespace

Install-WindowsFeature FS-DFS-Replication

Install-WindowsFeature FS-BranchCache

Install-WindowsFeature BranchCache
Start-Service BranchCache
Enable-BCHostedServer
```

```
Appcmd.exe add site /name:contoso /physicalPath:c:\contoso /bindings:http/*:80:www.
contoso.internal
```

```
Docker run -it --isolation=hyperv mcr.microsoft.com/windows/servercore cmd

Install-Module -Name DockerMsftProvider -Repository PSGallery -Force

Install-Package -Name docker -ProviderName DockerMsftProvider

Install-Package -Name docker -ProviderName DockerMsftProvider
```

```
{
    "authorization-plugins": [],
    "dns": [],
    "dns-opts": [],
    "dns-search": [],
    "exec-opts": [],
    "storage-driver": "",
    "storage-opts": [],
    "labels": [],
    "log-driver": "",
    "mtu": 0,
    "pidfile": "",

    "cluster-store": "",
    "cluster-advertise": "",
    "debug": true,
    "hosts": [],
    "log-level": "",
    "tlsverify": true,
    "tlscacert": "",
    "tlscert": "",
    "tlskey": "",
    "group": "",
    "default-ulimits": {},
    "bridge": "",
    "fixed-cidr": "",
    "raw-logs": false,
    "registry-mirrors": [],
    "insecure-registries": [],
    "disable-legacy-registry": false
}
```

```
docker pull mcr.microsoft.com/windows/servercore:ltsc2019

docker pull mcr.microsoft.com/windows/nanoserver:1909

docker pull -a mcr.microsoftcom/windows/servercore

Docker pull -a mcr.microsoft.com/windows/nanoserver

docker pull mcr.microsoft.com/windows/nanoserver:1909

docker pull mcr.microsoft.com/windows/servercore:ltsc2019
docker pull mcr.microsoft.com/windows/servercore:1909

docker pull mcr.microsoft.com/windows/iotcore:1809

docker pull mcr.microsoft.com/windows/servercore/iis:windowsservercore-ltsc2019
docker pull mcr.microsoft.com/windows/servercore/iis:windowsservercore-1909

docker pull mcr.microsoft.com/mssql/server:2019-latest

docker pull mcr.microsoft.com/azure-cli

docker pull mcr.microsoft.com/dotnet/core/aspnet:3.1

docker pull mcr.microsoft.com/powershell

Docker run -it mcr.microsoft.com/windows/servercore:ltsc2019 powershell.exe

Docker run -d -p 8080:80 mcr.microsoft.com/windows/servercore/
iis:windowsservercore-ltsc2019

wget -Uri "http://javadl.sun.com/webapps/download/AutoDL?BundleId=107944" -outfile
javaInstall.exe -UseBasicParsing
REG ADD HKLM\Software\Policies\Microsoft\Windows\Installer /v DisableRollback /t REG_
DWORD /d 1 | Out-Null ./javaInstall.exe /s INSTALLDIR=C:\Java REBOOT=Disable | Out-Null

FROM mcr.microsoftcom/windows/servercore:ltsc2019
RUN mkdir ExampleDirectory
RUN dism.exe /online /enable-feature /all /featurename:iis-webserver /NoRestart

New-CredentialSpace -Name MelbourneAlpha -AccountName MelbourneAlpha

docker run --security-opt "credentialspec=file://twt_webapp01.json" --hostname
webapp01 -it mcr.microsoft.com/windows/servercore:ltsc2019 powershell

Docker network inspect <network name>

Get-ContainerNetwork | Remove-ContainerNetwork
```

```
Docker network create -d nat --subnet=192.168.15.0/24 --gateway=192.168.15.1
CustomNat

Docker run -it --network=CustomNat <ContainerImage> <cmd>

Docker run -it -p 8080:80 microsoft/windowsservercore powershell.exe

Docker run -itd -p 80 mcr.microsoft.com/windows/servercore/iis:windowsservercore-
ltsc2019 powershell.exe

docker network create -d transparent TransparentNetworkName

Docker network create -d transparent --subnet=192.168.30.0/24 --gateway=192.168.30.1
TransNet

Docker run -it --network:TransNet --ip 192.168.30.101 microsoft/windowsservercore cmd.exe

docker network create --driver=overlay -subnet=192.168.50.0/24 OverlayNet

Docker network create -d l2bridge -subnet=192.168.88.0/24 --gateway=192.168.88.1
L2BridgeNet

{
    "experimental":true
}

[Environment]::SetEnvironmentVariable("LCOW_SUPPORTED", "1", "Machine")

Restart-service *docker*

docker run --platform linux busybox uname -a

Docker swarm init --advertise-addr <container_host_ip> --listen-addr
<container_host_ip>:2377

Docker swarm join-token worker -q

Docker swarm join-token manager -q

Docker swarm join --token <worker-join-token> <manager_container_host_ip>

Docker swarm join --token <manager-join-token> <manager_container_host_ip>

Docker network create --driver=overlay --subnet=10.0.10.0/24 <OverlayNetworkName>

Docker service create -name=SWARMFRONT --endpoint-mode=dnsrr --network=OverNet
WEBFRONTEND

Docker service scale SWARMFRONT=10

Docker service ps <servicename>
```

```
Install-WindowsFeature -ComputerName Node1 -Name MultiPath-IO

(Get-Cluster).AutoBalancerLevel = <value>

(Get-Cluster).AutoBalancerMode = <value>

Get-Cluster | Select ClusterFunctionalLevel

new-itemproperty -path HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System
-Name LocalAccountTokenFilterPolicy -Value 1

New-Cluster -name workgrpclst -node node1,node2 -AdministrativeAccessPoint DNS
```

```
Certutil.exe -dspublish -f SA_ROOT-CA.crt
```

```
Certutil -getkey <certificate_serial_number> outputblob
```

```
Certutil -recoverkey outputblob recoveredkey.pfx
```

```
[Version]
Signature= "$Windows NT$"
[Certsrv_Server]
RenewalKeyLength=4096
RenewalValidityPeriod=Years
RenewalValidityPeriodUnits=10
LoadDefaultTemplates=0
```

```
[Version]
Signature= "$Windows NT$"
[PolicyStatementExtension]
Policies = LegalPolicy, LimitedUsePolicy
[LegalPolicy]
OID = 1.1.1.1.1.1.1.1.1
URL = "http://www.contoso.com/Policy/LegalPolicy.asp"
URL = "ftp://ftp.contoso.com/Policy/LegalPolicy.txt"
[LimitedUsePolicy]
OID = 2.2.2.2.2.2.2.2.2
URL = "http://www.contoso.com/Policy/LimitedUsePolicy.asp"
URL = "ftp://ftp.contoso.com/Policy/LimitedUsePolicy.txt"
LoadDefaultTemplates=0
```

```
Initialize-ADFSDeviceRegistration
Enable-ADFSDeviceRegistration
```

```
Install-WindowsFeature -IncludeAllSubfeature RemoteAccess
Dnscmd /config /globalqueryblocklist wpad
```

```
adminPassword=&&Tailwind&&Tailwind

az vm create --resource-group 2019IO-RG --name 20190IO-C --image win2019datacenter
--vnet-name VNet1 --subnet Alpha --admin-username prime --admin-password $adminPassword

Enable-AzVMPSRemoting -Name 2019-IO-A -ResourceGroupName 2019-IO-RG -Protocol
https -OSType Windows

Enable-AzVMPSRemoting -Name 2019-IO-A -ResourceGroupName 2019-IO-RG -Protocol
https -OSType Windows

Enter-AzVM -name 2019-IO-A -ResourceGroupName 2019-IO-RG -Credential $cred

az backup vault create --resource-group VMResourceGroup --name
IaaSVMRecoveryServicesVault --location southeasteasia

az backup protection enable-for-vm --resource-group VMResourceGroup --vault-name
IaaSVMRecoveryServicesVault --vm NameOfVM --policy-name DefaultPolicy

az backup protection backup-now  --resource-group VMResourceGroup --vault-name
IaaSVMRecoveryServicesVault
--container-name NameOfVM --item-name NameOfVM --retain-until 12-12-2020

az keyvault create --name "KeyVaultName" --resource-group "VMResourceGroup"
--location AzureRegion --enabled-for-disk-encryption

az vm encryption enable -g VMResourceGroup --name VMName --disk-encryption-
keyvault KeyVaultName

Get-ADUser -Filter {UserPrincipalName -like "*@epistemicus.internal"} -SearchBase
"DC=epistemicus,DC=internal" |
ForEach-Object {
$UPN =
$_.UserPrincipalName.Replace("epistemicus.internal","epistemicus.onmicrosoft.com")
Set-ADUser $_ -UserPrincipalName $UPN
}

Start-ADSyncSyncCycle -PolicyType Delta

Start-ADSyncSyncCycle -PolicyType Initial
```

```
Enable-WindowsOptionalFeature -Online -FeatureName
Microsoft-Windows-Subsystem-Linux
```

```
Curl.exe -L -o Ubuntu.zip https://aka.ms/wsl-ubuntu-1804
```

```
Expand-Archive .\Ubuntu.zip .\Ubuntu
```

```
Enable-WindowsOptionalFeature -Online -FeatureName Microsoft-Windows-Subsystem-Linux
Enable-WindowsOptionalFeature -Online -FeatureName VirtualMachinePlatform
```

```
Wsl --set-default-version 2
```


```
Wsl --set-default-version 2
```

```
Test-ComputerSecureChannel -credential <domain>\<admin> -Repair
Add-KdsRootKey -EffectiveTime ((get-date).addhours(-10))
New-ADServiceAccount SYD-SVC1 -DNSHOSTNAME SYD-SVC1.adatum.com
Set-ADServiceAccount -Identity SYD-SVC1 -PrincipalsAllowedToRetrieveManagedPassword
SYD-SVRS
New-PSRoleCapabilityFile -Path .\DNSOps.psrc
New-PSSessionConfigurationFile -Path .\DNSOps.pssc -Full
Register-PSSessionConfiguration -Name DNSOps -Path .\DNSOps.pssc
Enter-PSSession -ComputerName MEL-DNS1 -ConfigurationName DNSOps
New-PAMUser -SourceDomain adatum.com -SourceAccountName Oksana
New-PAMGroup -SourceGroupName "CorpAdmins" -SourceDomain adatum.com
-SourceDC mel-dc1.adatum.com
Set-AdmPwdComputerSelfPermission -Identity “Adelaide”
Auditpol.exe /set /subcategory:"File System" /success:Enable /failure:Enable
Auditpol.exe /get /category:*
Auditpol.exe /get /category:"Object Access"
```

```
bcdedit /set vsmlaunchtype auto
Set-MpPreference -EnableControlledFolderAccess Enabled
Set-MpPreference -EnableControlledFolderAccess Disabled
Add-MpPreference -ControlledFolderAccessProtectedFolders “E:\fileshare”
Remove-MpPreference -ControlledFolderAccessProtectedFolders “E:\fileshare”
Add-MpPreference -ControlledFolderAccessAllowedApplications “c:\application\app.exe”
Set-ProcessMitigation -System -Enable DEP
Get-ProcessMitigation -System
Get-ProcessMitigation -RegistryConfigFilePath exportedconfig.xml
Set-ProcessMitigation -PolicyFilePath exportedconfig.xml
Uninstall-WindowsFeature -Name Windows-Server-Antimalware
```

```
Smigdeploy.exe /package /architecture amd64 /os WS12R2 /path <path to share>
```

```
Add-PSSnapin Microsoft.Windows.ServerManager.Migration
```

```
reg export HKLM\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration
<output file>.reg
```

```
dfsutil.exe root export <\\<DomainName>\NameSpace> <BackupFileName>
```

```
Send-SmigServerData -ComputerName MEL-FS-2K16 -SourcePath e:\fileshare -DestinationPath
f:\fileshare -Recurse -Include All -Force
```