VMware NSX: Install, Configure, Manage [V4.0]

Lecture Manual



VMware® Education Services VMware, Inc. www.vmware.com/education VMware NSX: Install, Configure, Manage [V4.0] Lecture Manual Part Number EDU-EN-NSXICM4-LECT (07-NOV-2022)

Copyright © 2022 VMware, Inc. All rights reserved. This manual and its accompanying materials are protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. VMware vSphere* vMotion*, VMware vSphere* Client*, VMware vSphere* 2015, VMware vSphere*, VMware vShield Endpoint*, VMware vCenter Server*, VMware vCenter*, VMware View*, VMware Horizon* View*, VMware Verify*, VMware vSphere® Distributed Switch™, VMware Pivotal Labs® Platform Deployment™, VMware Pivotal Labs® Navigator™, VMware NSX-T™, VMware NSX# Network Detection and Response™, VMware NSX# Manager™. VMware NSX# Gateway Firewall™. VMware NSX# Firewall for Bare Metal. VMware NSX# Firewall with Advanced Threat Prevention, VMware NSX* Firewall, VMware NSX* Edge*, VMware NSX* Distributed IDS/IPS**, VMware NSX* Distributed Firewall**, VMware NSX* Data Center Enterprise Plus. VMware NSX* Data Center, VMware NSX* Advanced Load Balancer Controller*, VMware NSX* Advanced Load Balancer*, VMware NSX* Advanced Load Balancer* - Basic Edition, VMware NSX*. VMware NSX* Professional, VMware NSX* for Remote Office Branch Office, VMware NSX* for Desktop. VMware NSX# Enterprise Plus. VMware Go™. VMware ESXi™, and VMware ACE™ are registered. trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

The training material is provided "as is," and all express or implied conditions, representations, and warranties, including any implied warranty of merchantability, fitness for a particular purpose or noninfringement, are disclaimed, even if VMware, Inc., has been advised of the possibility of such claims. This material is designed to be used for reference purposes in conjunction with a training course.

The training material is not a standalone training tool. Use of the training material for self-study without class attendance is not recommended. These materials and the computer programs to which it relates are the property of, and embody trade secrets and confidential information proprietary to, VMware, Inc., and may not be reproduced, copied, disclosed, transferred, adapted or modified without the express written approval of VMware, Inc.

Contents

Modu	le 1 Course Introduction	1
1-2	Course Introduction	1
1-3	Importance	1
1-4	Learner Objectives	
1-5	Course Outline	2
1-6	Typographical Conventions	3
1-7	References	3
1-8	VMware Online Resources	4
1-9	VMware Learning Overview	5
1-10	VMware Certification Overview	6
1-11	VMware Credentials Overview	7
Modu	ule 2 VMware Virtual Cloud Network and VMware NSX	9
2-2	Importance	9
2-3	Lesson 1: VMware Virtual Cloud Network and VMware NSX	10
2-4	Learner Objectives	10
2-5	Virtual Cloud Network Framework	11
2-6	NSX Portfolio	13
2-7	Product Name Changes to VMware NSX 4.x	15
2-8	Use Cases for NSX	15
2-9	NSX Features (1)	16
2-10	NSX Features (2)	
2-11	NSX 4.0.0.1 Deprecated Features	17
2-12	High-Level Architecture of NSX	18
2-13	Management and Control Planes	19
2-14	About the NSX Management Cluster	20
2-15	NSX Management Cluster with Virtual IP Address	22

2-16	NSX Management Cluster with Load Balancer	23
2-17	About the NSX Policy	24
2-18	About NSX Manager	25
2-19	NSX Policy and NSX Manager Workflow	26
2-20	About NSX Controller	27
2-21	Control Plane Components (1)	28
2-22	Control Plane Components (2)	29
2-23	Control Plane Change Propagation	30
2-24	Control Plane Sharding Function	31
2-25	Managing Controller Failure	32
2-26	About the Data Plane	33
2-27	Data Plane Functions	33
2-28	Data Plane Components	34
2-29	Data Plane Communication Channels	35
2-30	Review of Learner Objectives	36
2-31	Key Points	36
Modu	Ile 3 Preparing the NSX Infrastructure	37
3-2	Importance	37
3-3	Module Lessons.	37
3-4	Lesson 1: Deploying the NSX Management Cluster	38
3-5	Learner Objectives	38
3-6	Implementing NSX in vSphere	38
3-7	Considerations for Deploying NSX Manager	39
3-8	NSX Manager Node Sizing	39
3-9	Deploying NSX Manager from an OVF Template	40
3-10	Methods to Access NSX Manager	41
3-11	Accessing the NSX UI	42
3-12	Accessing the NSX CLI	43
3-13	Accessing NSX Manager with API	45
3-14	Registering vCenter Server with NSX Manager	
3-15	Verifying vCenter Server Registration with NSX Manager	47
3-16	Deploying Additional NSX Manager Instances (1)	47
3-17	Deploying Additional NSX Manager Instances (2)	48
3-18	Management Cluster Status: GUI (1)	
3-19	Management Cluster Status: GUI (2)	49
3-20	Configuring the Virtual IP Address	50

3-21	Management Cluster Status: CLI (1)	51
3-22	Management Cluster Status: CLI (2)	52
3-23	Replacing Self-Signed Certificates (1)	53
3-24	Replacing Self-Signed Certificates (2)	
3-25	Lab 1: Reviewing the Lab Environment and Topologies	55
3-26	Lab 2: Reviewing the Configuration of the Predeployed NSX Manager Instance	e55
3-27	Lab 3: (Simulation) Deploying a Three-Node NSX Management Cluster	56
3-28	Review of Learner Objectives	56
3-29	Lesson 2: Navigating the NSX UI	57
3-30	Learner Objectives	
3-31	NSX Manager Policy and Manager Views	57
3-32	User Interface Preferences for Policy and Manager Modes	58
3-33	About the Networking Tab	59
3-34	About the Security Tab.	
3-35	About the Inventory Tab	60
3-36	About the Plan & Troubleshoot Tab	61
3-37	About the System Tab	62
3-38	Review of Learner Objectives	
3-39	Lesson 3: Preparing the Data Plane	63
3-40	Learner Objectives	63
3-41	Data Plane Components and Functions	64
3-42	Overview of the Transport Node	
3-43	Transport Node Components and Architecture	65
3-44	Physical Connectivity of a Transport Node	66
3-45	About IP Address Pools	67
3-46	About Transport Zones (1)	68
3-47	About Transport Zones (2)	69
3-48	Transport Node Switch Configuration	70
3-49	About VDS.	71
3-50	About N-VDS	72
3-51	Creating Transport Zones	73
3-52	Reviewing the Transport Zone Configuration	74
3-53	VDS Operational Modes	75
3-54	Physical NICs, LAGs, and Uplinks	77
3-55	About Uplink Profiles	78
3-56	Default Uplink Profiles	79

3-57	About Teaming Policies	80
3-58	Teaming Policy Modes	81
3-59	About LLDP	82
3-60	About Transport Node Profiles	83
3-61	Benefits of Transport Node Profiles	84
3-62	Prerequisites for Transport Node Profile	85
3-63	Creating a Transport Node Profile	85
3-64	Attaching a Transport Node Profile to the vSphere Cluster	86
3-65	Reviewing the ESXI Transport Node Status (1)	87
3-66	Reviewing the ESXI Transport Node Status (2)	
3-67	Verifying the ESXI Transport Node by CLI	89
3-68	Lab 4: Preparing the NSX Infrastructure	90
3-69	Review of Learner Objectives	90
3-70	Lesson 4: DPU-Based Acceleration for VMware NSX	91
3-71	Learner Objectives	91
3-72	Traditional Infrastructure Challenges	92
3-73	Next-Generation Infrastructure with DPUs	93
3-74	DPU-Based Acceleration Use Cases	94
3-75	SmartNICs Ecosystem	95
3-76	Architecture Changes with DPUs	96
3-77	SmartNIC Hardware Components	97
3-78	Network Offloading with SmartNICs	98
3-79	Supported NSX Features	99
3-80	NSX Traffic Flow with DPUs	100
3-81	Networking Configurations: SmartNIC Only	101
3-82	Networking Configurations: SmartNIC with Standard NICs	102
3-83	Installing NSX with DPUs (1)	103
3-84	Installing NSX with DPUs (2)	104
3-85	Installing NSX with DPUs (3)	105
3-86	Installing NSX with DPUs (4)	106
3-87	Installing NSX with DPUs (5)	107
3-88	Installing NSX with DPUs (6)	108
3-89	Review of Learner Objectives	109
3-90	Key Points (1)	109
3-91	Key Points (2)	109

Modu	lle 4 NSX Logical Switching	111
4-2	Importance	111
4-3	Module Lessons	111
4-4	Lesson 1: Overview of Logical Switching Architecture	112
4-5	Learner Objectives	112
4-6	Use Cases for Logical Switching	113
4-7	Prerequisites for Logical Switching	114
4-8	Logical Switching Terminology	115
4-9	About Segments (1)	116
4-10	About Segments (2)	117
4-11	About Tunneling	118
4-12	About Geneve	119
4-13	Geneve Header Format	120
4-14	Logical Switching: End-to-End Communication	122
4-15	Review of Learner Objectives	123
4-16	Lesson 2: Configuring Segments	124
4-17	Learner Objectives	124
4-18	Segment Configuration Tasks	124
4-19	Creating Segments	125
4-20	Creating Segments Workflow	126
4-21	Viewing Configured Segments	127
4-22	Attaching a VM to a Segment	128
4-23	Workflow: Attaching a VM to a Segment (1)	129
4-24	Workflow: Attaching a VM to a Segment (2)	130
4-25	Verifying the Segment Port Status	131
4-26	About Network Topology	132
4-27	Using Network Topology to Validate the Segment Configuration	133
4-28	Review of Learner Objectives	133
4-29	Lesson 3: Configuring Segment Profiles	134
4-30	Learner Objectives	134
4-31	About Segment Profiles (1)	134
4-32	About Segment Profiles (2)	135
4-33	Default Segment Profiles	
4-34	Applying Segment Profiles to Segments	136
4-35	Applying Segment Profiles to Segment Ports	
4-36	SpoofGuard Segment Profile	138

4-37	Creating a SpoofGuard Segment Profile	139
4-38	IP Discovery Segment Profile	
4-39	Creating an IP Discovery Segment Profile	141
4-40	MAC Discovery Segment Profile	142
4-41	Segment Security Profile	144
4-42	QoS Segment Profile	
4-43	Review of Learner Objectives	146
4-44	Lesson 4: Logical Switching Packet Forwarding	147
4-45	Learner Objectives	147
4-46	NSX Controller Tables	
4-47	TEP Table Update (1)	148
4-48	TEP Table Update (2)	
4-49	TEP Table Update (3)	150
4-50	TEP Table Update (4)	15
4-51	MAC Table Update (1)	152
4-52	MAC Table Update (2)	
4-53	MAC Table Update (3)	154
4-54	MAC Table Update (4)	155
4-55	About the ARP Table	156
4-56	ARP Table Update (1)	156
4-57	ARP Table Update (2)	157
4-58	ARP Table Update (3)	158
4-59	ARP Table Update (4)	159
4-60	Unicast Packet Forwarding Across Hosts (1)	160
4-61	Unicast Packet Forwarding Across Hosts (2)	16
4-62	Unicast Packet Forwarding Across Hosts (3)	162
4-63	Unicast Packet Forwarding Across Hosts (4)	163
4-64	Overview of BUM Traffic	164
4-65	Managing BUM Traffic: Head Replication	166
4-66	Managing BUM Traffic: Hierarchical Two-Tier Replication	167
4-67	Lab 5: Configuring Segments	168
4-68	Review of Learner Objectives	168
4-69	Key Points	169
Modu	ile 5 NSX Logical Routing	
5-2	Importance	17
5-3	Module Lessons	

5-4	Lesson 1: Overview of Logical Routing	172
5-5	Learner Objectives	
5-6	Use Cases for Logical Routing	172
5-7	Prerequisites for Logical Routing	173
5-8	Logical Routing in NSX	174
5-9	Tier-0 and Tier-1 Gateways	
5-10	Single-Tier Topology	177
5-11	Multitler Topology	178
5-12	Edge Nodes and Edge Clusters	179
5-13	Tier-0 Gateway Uplink Connections	180
5-14	Gateway Components: Distributed Router and Service Router (1)	181
5-15	Gateway Components: Distributed Router and Service Router (2)	182
5-16	Realization of Distributed Routers and Service Routers	183
5-17	Gateway Components in a Single-Tier Topology	184
5-18	Gateway Components in a Multitler Topology (1)	185
5-19	Gateway Components in a Multitler Topology (2)	186
5-20	Gateway Interfaces	187
5-21	Review of Learner Objectives	188
5-22	Lesson 2: NSX Edge and Edge Clusters	189
5-23	Learner Objectives	189
5-24	About the NSX Edge Node	190
5-25	About the NSX Edge Cluster	191
5-26	NSX Edge Node Form Factors	192
5-27	NSX Edge VM Sizing Options	193
5-28	Prerequisites for Deploying the NSX Edge Node VM	194
5-29	Deployment Considerations for NSX Edge Node VM Interfaces	195
5-30	Deploying the NSX Edge Node VM with Multiple N-VDS	196
5-31	Deploying the NSX Edge Node VM with a Single N-VDS	197
5-32	NSX Edge Node VM Network Offloading with SmartNiCs	198
5-33	Requirements for the NSX Edge Bare-Metal Node	199
5-34	Prerequisites for Deploying the NSX Edge Bare-Metal Node	200
5-35	Deployment Methods for NSX Edge Nodes	200
5-36	Deploying NSX Edge Nodes from the NSX UI (1)	201
5-37	Deploying NSX Edge Nodes from the NSX UI (2)	202
5-38	Deploying NSX Edge Nodes from vCenter	203
5-39	Using PXE to Deploy NSX Edge Nodes from an ISO File	204

5-40	Installing NSX Edge on Bare Metal	205
5-41	Joining NSX Edge Bare Metal with the Management Plane	206
5-42	Verifying the Edge Transport Node Status	207
5-43	Changing the NSX Edge VM Resource Reservations	
5-44	Changing Node Settings	209
5-45	Enabling UPT Mode on NSX Edge Node VMs	
5-46	Postdeployment Verification Checklist	21
5-47	Creating an NSX Edge Cluster	
5-48	Lab 6: Deploying and Configuring NSX Edge Nodes	213
5-49	Review of Learner Objectives	213
5-50	Lesson 3: Configuring Tier-0 and Tier-1 Gateways	214
5-51	Learner Objectives	214
5-52	Gateway Configuration Tasks	215
5-53	Creating the Tier-1 Gateway	
5-54	Connecting Segments to the Tier-1 Gateway	217
5-55	Using Network Topology to Validate the Tier-1 Gateway Configuration	218
5-56	Testing East-West Connectivity	219
5-57	Creating the Uplink Segments	
5-58	Creating the Tier-O Gateway (1)	22
5-59	Creating the Tier-0 Gateway (2)	
5-60	Configuring Routing	222
5-61	Connecting the Tier-1 and Tier-0 Gateways	223
5-62	Enabling Route Advertisement in the Tier-1 Gateway	224
5-63	Configuring Route Redistribution on the Tier-O Gateway	225
5-64	Using Network Topology to Validate the Tier-O Gateway Configuration	226
5-65	Testing North-South Connectivity	227
5-66	Lab 7: Configuring the Tier-1 Gateway	228
5-67	Review of Learner Objectives	
5-68	Lesson 4: Configuring Static and Dynamic Routing	229
5-69	Learner Objectives	
5-70	Static and Dynamic Routing	230
5-71	Tier-O Gateway Routing Configurations (1)	
5-72	Tier-0 Gateway Routing Configurations (2)	232
5-73	Configuring Static Routes on a Tier-O Gateway (1)	
5-74	Configuring Static Routes on a Tier-0 Gateway (2)	234
5-75	Configuring Dynamic Routing with BGP on Tier-0 Gateways (1)	235

5-76	Configuring Dynamic Routing with BGP on Tier-0 Gateways (2)	236
5-77	Verifying the BGP Configuration of the Tier-0 Gateways	237
5-78	BGP Route Aggregation	238
5-79	Configuring Route Aggregation with BGP	
5-80	Configuring Dynamic Routing with OSPF on Tier-0 Gateways (1)	240
5-81	Configuring Dynamic Routing with OSPF on Tier-0 Gateways (2)	241
5-82	Configuring Dynamic Routing with OSPF on Tier-0 Gateways (3)	242
5-83	Verifying OSPF Configuration of the Tier-0 Gateways	243
5-84	OSPF Route Summarization	244
5-85	Configuring Route Summarization with OSPF	
5-86	Lab 8: Creating and Configuring a Tier-0 Gateway with OSPF	245
5-87	Lab 9: Configuring the Tier-O Gateway with BGP	246
5-88	Review of Learner Objectives	246
5-89	Lesson 5: ECMP and High Availability	247
5-90	Learner Objectives	247
5-91	About Equal-Cost Multipath Routing	248
5-92	Enabling ECMP in BGP	249
5-93	Enabling ECMP in OSPF	
5-94	About High Availability	250
5-95	Active-Active HA Mode	251
5-96	Active-Active Topology with BGP	252
5-97	Active-Active Topology with OSPF	253
5-98	Active-Standby HA Mode	254
5-99	Active-Standby Topology with BGP	255
5-100	Active-Standby Topology with OSPF	256
5-101	Fallover Detection Mechanisms	257
5-102	About BFD	258
5-103	Fallover Scenario with BFD	259
5-104	Failover Scenario with Dynamic Routing	260
5-105	Fallover Modes	261
5-106	Review of Learner Objectives	261
5-107	Lesson 6: Logical Routing Packet Walk	262
5-108	Learner Objectives	262
5-109	Single-Tier Routing: Egress to Physical Network (1)	262
5-110	Single-Tier Routing: Egress to Physical Network (2)	263
5-111	Single-Tier Routing: Egress to Physical Network (3)	264

5-112	Single-Tier Routing: Egress to Physical Network (4)	265
5-113	Single-Tier Routing: Egress to Physical Network (5)	
5-114	Single-Tier Routing: Egress to Physical Network (6)	267
5-115	Single-Tier Routing: Ingress from Physical Network (7)	
5-116	Single-Tier Routing: ingress from Physical Network (8)	269
5-117	Single-Tier Routing: Ingress from Physical Network (9)	
5-118	Single-Tier Routing: ingress from Physical Network (10)	27
5-119	Single-Tier Routing: Ingress from Physical Network (11)	272
5-120	Multitier Routing: Egress to Physical Network (1)	273
5-121	Multitier Routing: Egress to Physical Network (2)	
5-122	Multitier Routing: Egress to Physical Network (3)	275
5-123	Multitier Routing: Egress to Physical Network (4)	
5-124	Multitler Routing: Egress to Physical Network (5)	277
5-125	Multitier Routing: Egress to Physical Network (6)	
5-126	Multitier Routing: Egress to Physical Network (7)	279
5-127	Multitier Routing: Egress to Physical Network (8)	
5-128	Multitler Routing: Egress to Physical Network (9)	281
5-129	Multitier Routing: Ingress from Physical Network (10)	
5-130	Multitier Routing: Ingress from Physical Network (11)	283
5-131	Multitler Routing: Ingress from Physical Network (12)	284
5-132	Multitler Routing: Ingress from Physical Network (13)	285
5-133	Multitler Routing: Ingress from Physical Network (14)	
5-134	Multitler Routing: Ingress from Physical Network (15)	287
5-135	Multitler Routing: Ingress from Physical Network (16)	288
5-136	Review of Learner Objectives	288
5-137	Lesson 7: VRF Lite	
5-138	Learner Objectives	289
5-139	About VRF Lite	
5-140	VRF Lite Requirements and Limitations	291
5-141	Use Cases for VRF Lite	292
5-142	VRF Lite Topologies	293
5-143	VRF Lite Gateway Interfaces	
5-144	VRF Lite: Control and Data Planes	295
5-145	Configuring VRF Lite	
5-146	Deploying the Default Tier-0 Gateway	297
5-147	Adding Uplink Interfaces to the Default Tier-O Gateway	298

5-148	Configuring BGP for the Default Tier-0 Gateway	299
5-149	Adding the Uplink Trunk Segment for the VRF Gateway	
5-150	Deploying the VRF Gateway	301
5-151	Adding Uplink Interfaces to the VRF Gateway	302
5-152	Configuring the BGP for the VRF Gateway	303
5-153	Connecting a Tier-1 Gateway to the VRF Gateway	
5-154	VRF Lite Validation	305
5-155	Lab 10: Configuring VRF Lite	306
5-156	Review of Learner Objectives	306
5-157	Key Points (1)	307
5-158	Key Points (2)	307
Modu	lle 6 NSX Logical Bridging	309
6-2	Importance	309
6-3	Lesson 1: NSX Logical Bridging	310
6-4	Learner Objectives	310
6-5	Overview of Logical Bridging	310
6-6	Logical Bridging Use Cases	311
6-7	Routing and Bridging for Physical-to-Virtual Communication	312
6-8	Example of Virtual-to-Physical Routing	314
6-9	Example of Virtual-to-Physical Bridging	315
6-10	Logical Bridging Components	316
6-11	Using Multiple Bridge Profiles	317
6-12	Creating an Edge Bridge Profile	318
6-13	Creating a Layer 2 Bridge-Backed Segment	319
6-14	Monitoring the Bridged Traffic Statistics	320
6-15	Review of Learner Objectives	321
6-16	Key Points	321
Modu	le 7 NSX Firewalls	323
7-2	Importance	323
7-3	Module Lessons	323
7-4	Lesson 1: NSX Segmentation	324
7-5	Learner Objectives	324
7-6	Traditional Security Challenges	324
7-7	About Zero-Trust Security	325
7-8	About NSX Segmentation	326

7-9	Use Cases for NSX Segmentation	327
7-10	NSX Segmentation Benefits	
7-11	Enforcing Zero-Trust with NSX Segmentation	329
7-12	Step 1: Creating Virtual Security Zones	
7-13	Step 2: Identifying the Application Boundaries	33
7-14	Step 3: Implementing Micro-Segmentation	
7-15	Step 4: Securing Through Context	334
7-16	Review of Learner Objectives	
7-17	Lesson 2: NSX Distributed Firewall	336
7-18	Learner Objectives	
7-19	NSX Firewalls	337
7-20	Features of the Distributed Firewall	
7-21	Distributed Firewall: Key Concepts	339
7-22	Overview of a Security Policy	
7-23	Distributed Firewall Policy Categories	341
7-24	About Distributed Firewall Policies	
7-25	Distributed Firewall Rule Processing within a Policy	343
7-26	Applied To Field for the Policy	
7-27	Configuring Distributed Firewall Policy Settings	345
7-28	Configuring Time-Based Firewall Policies	346
7-29	Creating Distributed Firewall Rules	347
7-30	Configuring Distributed Firewall Rule Parameters	348
7-31	Specifying Sources and Destinations for a Rule	349
7-32	Creating Groups	349
7-33	Adding Members and Member Criteria for a Group	350
7-34	Creating Groups Based on Tags	35
7-35	Specifying Services for a Rule	352
7-36	Adding a Context Profile to a Rule	353
7-37	Configuring Context Profile Attributes	354
7-38	Custom FQDN Filtering	355
7-39	Setting the Scope of Rule Enforcement	356
7-40	Specifying the Action for a Rule	
7-41	Distributed Firewall Rule Settings	358
7-42	Blocking Malicious IPs in the Distributed Firewall	359
7-43	Enable or Disable Malicious IP Feeds	360
7-44	Default Malicious IP Group	361

7-45	Creating a Custom Malicious IP Group	362
7-46	Default Malicious IP Block Rules	
7-47	Configuring Rules to Block Malicious IPs	363
7-48	Saving and Viewing the Distributed Firewall Configuration	364
7-49	Rolling Back to a Saved Distributed Firewall Configuration	365
7-50	Distributed Firewall Configuration Export and Import	
7-51	Distributed Firewall Architecture	367
7-52	Distributed Firewall Architecture: ESXI	
7-53	Distributed Firewall Rule Processing: ESXI	370
7-54	Lab 11: Configuring the NSX Distributed Firewall	371
7-55	Review of Learner Objectives	371
7-56	Lesson 3: Use Case for Security in Distributed Firewall on VDS	372
7-57	Learner Objectives	372
7-58	About Distributed Firewall on VDS	
7-59	Supported Features	374
7-60	Distributed Firewall on VDS Requirements	375
7-61	Installation Workflow	376
7-62	Preparing the Cluster for Security	376
7-63	Validating the Security Cluster Preparation from the NSX UI	377
7-64	Transport Node Preparation	378
7-65	Autoconfigured Transport Node Profile	378
7-66	VLAN Transport Zones	379
7-67	Discovered Segments (1)	380
7-68	Discovered Segments (2)	381
7-69	Configuring Segment Profiles	382
7-70	Using VDS Attributes to Define NSX Groups	383
7-71	Review of Learner Objectives	383
7-72	Lesson 4: NSX Gateway Firewall	384
7-73	Learner Objectives	384
7-74	About the Gateway Firewall	385
7-75	Predefined Gateway Firewall Categories	387
7-76	Gateway Firewall Policy	388
7-77	Configuring Gateway Firewall Policy Settings	389
7-78	Configuring Gateway Firewall Rules	390
7-79	Configuring Gateway Firewall Rules Settings	391
7-80	Gateway Firewall Architecture	392

7-81	Gateway Firewall Rule Processing	393
7-82	Lab 12: Configuring the NSX Gateway Firewall	394
7-83	Review of Learner Objectives	394
7-84	Key Points	395
Modu	Ile 8 NSX Advanced Threat Prevention	397
8-2	Importance	397
8-3	Module Lessons	397
8-4	Lesson 1: NSX Intrusion Detection and Prevention	398
8-5	Learner Objectives	398
8-6	About NSX Distributed IDS/IPS	399
8-7	Use Cases for NSX Distributed IDS/IPS	400
8-8	About Behavior-Based IDS/IPS	401
8-9	Requirements for NSX Distributed IDS/IPS	402
8-10	About IDS/IPS Signatures	403
8-11	About IDS/IPS Profiles	404
8-12	About IDS/IPS Policies and Rules	405
8-13	IDS/IPS Signature Curation	406
8-14	NSX Distributed IDS/IPS Architecture	407
8-15	Configuring NSX Distributed IDS/IPS	408
8-16	Configuring IDS/IPS Signatures	409
8-17	Global Intrusion Signature Management	410
8-18	Configuring Custom IDS/IPS Profiles	412
8-19	Configuring IDS/IPS Rules	413
8-20	Monitoring IDS/IPS Events (1)	414
8-21	Monitoring IDS/IPS Events (2)	415
8-22	About North-South IDS/IPS	416
8-23	Use Cases for North-South IDS/IPS	416
8-24	Requirements for Configuring North-South IDS/IPS	416
8-25	North-South IDS/IPS Architecture	417
8-26	Configuring North-South IDS/IPS	
8-27	Configuring North-South IDS/IPS Rules	418
8-28	Monitoring North-South IDS/IPS Events	419
8-29	Lab 13: Configuring Distributed Intrusion Detection and Prevention	420
8-30	Review of Learner Objectives	420
8-31	Lesson 2: NSX Application Platform	421
8-32	Learner Objectives	

8-33	About NSX Application Platform	421
8-34	Prerequisites for NSX Application Platform Deployment	422
8-35	Setting Up a Private Harbor Registry	423
8-36	NSX Application Platform Form Factors	424
8-37	NSX Application Platform Deployment (1)	425
8-38	NSX Application Platform Deployment (2)	426
8-39	NSX Application Platform Predeployment Checks	427
8-40	NSX Application Platform Deployment Validation	428
8-41	NSX Application Platform Services	429
8-42	Objects Created During the NSX Application Platform Deployment	431
8-43	Basic kubecti Commands	432
8-44	Namespaces Available After NSX Application Platform Deployment	433
8-45	Pods Available After NSX Application Platform Deployment	434
8-46	Lab 14: (Simulation) Deploying NSX Application Platform	435
8-47	Review of Learner Objectives	435
8-48	Lesson 3: NSX Malware Prevention	436
8-49	Learner Objectives	436
8-50	About Malware Prevention	437
8-51	About East-West Malware Prevention	439
8-52	Use Cases for East-West Malware Prevention	440
8-53	Requirements for East-West Malware Prevention	441
8-54	East-West Malware Prevention Architecture	442
8-55	NSX Application Platform Components	443
8-56	ESXI Host Components.	445
8-57	East-West Malware Prevention Packet Flow for Known Files	447
8-58	East-West Malware Prevention Packet Flow for Unknown Files	448
8-59	Activating Malware Prevention on NSX Application Platform	449
8-60	Setting the Cloud Region	450
8-61	Validating the Malware Prevention Installation	451
8-62	Service Registration	452
8-63	Service Registration Validation from the NSX UI	452
8-64	Service Deployments	453
8-65	Service Deployment Validation from the NSX UI	454
8-66	Service Deployment Validation from vCenter Server	455
8-67	Creating East-West Malware Prevention Profiles	456
8-68	Creating Rules for East-West Malware Prevention	457

8-69	About North-South Malware Prevention	458
8-70	Use Cases for North-South Malware Prevention	459
8-71	Requirements for North-South Malware Prevention	460
8-72	North-South Malware Prevention Architecture	461
8-73	NSX Edge Components	462
8-74	North-South Malware Prevention Packet Flow for Known Files	
8-75	North-South Malware Prevention Packet Flow for Unknown Files	464
8-76	Enabling Malware Prevention on Tier-1 Gateways	466
8-77	Creating North-South Malware Prevention Profiles	466
8-78	Creating Rules for North-South Malware Prevention	467
8-79	Malware Prevention Dashboard (1)	468
8-80	Malware Prevention Dashboard (2)	469
8-81	About the Allowlist	470
8-82	Lab 15: (Simulation) Configuring Malware Prevention for East-West Traffic	471
8-83	Review of Learner Objectives	471
8-84	Lesson 4: NSX Intelligence	472
8-85	Learner Objectives	472
8-86	About NSX Intelligence	473
8-87	Use Cases for NSX Intelligence	474
8-88	NSX Intelligence Requirements	474
8-89	NSX Intelligence Activation	475
8-90	Validating the NSX Intelligence Deployment	476
8-91	Granular Data Collection	477
8-92	NSX Intelligence Visualization (1)	478
8-93	NSX Intelligence Visualization (2)	479
8-94	NSX Intelligence Recommendations (1)	480
8-95	NSX Intelligence Recommendations (2)	481
8-96	NSX Intelligence Recommendations (3)	483
8-97	NSX Intelligence Recommendations (4)	484
8-98	NSX Suspicious Traffic Detection	485
8-99	Configuring Detector Definitions	487
8-100	Visualizing Detected Threats (1)	488
8-101	Visualizing Detected Threats (2)	489
8-102	Review of Learner Objectives	489
8-103	Lesson 5: NSX Network Detection and Response	490
8-104	Learner Objectives	490

8-105	About NSX Network Detection and Response	491
8-106	NSX Network Detection and Response Use Cases	
8-107	NSX Network Detection and Response High-Level Architecture	493
8-108	NSX Network Detection and Response in NSX Deployments	494
8-109	NSX Network Detection and Response Architecture (1)	495
8-110	NSX Network Detection and Response Architecture (2)	
8-111	NSX Network Detection and Response Requirements	498
8-112	NSX Network Detection and Response Activation (1)	
8-113	NSX Network Detection and Response Activation (2)	501
8-114	Validating the NSX Network Detection and Response and NSX Cloud Deployments	
8-115	Visualizing and Mitigating Attacks	
8-116	Accessing the NSX Network Detection and Response UI	
8-117	Campaign Overview: Active Threats and Attack Stages	
8-118	Campaign Blueprint	
8-119	Campaign Timeline	506
8-120	Reviewing Events	
8-121	Lab 16: (Simulation) Using NSX Network Detection and Response to D 507	etect Threats
8-122	Review of Learner Objectives	507
8-123	Key Points (1)	508
8-124	Key Points (2)	
Modu	le 9 NSX Services	509
9-2	Importance	509
9-3	Module Lessons.	509
9-4	Lesson 1: Configuring NAT	510
9-5	Learner Objectives	510
9-6	About NAT	511
9-7	About SNAT	512
9-8	About DNAT	
9-9	About Reflexive NAT	514
9-10	Configuring SNAT and DNAT	
9-11	Configuring Reflexive NAT	517
9-12	Stateful Active-Active Services	
9-13	Stateful Active-Active Interface Group	521
9-14	Stateful Active-Active SNAT	522

9-15	Stateful Active-Active Configuration	523
9-16	Configuring Interface Groups	
9-17	Configuring Stateful SNAT and DNAT	525
9-18	About NAT64	
9-19	Configuring NAT64 Rules	527
9-20	Lab 17: Configuring Network Address Translation	528
9-21	Review of Learner Objectives	528
9-22	Lesson 2: Configuring DHCP and DNS Services	529
9-23	Learner Objectives	529
9-24	About DHCP	
9-25	About DHCP Relay	531
9-26	DHCP in NSX	
9-27	DHCP Local Server	532
9-28	DHCP Relay	
9-29	Gateway DHCP	534
9-30	DHCP Workflow	
9-31	Creating a DHCP Profile	535
9-32	Assigning the DHCP Profile to a Segment	
9-33	Assigning the DHCP Profile to a Gateway	537
9-34	Configuring DHCP Services	538
9-35	About DNS Services	539
9-36	About DNS Forwarder	
9-37	DNS Workflow	540
9-38	Creating DNS Zones.	
9-39	Creating DNS Forwarder Services	542
9-40	Establishing Forwarder Connectivity	
9-41	Configuring DHCP to Allocate DNS	544
9-42	Review of Learner Objectives	544
9-43	Lesson 3: Configuring NSX Advanced Load Balancer	545
9-44	Learner Objectives	545
9-45	About NSX Advanced Load Balancer	545
9-46	Benefits of NSX Advanced Load Balancer	
9-47	NSX Advanced Load Balancer Feature Edition Comparison (1)	547
9-48	NSX Advanced Load Balancer Feature Edition Comparison (2)	548
9-49	NSX Advanced Load Balancer Architecture	549
9-50	NSX Advanced Load Balancer Deployment Workflow	

9-51	NSX Advanced Load Balancer Consumption Workflow	551
9-52	Requirements for NSX Advanced Load Balancer	552
9-53	Deploying the NSX Advanced Load Balancer Controller Cluster	553
9-54	Service Engines Deployment and Connectivity	554
9-55	Creating a Cloud Connector (1)	555
9-56	Creating a Cloud Connector (2)	
9-57	Creating a Service Engine Group	557
9-58	NSX Advanced Load Balancer Components	
9-59	NSX Advanced Load Balancer Topologies	560
9-60	VIP Placement and Route Redistribution	
9-61	North-South Traffic	562
9-62	East-West Traffic (1)	
9-63	East-West Traffic (2)	564
9-64	Accessing the NSX Advanced Load Balancer UI	
9-65	Creating a Virtual IP Address	566
9-66	Creating a Server Pool	
9-67	Configuring Load-Balancing Algorithms	570
9-68	Configuring Health Monitor Profiles	
9-69	Configuring Persistence Profiles	572
9-70	Configuring Server Pool Security Settings	573
9-71	Creating a Virtual Service	574
9-72	Validating Virtual Services and Server Pools	
9-73	Lab 18: Configuring NSX Advanced Load Balancer	577
9-74	Review of Learner Objectives	577
9-75	Lesson 4: IPSec VPN	578
9-76	Learner Objectives	
9-77	Use Cases for IPSec VPN	578
9-78	IPSec VPN Protocols and Algorithms	579
9-79	IPSec VPN Methods	580
9-80	IPSec VPN Modes	581
9-81	IPSec VPN Types	582
9-82	NSX IPSec VPN Deployment	
9-83	IPSec VPN: High Availability	584
9-84	IPSec VPN Workflow	585
9-85	Configuring an IPSec VPN Service	586
9-86	Configuring DPD Profiles	587

9-87	Configuring IKE Profiles	588
9-88	Configuring IPSec Profiles	
9-89	Adding a Local Endpoint	590
9-90	Configuring IPSec VPN Sessions	
9-91	Configuring Policy-Based IPSec VPN Sessions (1)	592
9-92	Configuring Policy-Based IPSec VPN Sessions (2)	
9-93	Configuring Route-Based IPSec VPN Sessions	594
9-94	Traceflow and Live Traffic Analysis Support for VPN	
9-95	Review of Learner Objectives	596
9-96	Lesson 5: L2 VPN	
9-97	Learner Objectives	597
9-98	About Layer 2 VPN	
9-99	L2 VPN Architecture	599
9-100	L2 VPN Edge Packet Flow	600
9-101	L2 VPN Considerations	601
9-102	Recommended L2 VPN Clients	602
9-103	About Autonomous Edge	603
9-104	Sample L2 VPN Network Topology	604
9-105	L2 VPN Server Workflow	604
9-106	Configuring an IPSec VPN Service	605
9-107	Adding a Local Endpoint	606
9-108	Adding an L2 VPN Server Service	607
9-109	Adding an L2 VPN Server Session	608
9-110	Attaching Segments to the L2 VPN (1)	609
9-111	Attaching Segments to the L2 VPN (2)	610
9-112	L2 VPN Client Workflow	610
9-113	Configuring an L2 VPN Client Service	611
9-114	Adding an L2 VPN Client Session (1)	612
9-115	Adding an L2 VPN Client Session (2)	613
9-116	Attaching Segments to the L2 VPN Client	614
9-117	Lab 19: Deploying Virtual Private Networks	615
9-118	Review of Learner Objectives	615
9-119	Key Points (1)	616
9-120	Key Points (2)	616
Modu	le 10 NSX User and Role Management	617
10-2	Importance	
110000	51 (** VATA 1047) \$ 145	

10-3	Module Lessons	617
10-4	Lesson 1: Integrating NSX with VMware Identity Manager	618
10-5	Learner Objectives	618
10-6	About VMware Identity Manager	619
10-7	Benefits of Integrating VMware Identity Manager with NSX	620
10-8	Prerequisites for VMware Identity Manager Integration	621
10-9	Configuring VMware Identity Manager	622
10-10	Overview of the VMware Identity Manager and NSX Integration	623
10-11	Creating an OAuth Client	623
10-12	Obtaining the SHA-256 Certificate Thumbprint	624
10-13	Configuring the VMware Identity Manager Details in NSX	625
10-14	Verifying the VMware Identity Manager Integration	626
10-15	Default UI Login	627
10-16	UI Login with VMware Identity Manager	628
10-17	Local Login with VMware identity Manager	629
10-18	Review of Learner Objectives	629
10-19	Lesson 2: Integrating NSX with LDAP	630
10-20	Learner Objectives	630
10-21	About LDAP	630
10-22	Benefits of Integrating LDAP with NSX	631
10-23	Authentication with LDAP	631
10-24	Adding an Identity Source	632
10-25	Configuring the LDAP Server	633
10-26	UI Login with LDAP	634
10-27	Review of Learner Objectives	634
10-28	Lesson 3: Managing Users and Configuring RBAC	635
10-29	Learner Objectives.	635
10-30	NSX Users	635
10-31	Activating Guest Users	636
10-32	Using Role-Based Access Control	637
10-33	Built-in Roles (1)	638
10-34	Bullt-In Roles (2)	638
10-35	Custom Role-Based Access Control.	639
10-36	Creating Custom Roles (1)	640
10-37	Creating Custom Roles (2)	
10-38	Multitenancy Hierarchy Model in NSX 4.0.1	642

10-39	Object-Based RBAC in a Multitenancy Environment	643
10-40	Role Assignment	
10-41	Lab 20: Managing Users and Roles	645
10-42	Review of Learner Objectives	
10-43	Key Points	645
Modul	e 11 NSX Federation	
11-2	Importance	647
11-3	Module Lessons.	647
11-4	Lesson 1: Federation Architecture	648
11-5	Learner Objectives.	648
11-6	About NSX Federation	649
11-7	NSX Federation Components: Global Manager	650
11-8	NSX Federation Components: Local Manager	650
11-9	NSX Federation Components: Global Manager and Local Manager Clusters	651
11-10	Federation Configuration Types	652
11-11	Ownership of Logical Configuration (1)	652
11-12	Ownership of Logical Configuration (2)	653
11-13	Infrastructure Ownership	654
11-14	Global Configuration	655
11-15	Local Configuration	656
11-16	Federation Configuration Example	657
11-17	Federation Configuration Example Workflow (1)	658
11-18	Federation Configuration Example Workflow (2)	659
11-19	Review of Learner Objectives	
11-20	Lesson 2: Installing and Onboarding Federation	660
11-21	Learner Objectives	660
11-22	NSX Federation: Prerequisites	660
11-23	Onboarding Process	661
11-24	Active Global Manager Configuration	662
11-25	Adding Standby Global Manager (1)	663
11-26	Adding Standby Global Manager (2)	664
11-27	Adding a Location	665
11-28	Validating the Local Manager	666
11-29	Location Onboarding	
11-30	Onboarding Preparation	669
11-31	Verifying Onboarding	
XXIV		

11-32	Review of Learner Objectives	670
11-33	Lesson 3: Federation Networking	
11-34	Learner Objectives	671
11-35	Stretched Networking (1)	
11-36	Stretched Networking (2)	672
11-37	Tier-O and Tier-1 Gateways: Logical Topologies (1)	
11-38	Tier-0 and Tier-1 Gateways: Logical Topologies (2)	674
11-39	Tier-O and Tier-1 Gateways: Logical Topologies (3)	
11-40	Single-Location Tier-O Gateway Deployments	676
11-41	Single-Location Tier-1 Gateway Deployments	
11-42	Multilocation Tier-0 and Tier-1 Gateway Deployments (1)	678
11-43	Multilocation Tier-0 and Tier-1 Gateway Deployments (2)	679
11-44	Multilocation T0-Stretched Gateway Modes (1)	680
11-45	Multilocation T0-Stretched Gateway Modes (2)	
11-46	Multilocation T1-Stretched Gateway Modes (1)	682
11-47	Multilocation T1-Stretched Gateway Modes (2)	
11-48	About RTEP	684
11-49	Stretched Layer-2 Network	
11-50	Stretched L2: VNI Mapping	686
11-51	About Federation L2 Bridging	687
11-52	Components of Federation L2 Bridging	688
11-53	Federation L2 Bridge Communication	
11-54	Federation Traceflow	690
11-55	Review of Learner Objectives	690
11-56	Lesson 4: Federation Security	691
11-57	Learner Objectives	691
11-58	Stretched Security in NSX Federation	692
11-59	Use Cases for Security	693
11-60	Security Configuration Workflow	694
11-61	About Regions	695
11-62	About Groups	696
11-63	GM Groups and Span Example (1)	697
11-64	GM Groups and Span Example (2)	698
11-65	Group Membership Criteria	699
11-66	Group Membership Based on the VM Tag (1)	700
11-67	Group Membership Based on the VM Tag (2)	701

11-68	About GM-Based Policy	702
11-69	About GM-Based Rules	703
11-70	Overlap of GM and LM Sections	704
11-71	GM DFW Management Enhancement	
11-72	Global DFW Exclusion List Enhancement	706
11-73	Time-Based DFW Rules Enhancement	707
11-74	Review of Learner Objectives	708
11-75	Key Points	

Module 1

Course Introduction

1-2 Course Introduction

1-3 Importance

NSX is the network virtualization and security platform that enables the virtual cloud network. The virtual cloud network is a software-defined approach to networking that extends across data centers, clouds, and application frameworks. An application might run on virtual machines, containers, or bare metal. NSX brings networking and security closer to the location where the application runs. The application framework that you create can support virtual machines running on ESXI hosts, containers, bare-metal servers, and public clouds.

In an NSX environment, you can select the technologies that best sult your applications. You can also perform your daily operational and management tasks with various tools supported by NSX.

1-4 Learner Objectives

- Describe the architecture and main components of NSX
- Explain the features and benefits of NSX
- Deploy the NSX Management cluster and NSX Edge nodes
- Prepare ESXI hosts to participate in NSX networking
- Create and configure segments for layer 2 forwarding
- Create and configure Tier-0 and Tier-1 gateways for logical routing
- Use distributed and gateway firewall policies to filter east-west and north-south traffic in NSX
- Configure Advanced Threat Prevention features
- Configure network services on NSX Edge nodes
- Use VMware Identity Manager and LDAP to manage users and access
- Explain the use cases, importance, and architecture of Federation

1-5 Course Outline

- Course Introduction
- 2. VMware Virtual Cloud Network and VMware NSX
- 3. Preparing the NSX Infrastructure
- 4. NSX Logical Switching
- 5. NSX Logical Routing
- 6. NSX Logical Bridging
- NSX Firewalls
- 8. NSX Advanced Threat Prevention
- 9. NSX Services
- 10. NSX User and Role Management
- 11. NSX Federation

1-6 Typographical Conventions

The following typographical conventions are used in this course.

Conventions	Usage and Examples
Monospace	Identifies command names, command options, parameters, code fragments, error messages, filenames, folder names, directory names, and path names:
	Run the esxtop command.
	 found in the var/log/messages file.
Monospace Bold	Identifies user inputs:
	 Enter ipconfig /release.
Boldface	Identifies user interface controls:
	Click the Configuration tab.
Italic	Identifies book titles:
	vSphere Virtual Machine Administration
<>	Indicates placeholder variables:
	 <esxi_host_name></esxi_host_name>
	 the Settings/<your_name>.txt file</your_name>

1-7 References

Title	Location
NSX Administration Guide	https://docs.vmware.com/en/VMware-NSX/4.0/administration/GUID FBFD577B-745C-4658-B713-A3016D18CB9A.html
NSX Installation Guide	https://docs.vmware.com/en/VMware-NSX/4.0/installation/GUID- 3E0C4CEC-D593-4395-84C4-150CD6285963.html
NSX Upgrade Guide	https://docs.vmware.com/en/VMware-NSX/4.0/upgrade/GUID- E04242D7-EF09-4601-8906-3FA77FBB06BD.html

1-8 VMware Online Resources

Documentation for NSX: https://docs.vmware.com/en/VMware-NSX/index.html

Documentation for VMware Validated Solutions: https://core.vmware.com/vmware-validatedsolutions.

VMware Communities: http://communities.vmware.com

- Start a discussion.
- Access the knowledge base.
- Access documentation, technical papers, and compatibility guides.
- Access communities.
- Access user groups.

VMware Support: http://www.vmware.com/support

VMware Hands-on Labs: http://labs.hol.vmware.com

VMware Learning: http://www.vmware.com/learning

- Access course catalog.
- Access worldwide course schedule.


1-9 VMware Learning Overview

You can access the following Education Services:

- VMware Learning Paths:
 - Help you find the course that you need based on the product, your role, and your level of experience
 - Can be accessed at https://vmware.com/learning
- VMware Customer Connect Learning, which is the official source of digital training, includes the following options:
 - On Demand Courses: Self-paced learning that combines lecture modules with hands-on practice labs
 - VMware Lab Connect: Self-paced, technical lab environment where you can practice skills learned during instructor-led training
 - Certification Exam Prep: Comprehensive video-based reviews of exam topics and objectives to help you take your certification exam
- For more information, see https://vmware.com/learning/connect-learning.

1-10 VMware Certification Overview

VMware certifications validate your expertise and recognize your technical knowledge and skills with VMware technology.



VMware certification sets the standards for IT professionals who work with VMware technology. Certifications are grouped into technology tracks. Each track offers one or more levels of certification (up to four levels).

For the complete list of certifications and details about how to attain these certifications, see https://vmware.com/certification.

1-11 VMware Credentials Overview

VMware badges are digital emblems of skills and achievements. Career certifications align to job roles and validate expertise across a solution domain. Certifications can cover multiple products in the same certification.











Specialist certifications and skills badges align to products and verticals and show expanded expertise.





Digital badges have the following features:

- Easy to share in social media (Linkedin, Twitter, Facebook, blogs, and so on)
- Validate and verify achievement
- Contain metadata with skill tags and accomplishments
- Based on Mozilla's Open Badges standard

For the complete list of digital badges, see http://www.pearsonvue.com/vmware/badging.

Module 2

VMware Virtual Cloud Network and VMware NSX

2-2 Importance

As a network administrator, you must understand the VMware Virtual Cloud Network framework and the solutions that it offers for addressing challenges in your data center. You must also understand the VMware NSX architecture and components so that you can properly design, deploy, and manage a data center that meets your business requirements.

2-3 Lesson 1: VMware Virtual Cloud Network and VMware NSX

2-4 Learner Objectives


- Describe the purpose of VMware Virtual Cloud Network and its framework
- Identify the benefits and recognize the use cases for NSX
- Describe how VMware NSX fits Into the NSX product portfolio
- Recognize features and the main elements in the NSX architecture
- Enumerate the deprecated features in NSX 4.0.x
- Describe the NSX policy and centralized policy management.
- Describe the NSX management cluster and the management plane
- Identify the functions of control plane components, data plane components, and communication channels

2-5 Virtual Cloud Network Framework

Virtual Cloud Network is the VMware framework for connecting and protecting different types of workloads running across various environments.

The Virtual Cloud Network framework is built on the NSX technology.

Virtual Cloud Network is a software layer. This layer provides connectivity between data center, cloud, and edge infrastructure with data visibility and security.



Virtual Cloud Network connects and protects applications and data, regardless of their physical locations. Virtual Cloud Network also connects and protects workloads running across any environment. Workloads might be running on premises in a customer data center, in a branch, or in a public cloud such as Amazon AWS or Microsoft Azure.

Virtual Cloud Network enables organizations to embrace cloud networking as the softwaredefined architecture for connecting components in a distributed world.

Virtual Cloud Network is a ubiquitous software layer that provides maximum visibility into, and context for, the interaction among various users, applications, and data. NSX supports various types of endpoints.


The VMware software-based approach delivers a networking and security platform that enables customers to connect, secure, and operate an end-to-end architecture to deliver services to applications. The VMware software-based approach provides the following benefits:

- Enables you to design and build the next-generation policy-driven data center.
 - This data center connects, secures, and automates traditional hypervisors and new microservices-based (container) applications across a range of deployment targets such as the data center, cloud, and so on.
- Embeds security in the platform by compartmentalizing the network through microsegmentation, encrypting in-flight data, and automatically detecting and responding to security threats.
- Delivers a WAN solution that provides full visibility, metrics, control, and automation of all endpoints.

2-6 NSX Portfolio

NSX provides consistent networking and security across the entire IT environment.

Virtual Cloud Network is based on a robust portfolio of products built on the foundations of the concept of any infrastructure, any cloud, any application, any platform, and any device. Virtual Cloud Network includes several key solutions that provide security, integration, extensibility, automation, and elasticity.



VMware Virtual Cloud Network enables you to run your applications everywhere.

NSX supports cloud-native applications, bare-metal workloads, ESXI hypervisors, public clouds, and multiple clouds.

You can bring key capabilities from one central control point to wherever your applications run.

Network virtualization and security includes the following solutions:

- NSX is the industry's only complete layer 2 to layer 7 software-defined networking stack, including networking, load balancing, security, and analytics. With NSX, you can provision networking and security services across ESXI hypervisors and bare-metal servers.
- NSX Cloud extends the networking and security capabilities of NSX to the public cloud. You
 can provide your workloads running natively on Amazon AWS or Microsoft Azure with
 consistent networking and security policies, helping you improve scalability, control, and
 visibility.
- NSX Distributed IDS/IPS is an advanced threat detection engine built to detect lateral threat
 movement on east-west network traffic across multicloud environments.

- NSX Advanced Load Balancer enables you to deliver multicloud application services such as load balancing, application, security, autoscaling, container networking, and web application firewall.
- VMware SD-WAN virtualizes WAN connections to deliver high-performance, reliable branch access to cloud services, private data centers, and SaaS-based enterprise applications.
- VMware HCX makes it easy to migrate thousands of virtual machines within and across data centers or clouds, without requiring a reboot.

The NSX portfolio supports management and automation tools:

- NSX Intelligence is a distributed analytics solution that provides visibility and dynamic security policy enforcement for NSX environments. NSX Intelligence enables network and application security teams to deliver a granular security posture, simplify compliance analysis, and enable proactive security. It also supports network traffic analysis to help identify advanced threats in the environment.
- VMware Aria Operations for Networks (former name: VMware vRealize Network Insight) provides visibility across virtual and physical networks. It helps with operations management for NSX and NSX Cloud.
- VMware Aria Automation (former name: vRealize Automation) is the VMware infrastructure automation platform for the modern software-defined data center. When used with NSX, it automates an application's network connectivity, security, performance, and availability.
- Tanzu Service Mesh Advanced, built on VMware NSX, is the VMware enterprise-class service mesh solution that provides consistent control and security for microservices, end users, and data across the most demanding multicluster and multicloud environments.

2-7 Product Name Changes to VMware NSX 4.x

NSX-T Data Center is now called NSX.



The new name better reflects the multifaceted value that NSX brings to customers. The update is apparent in the product graphical user interface as well as documentation. The change neither affects the functionality of the product nor changes the API to affect compatibility with previous releases.

2-8 Use Cases for NSX

NSX can be used in several ways.



You can use NSX for the following purposes:

- Security: Delivers application-centric security at the workload level to prevent the lateral spread of threats
- Multicloud networking: Brings consistency in networking and security across varied sites and streamlines multicloud operations
- Automation: Enables faster deployment through automation by reducing manual, errorprone tasks
- Cloud-native applications: Enables native networking and security for containerized workloads across application frameworks

2-9 NSX Features (1)

Platform

- · Policy-driven configuration
- · Bare-metal server support
- . Support for VMs and containers
- Support for Amazon AWS and Microsoft Azure instances
- HTML 5 UI for management plane
- Data Plane Development Kit (DPDK) based NSX Edge nodes: VM or bare-metal form factor
- Multisite/Federation
- V2T migration
- DPU-Based Acceleration

Networking

- Overlay or VLAN-backed logical switching
- · Layer 2 bridging and QoS
- Distributed routing
- Static routing and equal-cost multipath (ECMP) routing
- BGP and OSPF support
- · Duplicate IP detection
- Bidirectional Forwarding (BFD) for fast convergence
- . VRF Lite and EVPN
- · Rate limiting
- · Layer 3 multicast

Security and Services

- · Gateway firewall
- · Distributed firewall
- · IDPS
- · Malware Prevention
- URL Filtering and FODN Analysis
- · TLS Inspection
- NSX Intelligence
- NSX Network Detection and Response
- NAT and NAT64
- . DNS and DHCP
- · Load Balancing
- . L2 VPN and IPSec VPN

2-10 NSX Features (2)

Automation

- REST/JSON API support
- Upstream OpenStack support and partner ecosystem
- Tag-based security grouping
- · Inventory support
- Management Plane to Policy Promotion Tool

Operations

- · Getting Started wizards
- · Dashboards
- · Role-based access control
- . Object-based access control
- · Upgrade coordinator
- · Backup and restore
- · vRealize Log Insight

Troubleshooting

- IP Flow Information Export (IPFIX)
- · Port mirroring
- Traceflow
- · Network Topology views
- · Selective technical support logs
- Monitoring Dashboard and statistics
- Alarms

2-11 NSX 4.0.0.1 Deprecated Features

The following NSX features are deprecated in NSX 4.0.0.1:

- N-VDS host switch
- KVM hypervisor
- NSX Advanced Load Balancer Policy API and UI

Since NSX 4.0.0.1, N-VDS host switch will not be supported for ESXI hosts. In a brownfield environment, you must migrate the ESXI hosts to VDS before upgrading to NSX 4.0.0.1.


N-VDS will remain the supported virtual switch on NSX Edge nodes, native public cloud NSX agents, and bare-metal workloads.

NSX 4.0.0.1 no longer supports KVM-based hypervisors. In a brownfield environment, you must remove the KVM transport nodes from NSX before upgrade.

NSX Advanced Load Balancer UI and Policy API are deprecated from NSX 4.0.0.1. All configuration tasks for load balancers integrated with NSX environments should be performed directly through the NSX Advanced Load Balancer UI and API. For all orchestration integrations use cases, use the NSX Advanced Load Balancer API or UI.

2-12 High-Level Architecture of NSX

The three main elements of NSX architecture are the management, control, and data planes. This architectural separation enables scalability without affecting workloads.



Each plane has its own components:

- Management plane: The management plane is designed with advanced clustering technology, which allows the platform to process large-scale concurrent API requests. NSX Manager provides the REST API and a web-based UI interface entry point for all user configurations.
- Control plane: The control plane manages computing and distributing the runtime virtual
 networking and security state of the NSX environment. The control plane includes a central
 control plane (CCP) and a local control plane (LCP). This separation significantly simplifies
 the work of the CCP and enables the platform to extend and scale for various endpoints.
 The management plane and control plane are converged. Each manager node in NSX is an
 appliance with converged functions, including management, control, and policy.


- Data plane: The data plane includes a group of ESXI hosts and NSX Edge nodes. The group
 of servers and edge nodes prepared for NSX are called transport nodes. Transport nodes
 manage the distributed forwarding of network traffic. The ESXI hosts managed by vCenter
 Server use vSphere Distributed Switch (VDS) for the traffic forwarding.
- Consumption plane: Although the consumption plane is not part of NSX, it provides integration into any CMP through the REST API and integration with VMware cloud management planes, such as vRealize Automation:
 - The consumption of NSX can be driven directly through the NSX UI.
 - Typically, end users tie network virtualization to their cloud management plane for deploying applications.
 - Integration is also available through OpenStack (Red Hat, VMware Integrated OpenStack, and so on), Kubernetes, and Tanzu Application Service.

The management plane performs all operations. These operations include create, read, update, and delete (CRUD).

2-13 Management and Control Planes

In NSX, the management plane and control plane are part of a single NSX management cluster:

- The management plane provides the REST API and web-based UI interface for all user configurations.
- The control plane manages computing and distributing the network runtime state.



2-14 About the NSX Management Cluster


The NSX management cluster is formed by a group of three NSX Manager nodes for high availability and scalability.

The NSX Manager appliance has the built-in policy, manager, and controller roles:

- The management plane includes the policy and manager roles.
- The central control plane (CCP) includes the controller role.

The desired state is replicated in the distributed persistent database, providing the same configuration view to all nodes in the cluster.

The NSX Manager appliance is available in different sizes for different deployment scenarios.



NSX Manager is a standalone appliance. It includes the manager, controller, and policy roles. As a result of this integrated approach, users do not need to install the manager, controller, and policy roles as separate VMs.

The diagram shows that the manager and controller instances run on all three nodes and provide resiliency. Three manager nodes can handle requests from users through the API or UI, resulting in shared workloads and efficiency.

Although the three services are merged on each node in the cluster, separate resources (CPU, memory, and so on) are allocated for each of the services.

The distributed persistent database runs across all three nodes, providing the same configuration view to each node. A manager or controller running on one node has the same view of the configuration topology as managers or controllers running on the other two nodes.

NSX Manager is available in different sizes for different deployment scenarios:


- A small appliance for lab or proof-of-concept deployments
- A medium appliance for deployments up to 64 hosts
- A large appliance for customers who deploy to a large-scale environment

For Information, see VMware Configuration Maximums at https://configmax.vmware.com.

2-15 NSX Management Cluster with Virtual IP Address

The NSX management cluster is highly available. It is configured in the following way:

- All managers must be on the same subnet.
- One manager node is elected as the leader.
- The cluster's virtual IP address is attached to the leader manager.
- Traffic is not load balanced across the managers while using VIP.
- The cluster virtual IP address is used for traffic destined for NSX Manager nodes.
- Traffic destined for any transport node uses the management IP of the node.
- A single virtual IP address is used for API and GUI client access.



The API and GUI are available on all three manager nodes in the cluster. When a user request is sent to the virtual IP address, the active manager (the leader that has the virtual IP address attached) responds to the request. If the leader falls, the two remaining managers elect a new leader. The new leader responds to the requests sent to that virtual IP address.

Load-balancing requests and traffic are not balanced across managers while using VIP.


If the leader node that owns VIP fails, a new leader is elected. The new leader sends a GARP request to take the ownership of the VIP. The new leader node then receives all new API and UI requests from users.

The diagram shows an administrator's perspective, where a single IP address (the virtual IP address) is always used to access the NSX management cluster.

2-16 NSX Management Cluster with Load Balancer

A load balancer provides high availability to the NSX management cluster:

- All nodes are active.
- GUI and API are available on all managers.
- Traffic to the virtual IP address is load balanced to multiple manager nodes.
- The manager nodes can be in different subnets.



The diagram shows how a traditional load balancer can balance the traffic across multiple manager nodes.

2-17 About the NSX Policy

The policy role performs several functions:


- Provides a centralized location for configuring networking and security across the environment
- Enables users to enter the intended configuration in the NSX UI
- Enables users to specify the final desired state of the system without being concerned about the current state or underlying implementation



2-18 About NSX Manager

NSX Manager performs several functions:


- Installs and prepares the data plane components
- Receives and validates the configuration from NSX policy
- Publishes the configuration to the CCP
- · Retrieves the statistical data from the data plane components



2-19 NSX Policy and NSX Manager Workflow

The components of an NSX Manager node Interact with each other:

- Reverse proxy is the entry point to NSX Manager.
- The policy role manages all networking and security policies and enforces them in the manager role.
- Proton is the core component of the NSX Manager node. Proton manages various functionalities such as logical switching, logical routing, distributed firewall.
- Both NSX policy and Proton persist data in CorfuDB.



Reverse proxy has authentication and authorization capabilities.

NSX Policy Manager and Proton are Internal web applications that communicate with each other through HTTP.


CorfuDB is a distributed persistent in-memory object store. Persistence is achieved by writing each transaction in a shared transaction log file. Queries are served from memory and provide better performance and scalability.

2-20 About NSX Controller

NSX Controller maintains the realized state of the system and configures the data plane.

The main functions of NSX Controller include:

- Providing control plane functionality, such as logical switching, routing, and distributed firewall
- Computing all ephemeral runtime states based on the configuration from the management plane
- Disseminating topology information reported by the data plane elements
- Pushing stateless configurations to forwarding engines




2-21 Control Plane Components (1)

In NSX, the control plane is divided into the CCP and local control plane (LCP).

The CCP exists as part of the NSX Manager nodes and is offered by the NSX Controller role.


The LCP exists on host transport nodes or on NSX Edge transport nodes.



2-22 Control Plane Components (2)

The CCP and LCP perform different functions:

- The CCP:
 - Computes the ephemeral runtime state based on the configuration from the management plane
 - Disseminates information reported by the data plane elements by using the LCP
- The LCP:
 - Monitors the local link status
 - Computes local ephemeral runtime states based on updates from the data plane and the CCP
 - Pushes stateless configurations to forwarding engines



The CCP computes and disseminates the ephemeral runtime state based on the configuration from the management plane and topology information reported by the data plane elements. The LCP runs on the compute endpoints. It computes the local ephemeral runtime state for the endpoint based on updates from the CCP and local data plane information. The LCP pushes stateless configurations to forwarding engines in the data plane and reports the information back to the CCP. This process simplifies the work of the CCP significantly and enables the platform to scale to thousands of different types of endpoints (hypervisor, container host, bare metal, or public cloud).

The NSX-RPC messaging protocol is a messaging solution for all communications between the management plane, CCP, and data plane.


Remote procedure call (RPC) is a protocol that one program can use to request a service from another program on another computer without understanding the network details.

2-23 Control Plane Change Propagation

The CCP receives the configuration information from NSX Manager and propagates the information to the LCP of the transport nodes.

The LCP on each transport node Interacts with the CCP.

If a change occurs, the LCP on the transport node notifies its assigned CCP, which further propagates these changes to the transport nodes.




The LCP on the transport node reports local runtime changes to its designated CCP node. The designated CCP node receives the changes and propagates the changes to other controllers in the cluster. All controllers propagate the changes to the transport nodes that they manage.

2-24 Control Plane Sharding Function

The NSX management cluster includes a three-node CCP.

The control plane uses a sharding mechanism to distribute workloads:


- Each transport node is assigned to a controller for L2 and L3 configuration and distributed firewall rule distribution.
- Each controller receives configuration updates from the management and data planes but maintains only the relevant information on the nodes that it is assigned to.



2-25 Managing Controller Failure

When a controller falls, its load is redistributed:

- The sharding table is recalculated to redistribute the load among the remaining controller nodes.
- This recalculation provides high availability and dual-active prevention.
- The traffic in the data plane continues to flow without being affected.




In the diagram, controller 3 is assigned to two transport nodes. When controller 3 falls, the nodes are moved to controllers 1 and 2.

2-26 About the Data Plane

The data plane has several components and functions:

- Includes multiple endpoints (ESXI hosts, bare-metal servers, and NSX Edge nodes)
- Contains various workloads, such as VMs, containers, and applications running on bare-metal servers
- Forwards data plane traffic
- Uses a scale-out distributed forwarding model
- Implements logical switching, distributed and centralized routing, and firewall filtering



2-27 Data Plane Functions

The data plane forwards packets based on configurations populated by the control plane and reports topology information to the control plane.

The data plane has the following responsibilities:

- Maintains the status of and manages fallover between multiple links or tunnels
- Performs stateless forwarding based on tables and rules populated by the control plane
- Maintains packet-level statistics

2-28 Data Plane Components

Types of data plane components, called transport nodes, include:

- Hypervisor transport nodes:
 - Act as forwarding plane for the VM and container traffic
 - Support ESXI hypervisors
- Bare-metal transport nodes:
 - Include Linux-based and Windows workloads running on bare-metal servers
 - Include containers running on bare-metal servers without a hypervisor
- NSX Edge cluster:
 - Contains edge transport nodes (VM or bare-metal form factors)
 - Provides stateful and gateway services




Data Plane

2-29 Data Plane Communication Channels

Appliance Proxy Hub (APH) acts as a communication channel between NSX Manager and the transport node.

APH offers the following functions:

- Runs as a service on NSX Manager
- Provides secure connectivity based on NSX-RPC
- Uses port 1234 for communication between the management plane and transport node
- Uses port 1235 for communication between the CCP and transport node



The NSX Manager management plane communicates with the transport nodes by using APH Server over NSX-RPC/TCP through port 1234.

CCP communicates with the transport nodes by using APH Server over NSX-RPC/TCP through port 1235.

The NSX-Proxy on the transport node receives the NSX-RPC messages from NSX Manager and CCP.

The nsx-appl-proxy APH service runs on NSX Manager.

2-30 Review of Learner Objectives

- Describe the purpose of VMware Virtual Cloud Network and its framework
- Identify the benefits and recognize the use cases for NSX
- Describe how VMware NSX fits into the NSX product portfolio
- Recognize features and the main elements in the NSX architecture
- Enumerate the deprecated features in NSX 4.0.x
- Describe the NSX policy and centralized policy management
- Describe the NSX management cluster and the management plane
- Identify the functions of control plane components, data plane components, and communication channels

2-31 Key Points

- VMware Virtual Cloud Network is a ubiquitous software layer that connects and protects any workload across any environment.
- The NSX family is a portfolio of various offerings, including NSX, vRealize Network Insight, NSX Cloud, NSX Intelligence, NSX Distributed IDS/IPS, NSX Advanced Load Balancer, Tanzu Service Mesh Advanced. VMware SD-WAN, and VMware HCX.
- In an NSX management cluster, each node performs the management, control, and policy roles.
- NSX policy provides consistency in networking and security configuration across the NSX environment.
- The data plane in NSX forwards packets based on tables populated by the control plane and reports topology information to the control plane.

Questions?

Module 3

Preparing the NSX Infrastructure

3-2 Importance

As the network administrator, you must plan and deploy a network infrastructure that meets the business requirements and growth of users and applications. You must thoroughly understand the function and configuration of the NSX management cluster and transport nodes to ensure a fully prepared environment for NSX.

3-3 Module Lessons


- Deploying the NSX Management Cluster
- Navigating the NSX UI
- Preparing the Data Plane
- DPU-Based Acceleration for VMware NSX

3-4 Lesson 1: Deploying the NSX Management Cluster

3-5 Learner Objectives

- Explain the deployment workflows for the NSX infrastructure
- Deploy NSX Manager on ESXI hosts
- Verify the deployment status of NSX Manager nodes and the NSX management cluster
- Explain the procedure to import and replace certificates

3-6 Implementing NSX in vSphere



The vSphere environment infrastructure preparation workflow is as follows:

- Deploy an NSX Manager node from an OVF template.
- Access the NSX UI.
- Register vCenter Server with NSX Manager.
- Deploy additional NSX Manager nodes to form an NSX management cluster.
- Preconfigure transport nodes, including transport zones, IP pools, and uplink profiles.
- Prepare ESXI hosts as transport nodes.
- Deploy the NSX Edge node (VM or bare metal):
 - Transport node preparation
 - NSX Edge cluster formation

3-7 Considerations for Deploying NSX Manager

You should consider the following factors before deploying NSX Manager instances:

- NSX Manager can be deployed on ESXI hosts managed by vCenter Server or on standalone ESXI hosts.
- Automated deployment of additional NSX Manager instances by using the UI or API is supported only on ESXI hosts managed by vCenter Server.
- Implementing NSX Manager in KVM is no longer supported.

NSX Manager combines the roles of the policy, manager, and controller in a single node (virtual appliance).

For supported ESXI hypervisor versions, see VMware Product Interoperability Matrix at https://www.vmware.com/resources/compatibility/sim/interop_matrix.php.

3-8 NSX Manager Node Sizing

NSX Manager supports the small, medium, and large form factors.


Form Factor	Number of CPUs	Memory (GB)	Hard Disk (GB)
Small	4	16	300
Medium	6	24	300
Large	12	48	300

NSX Manager is available in different sizes for different deployment scenarios:

- A small appliance for lab or proof-of-concept deployments
- A medium appliance for deployments with up to 64 hosts
- A large appliance for customers who deploy large-scale environments

3-9 Deploying NSX Manager from an OVF Template

In vSphere, use the OVF Template to deploy the first NSX Manager Instance.



You can select an OVF template from a remote URL or your local file system.

You can select the size of NSX Manager to be Small, Medium, or Large.

In NSX 4.0.0.1, you can deploy NSX Manager using an IPv6 address.

You customize the OVF template with the NSX Manager configuration. The Customize Template step has several sections:

- Application: Configure the passwords for the GRUB root user, system root user, admin user, and audit user.
- Network Properties: Configure the host name, IPv4, and IPv6 network configurations.
- DNS: Configure your DNS server information.
- Services Configuration: Configure NTP server and SSH settings.

By default, any password must meet the following requirements:

- Minimum of 12 characters in length
- At least one uppercase character
- At least one lowercase character
- At least one numeric character
- At least one special character
- At least five unique characters

3-10 Methods to Access NSX Manager


After you install and configure NSX Manager, you can access and manage it in the following ways:

- Log In to the NSX UI through a supported browser.
- Log In to NSX Manager through the CLI.
- Access NSX Manager with API.

For information about NSX Manager browser support, see NSX Installation Guide at https://docs.vmware.com/en/VMware-NSX/4.0/installation/GUID-AECA2EE0-90FC-48C4-8EDB-66517ACFE415.html.

3-11 Accessing the NSX UI

You enter the FQDN or IP address of the newly deployed NSX Manager Instance, Log in with the admin user name and password.



3-12 Accessing the NSX CLI

You can access the NSX CLI in the following ways:

- Open an SSH session and enter the admin user credentials that were configured during the NSX Manager installation.
- You can also use the NSX Manager virtual machine console to access the NSX CLI.

The NSX CLI is also available on all types of transport nodes and provides a consistent view of the NSX configuration across the environment.

Use the list command to retrieve all available commands to query and configure the environment.

```
_ O X
                                                          sa naxmer 01
@ Weing usermane "admin".
 apt lugam: Fri Det 14 U8:57:55 2022 Jrum 172.10.10.80
MSX CLI (Manager, Policy, Controller 6.0.1.3.0.20598732). Freze 7 for command list or enter: help
ea-nexpor-01> list.
olear auth-policy vidm enabled
clear acts policy wide it extern combled
-luar baccont
plear cluster apt-service redirect-bost
clear cluster vip
clear cluster vip all
pleas bluster wie apve
plear logging-servers
clear hame-servers
olear password-complexity
clear service http redirect-hast
clear service install-upgrade enabled
olear service manager logging-level
clear dervice hig start on boot
clear service samp start-om-boot
clear ash conc-login
olear daer dusername> password-expiration
oldar uses (username) password expiration weening
copy core-dusp ofalspath> url durl>
comp file ofilenames not curls
copy url <url> [file <filename>]
copy ari casis thumbyrint cthumbgrints [file <filenames]
destrivers cluster
del all capture sessions
del capture session seession-id-argo
del core-dump (allicopre-dump-file-appy)
del file Chilepatterny
del logging-server (hostname-or-ip-address[:port]) proto oproto) level (level) [famility (famility)] [messageid (messageid)]
del mane-servey ctm-eddresso
del ntp-server (hostname-or-in-address)
dol route grafia (profix) (gateway (gateway ig)) [interface (interface mass)] [motate (ig route metric)]
ciol magnety-describe distracted
```

You can access the CLI mode of NSX Manager either through the NSX Manager virtual machine console or by connecting remotely over SSH by using PuTTY.

You can access the NSX CLI mode through SSH only if the SSH mode is enabled in the NSX Manager appliance. By default, the following levels of access are available when connecting through the CLI mode:

- Root: This role can view, deploy, and configure any component in NSX by using the Linuxbased commands.
 - Root is an admin-level role. You must access this role only with technical support assistance.
- Admin: This role can view, deploy, and configure any component in NSX by using the CLI commands.
- Audit: This role can view settings, events, and reports. This role is read-only.

To access the admin mode, log in as the admin user through SSH or through the NSX Manager console by using the admin credentials.

To access the audit mode:

- Start a PuTTY session or open the NSX Manager virtual machine console.
- Log in by using the audit user name and password.

Use the list command to retrieve all available commands to query and configure the environment.

For example, you can use the get command to query information, for example, the get services command.

3-13 Accessing NSX Manager with API

REST APIs are used when you cannot use the GUI or when you want to automate by using scripting or other tools.


NSX Manager provides the following functions with API:

- NSX Manager accepts API requests on TCP port 443 over HTTPS application protocol to programmatically create, retrieve, modify, or delete NSX objects.
- You can use the following commands:
 - HTTPS GET commands to read and retrieve objects
 - HTTPS PUT, PATCH, or POST to create or update objects
 - HTTPS DELETE to delete objects
- To use the NSX API, you must configure a client and verify that the required ports are open between your client and NSX Manager.

For information about the various API calls and functionality, see NSX API Guide at https://developer.vmware.com/apis/1487/nsx.

3-14 Registering vCenter Server with NSX Manager

You register vCenter Server to NSX.



You add the configuration details to register the vCenter Server system to NSX.

You turn on the **Create a Service Account** toggle for features, such as vSphere Lifecycle Manager, which must authenticate with NSX APIs. During registration, the compute manager creates a service account.

The Enable Trust feature for vCenter Server 7.0 or later enables vCenter Server to execute tasks on NSX Manager.

The following access level options are available:

- Full Access to NSX: This access level ensures that vSphere with Tanzu and vSphere Lifecycle Manager can communicate with NSX. This access level is selected by default.
- Limited Access to NSX: This access level ensures that vSphere Lifecycle Manager can communicate with NSX.


3-15 Verifying vCenter Server Registration with NSX Manager

You verify that vCenter Server is successfully registered and the connection status appears as Up.



3-16 Deploying Additional NSX Manager Instances (1)

You can deploy the second and third NSX Manager nodes from the NSX UL



The screenshot shows the process for automatically deploying additional NSX Manager instances from the NSX UI.

To automatically deploy additional NSX Manager Instances from the NSX UI, you must register a vCenter Server to NSX.

3-17 Deploying Additional NSX Manager Instances (2)



For information about manually joining the NSX Manager nodes to form a cluster, see NSX Installation Guide at https://docs.vmware.com/en/VMware-NSX/4.0/installation/GUID-9F3C8273-FA5F-41C8-85CA-436F8D34977D.html.

3-18 Management Cluster Status: GUI (1)

The NSX management cluster is formed automatically.

You can view the status of the nodes in the cluster from the NSX UI.



You can check the status of nodes by selecting **Home > Monitoring Dashboard > System** in the NSX UI.

Different colors indicate different messages. Red indicates degraded performance, for example, memory usage of NSX Manager is consistently higher than 90 percent for the past 5 minutes.

3-19 Management Cluster Status: GUI (2)

On the Overview page in the NSX UI, you can view the status of the management cluster and its nodes.



The following parameters describe the NSX Manager details:

- Version: Describes the current running version of NSX Manager.
- System Load: Includes CPU, memory, and storage usage details.

The following functions are performed by the most relevant operational components:

- Cluster_Boot_Manager: Cluster Boot Manager is responsible for orchestrating the clustering services on the NSX Manager nodes.
- Datastore: Storage containers for files.
- Controller: Controller is an advanced distributed state management system that provides control plane functions for NSX logical switching and routing functions.
- Manager: Manager provides a GUI and REST APIs for creating, configuring, and monitoring NSX components such as logical switches, logical routers, firewall, and so on:
- HTTPS: The main HTTPS endpoint where API is involved. The component distributes incoming calls among components of NSX Manager.

3-20 Configuring the Virtual IP Address

You can manually configure a virtual IP address for the NSX management cluster. You access the GUI by using the configured virtual IP.



You can configure a virtual IP address for the management cluster to provide the availability among the management nodes:

- The virtual IP address is not set by default.
- When assigning Virtual IP, all the NSX Manager VMs in the cluster must be configured in the same subnet.
- The leader node of the NSX management cluster assumes ownership of the VIP to service any API and UI request. Any API and UI request coming in from clients is directed to the leader node. If the leader node that owns VIP becomes unavailable, NSX elects a new leader. The new leader owns the VIP.
- You might need to wait a few minutes for the newly configured address to take effect.

To add the virtual IP address, click SET VIRTUAL IP.

3-21 Management Cluster Status: CLI (1)

Run get cluster status to query the NSX management cluster status.

```
es-manage-Cio det cineter etamma
Tue Bep 27 2023 UTC 00(00120,000
Cinetes Id) Ffcf67es-f606-6002-0008-601562890061
Overell Status: STAGE
Broug Type: DataStone
Droop Finture STABLE
                                                                                                                                          STATUS
                                                                                                                   172,20,10,41
     19901012 0021 Work Your 28ch03800101
                                                           on number 61
                                                           se namer CV. -class level se namer CV. -class.level
     Plensidi 1819 Stor Wibe confectorary
bestites but them byth serves that a
                                                                                                                   272.20.20.32
                                                                                                                   179,30,10,55
Group Type: CLUSTER_SOUT_REPAIRS
Cabuy Francis Tintes
Buckey days
                                                           PECH
                                                                                                                                         EVALUE
     00112
     SYPHIGH -021 dies ber Zeebisschel - er manne Ci
Zitzebil 1615 fürz für erführtetigt - er manne Ci
                                                                                                                  192,20,10,81
    Figure 1617 for fine entrustricty on names UZ. volume.local before the total fire attractional on names UZ. volume.com
                                                                                                                  179.20.10.42
Group Type: COMPRISABLE
Druge Flatture Flance
Name of Co.
     101120
                                                           FOOM
                                                                                                                                         WEATHER.
     13791402-0844-4823-4504-9-2608145708
                                                                                                                    177.20.10.81
                                                           ля-плания-Cl
                                                                                                                                         100
     flerofth-oaff-4524-8555-444843838845
                                                           se-names-02, volume, local
                                                                                                                    177,20,15,45
```

You connect to an appliance in the cluster and enter the get cluster status command. The number and status of the nodes in the cluster appear.

The example output lists the Cluster Boot Manager, controller, and manager groups. It also shows each group's status with its members and member status.

3-22 Management Cluster Status: CLI (2)

Run get services to query the NSX Manager services status.

```
sa-naxmgr-01> get services
Tue Sep 27 2022 UTC 09:05:51.746
Service name:
                                         applianceproxy
Service state:
                                        running
Service name:
                                         async replicator
Service state:
                                         running
Logging level:
                                         info
Service name:
                                         auth
                                         running
Service state:
Logging level:
                                         info
Service name:
                                         cluster manager
Service state:
                                         running
Service name:
                                        cm-inventory
Service state:
                                        running
Service name:
                                        controller
Service state:
                                        running
Listen address:
```

Run get interfaces to get information about the NSX Manager network interfaces.

```
sa-naxmgr-01> get interfaces
Tue Sep 27 2022 UTC 09:09:10.359
Interface: eth0
 IPv4 Address:
   Address: 172.20.10.41/24
   MAC address: 00:50:56:06:ad:af
   MTU: 1500
   Default gateway: fd01:0:106:3:0:78ff:fe6e:5
   Broadcast address: 172,20,10,255
 IPv6 Addresses:
   - Address: fd01:0:106:3:0:78ff:fe6e:45/64
   - Address: fe80::250:56ff:fe06:ada8/64
 Link status: up
 Admin status: up
 RX packets: 3393029
 RX byses: 760449004
  RX errors; 0
 EX dropped: 599
 TX packets: 2346584
 TX bytes: 22819189293
 TX errors: 0
 TX dropped: 0
```

The following common misconfigurations might exist when troubleshooting an NSX Manager node:


- ESXI host with insufficient resources (CPUs, memory, or hard disk)
- Incorrect network details, such as the gateway address, network mask, DNS, and so on

3-23 Replacing Self-Signed Certificates (1)

After you Install NSX, the manager nodes and cluster have self-signed certificates.

Perform the following tasks if your security requirements do not include the use of self-signed certificates:

- Replace the self-signed certificates with CA-signed certificates.
- Use a different certificate for each node.



NSX Manager requires a signed certificate to authenticate the identity of the NSX Manager web service and encrypt information sent to the NSX Manager web server. As a security recommendation, you use CA-signed certificates instead of self-signed certificates.

To replace the self-signed certificates with CA-signed certificates:

- Generate a certificate signing request (CSR).
- You send the CSR file to a certificate authority (CA) to apply for a digital identity certificate.
- Import the certificate in NSX Manager.
- Replace the self-signed certificates with CA-signed certificates

To generate a CSR:

- Log in with admin privileges to an NSX Manager instance.
- Select System > Settings > Certificates.
- Click the CSRs tab.

- Select GENERATE CSR > GENERATE CSR and enter the CSR details.
- Click SAVE.

To import a CA-signed certificate:

- Log in with admin privileges to an NSX Manager instance.
- Select System > Settings > Certificates.
- Select Import > Certificate and enter the certificate details.
- Click SAVE.

3-24 Replacing Self-Signed Certificates (2)

You can replace the certificate for a manager node or the manager cluster virtual IP (VIP) by making an API call:

To replace the certificate of a manager node, use the POST API call:

```
https://<nsx-
mgr>/api/v1/node/services/http?action=apply_certificate&cert
ificate id=<certificate id>
```

To replace the certificate of the manager cluster VIP, use the POST API call:

```
https://<nsx-mgr>/api/v1/cluster/api-
certificate?action=set_cluster_certificate&certificate_id=<c
ertificate id>
```

After importing the CA-signed certificate in NSX Manager, replace the certificate:

- In your browser, log in with admin privileges to an NSX Manager instance at https://<nsx-manager-lp-address>.
- Select System > Settings > Certificates.
- 3. Expand the certificate to show its details and copy the certificate ID.

Ensure that the Service Certificate option was set to No when this certificate was imported.

 To replace the certificate of a manager node, use the POST /api/v1/node/services/http?action=apply_certificate API call.

```
For example, POST https://<nsx-
mgr>/api/v1/node/services/http?action=apply_certificate&cert
ificate_id=e61c7537-3090-4149-b2b6-19915c20504f
```

 To replace the certificate of the manager cluster VIP, use the POST /api/v1/cluster/apicertificate?action=set cluster certificate API call.

For example, POST https://<nsx-mgr>/api/v1/cluster/apicertificate?action=set_cluster_certificate&certificate_id=d6 0c6a07-6e59-4873-8edb-339bf75711ac

This step is not necessary if you did not configure VIP. For more information, see NSX REST API at https://developer.vmware.com/apis/1487/nsx.

3-25 Lab 1: Reviewing the Lab Environment and Topologies

Review the lab environment and network topologies:

- Use the Lab Environment
- Review the Networking Topologies

3-26 Lab 2: Reviewing the Configuration of the Predeployed NSX Manager Instance

Verify the NSX Manager appliance settings:

- Access the Lab Environment.
- Prepare for the Lab
- Verify the Licensing for vCenter Server and ESXI Hosts
- Verify the NSX Manager Configuration and Licensing
- 5. Use the NSX CLI to Review the NSX Management Cluster Information
- Register vCenter Server as a Compute Manager

3-27 Lab 3: (Simulation) Deploying a Three-Node NSX Management Cluster

Deploy a three-node NSX Management cluster from the NSX UI;

- Prepare for the Lab
- Deploy the Second NSX Manager Instance
- Deploy the Third NSX Manager Instance
- Configure the Virtual IP Address
- Review the NSX Management Cluster Information from the NSX CLI

3-28 Review of Learner Objectives

- Explain the deployment workflows for the NSX infrastructure
- Deploy NSX Manager on ESXI hosts
- Verify the deployment status of NSX Manager nodes and the NSX management cluster
- Explain the procedure to import and replace certificates

3-29 Lesson 2: Navigating the NSX UI

3-30 Learner Objectives

- Distinguish between the Policy and the Manager UI
- Navigate the Policy and the Manager UI main configuration window

3-31 NSX Manager Policy and Manager Views



The NSX Manager Interface provides the following modes for configuring resources:


- Policy mode:
 - Recommended mode for configuring NSX
 - Default mode for new installations
 - Supported by Federation Global Manager
 - Does not support some features
- Manager mode

Use the Manager mode UI in the following instances:

- When NSX is integrated with cloud management platforms, for example, VMware Aria Automation (former name: vRealize Automation), VMware Integrated OpenStack, and so on.
- To work with objects that are not available in the Policy mode.

3-32 User Interface Preferences for Policy and Manager Modes

By default, new installations display the UI in Policy mode, and the **UI Mode** toggle is hidden. You can change the visibility of the UI mode toggle under **System > Settings > General Settings > User Interface**.



Environments that contain objects created through the Manager mode, such as from NSX upgrades or cloud management platforms, display the **UI Mode** toggle by default in the top-right corner of the UI.

3-33 About the Networking Tab

On the **Networking** tab, you can configure functions such as switching, routing, and layer 3 services. Layer 3 services include NAT, VPN, load balancing, and so on. A Policy view and a Manager view are available.




A segment in the Policy view is called a logical switch in the Manager view.

The Tier-0 or Tier-1 gateways in the Policy view are called Tier-0 or Tier-1 logical routers in the Manager view.

3-34 About the Security Tab

On the Security tab, you can create firewall policies and endpoint security policies.




3-35 About the Inventory Tab

On the **Inventory** tab, you can review information about services, groups, VMs, containers, physical servers, and context profiles.



3-36 About the Plan & Troubleshoot Tab

On the Plan & Troubleshoot tab, you can select IPFIX, Port Mirroring, Traceflow, and Consolidated Capacity for monitoring and troubleshooting.



The Manager view does not have the Discover & Plan functionality, which is only found in the policy view.

3-37 About the System Tab

On the **System** tab, you can deploy the transport node and management cluster, add licenses, register compute managers, and so on:



The Overview page shows the number and details of the management nodes and the cluster.

The System tab does not have separate Policy and Manager views.

3-38 Review of Learner Objectives

- · Distinguish between the Policy and the Manager UI
- Navigate the Policy and the Manager UI main configuration window

3-39 Lesson 3: Preparing the Data Plane

3-40 Learner Objectives

- Describe the functions of transport zones, transport nodes, N-VDS, and VDS.
- Explain the relationships between transport nodes, transport zones, N-VDS, and VDS.
- Create IP pools for the node address assignment
- Describe and configure uplink profiles
- Prepare ESXI hosts to participate in NSX networking
- Verify the status of ESXI transport nodes and VIB installation

3-41 Data Plane Components and Functions

The NSX data plane contains several types of endpoints with different functions:

- Performs stateless forwarding and encapsulation or decapsulation of packets based on tables populated by the control plane.
- Uses a scale-out distributed forwarding model and carries data over designated transport networks in the physical network.
- Performs logical switching, distributed and centralized routing, and firewall filtering.



Various Types of Endpoints Supported in the Data Plane

3-42 Overview of the Transport Node

NSX requires transport nodes to perform networking (overlay or VLAN) and security functions.

A transport node is responsible for forwarding the data plane traffic that originates from VMs, containers, or applications running on bare-metal servers.

NSX supports the following types of transport nodes:

- ESXI Hypervisor
- Bare-metal server (RHEL, CentOS, Ubuntu, SLES, Oracle Linux, Windows)
- NSX Edge

3-43 Transport Node Components and Architecture

Each transport node contains:

- Virtual distributed switch: A vSphere Distributed Switch (VDS) or virtual distributed switch managed by NSX (N-VDS). It is the core data plane component on the transport nodes.
 - VDS is the only virtual switch supported on ESXI hosts.
 - N-VDS is the virtual switch supported on NSX Edge nodes and bare-metal workloads.
- NSX-Proxy: It is an agent running on all transport nodes that receives configuration and control plane data from CCP.

ESXi Transport Node




Since NSX 4.0.0.1, the N-VDS host switch is not supported for ESXI hosts. In a brownfield environment, you must migrate the ESXI hosts to VDS before upgrading to NSX 4.0.0.1.

3-44 Physical Connectivity of a Transport Node

For the physical connectivity of a transport node, you can select one of these options:

- Use dedicated physical NICs for management and transport (overlay or VLAN) traffic.
- Share the physical NIC for both management and transport traffic.



3-45 About IP Address Pools

Each transport node has a tunnel endpoint (TEP). The tunnel endpoint (TEP) enables transport nodes to participate in the NSX overlay.

Each TEP requires an IP address. You can create one or more IP address pools to assign addresses to TEPs

One or more TEPs can be assigned to each transport node.



An IP pool is a container created for assigning IP addresses to tunnel endpoints (TEPs).

You can manually configure IP address pools.


Each transport node has a TEP. Each TEP has an IP address. These IP addresses can be in the same subnet or in different subnets, depending on the IP pools or DHCP configured for the transport nodes.

3-46 About Transport Zones (1)

A transport zone defines the span of a logical network over the physical infrastructure.

Transport nodes can participate in the following transport zones:

- Overlay:
 - Used as the Internal tunnel between NSX hosts and NSX Edge transport nodes
 - Carries Geneve-encapsulated traffic
- VLAN:
 - Used at NSX Edge uplinks to establish northbound connectivity
 - Can carry the 802.1Q tagged traffic



A transport zone defines a collection of transport nodes that can communicate with each other across a physical infrastructure over one or more interfaces (TEPs).

A transport zone can accommodate either overlay or VLAN traffic.

The VLAN transport zones are used to connect the NSX Edge uplinks and the upstream physical routers to establish north-south connectivity.

Transport nodes are ESXI hosts, NSX Edge nodes, and bare-metal nodes that participate in an NSX overlay.

3-47 About Transport Zones (2)

Transport zones determine which hosts can participate in a particular network and have the following characteristics:


- A single transport zone can have all types of transport nodes (ESXI, bare-metal servers, and NSX Edge).
- A transport zone does not represent a security boundary.
- A transport node can belong to multiple transport zones: one overlay transport zone and multiple VLAN transport zones.
- A segment can belong to only one transport zone.

3-48 Transport Node Switch Configuration

A transport node uses N-VDS or VDS to connect to the transport zone.

A transport node switch can be configured in the following ways:

- ESXI transport nodes can use only VDS switches.
- N-VDS is the virtual switch supported on NSX Edge nodes, native public cloud NSX agents, and bare-metal workloads.
- Each N-VDS or VDS switch can be associated with multiple transport zones.
- Each N-VDS or VDS switch requires its own physical NICs.



Since NSX 4.0.0.1, an N-VDS host switch is not supported for ESXI hosts.

An edge transport node only has one N-VDS attached to an overlay transport zone.


If multiple VLAN segments are backed by the same VLAN ID, only one of those segments is realized.

To validate the versions of vSphere compatible with NSX 4.0.0.1, see VMware Product Interoperability Matrix at https://interopmatrix.vmware.com/Interoperability.

3-49 About VDS

The ESXI hosts managed by vCenter Server are configured to use only VDS during the transport node preparation.

The segments from NSX Manager are realized as distributed port groups in vCenter Server.



Configuring NSX on a standalone ESXI host is no longer supported because VDS is the only supported virtual switch and an N-VDS host switch is not supported for ESXI hosts since NSX 4.0.0.1.


The VDS MTU size must be set to 1,600 bytes or greater from the vSphere Client to utilize it with NSX.

VDS can attach to only a single overlay transport zone and multiple VLAN transport zones.

3-50 About N-VDS

N-VDS is a software logical switch that provides the forwarding service on a transport node. NSX Manager creates and centrally manages each N-VDS instance.

N-VDS instances are created and distributed across NSX Edge and bare-metal transport nodes with a consistent configuration.



N-VDS performs the switching functionality on a transport node:

- N-VDS typically owns several NICs of the bare-metal transport node.
- N-VDS instances are created on a host or NSX Edge transport nodes.
- N-VDS instances configured on different transport nodes are independent.
- N-VDS has a name assigned for grouping and management.


For example, the diagram shows two N-VDS instances that are configured on the edge node and bare-metal transport nodes: one N-VDS named Lab and another N-VDS named Prod (production).

3-51 Creating Transport Zones

Transport zones have the following characteristics:

- Dictate which transport nodes and which workloads can participate in a network.
- When using ESXI transport nodes, transport zones can span one or more vSphere clusters.

When creating a transport zone, you must specify the traffic type.



Transport zones dictate which transport nodes and which workloads can participate in a network:

- The overlay transport zone is used by ESXI transport nodes, NSX Edge nodes, and baremetal transport nodes for the overlay traffic.
- The VLAN transport zone is used by NSX Edge and host transport nodes for their VLAN uplinks.

An NSX environment can contain one or more transport zones, depending on your requirements.

By default, NSX Manager has the following preconfigured transport zones:

- nsx-overlay-transportzone for the overlay traffic type
- nsx-vlan-transportzone for the VLAN traffic type

NSX does not allow VMs in different transport zones in the layer 2 network to connect. The span of a segment is limited to a transport zone. So virtual machines in different transport zones cannot be on the same layer 2 network.

3-52 Reviewing the Transport Zone Configuration

You can verify the configuration and status of your transport zones from the NSX UI.




In the example, two transport zones are created. PROD-Overlay-TZ is an overlay transport zone used by ESXI transport nodes and NSX Edge nodes. PROD-VLAN-TZ is the VLAN transport zone used by NSX Edge and ESXI transport nodes for its VLAN uplinks.

3-53 VDS Operational Modes

VDS switches can be configured in one of three modes based on performance requirements:

- Standard datapath: Configured for regular workloads, where normal workload traffic throughput is expected.
- Enhanced datapath: Configured for telecom workloads, where high traffic throughput is expected on the workloads.
- Enhanced Network Stack (ENS) Interrupt: An Interrupt-driven version of Enhanced datapath.



The Enhanced Data Path virtual switch is optimized for Network Function Virtualization (NFV), where the workloads typically perform networking functions with very demanding requirements in terms of latency and packet rate. To benefit from this mode, workloads must be compiled with DPDK and will use VMXNET3 for their vNIC. This mode is only available to ESX hypervisor (6.7 and later, recommended 6.7 U2 and later) and unavailable on NSX Edge nodes and Public Cloud Gateway.

ENS Interrupt and Enhanced datapath requires compatible NICs and CPU cores dedicated to packet processing to support Telco-type environments with high packet count and small packet size (64 bytes).

For information about identifying suitable hardware components, see VMware Compatibility Guide at https://www.vmware.com/resources/compatibility/search.php.


Telecom service providers use SDN to deploy network function virtualization (NFV), which virtualizes a physical server or other dedicated hardware.

3-54 Physical NICs, LAGs, and Uplinks

A host can have several physical ports called physical NICs. Several physical NICs can be bundled to form an aggregated link called LAG.

Uplinks are logical interfaces on VDS.

With a VDS implementation, the ESXI host uplinks must be configured to carry the overlay and VLAN traffic.



- · Physical NICs on the host: p1, p2, and p3
- . Uplinks on the VDS: u1, u2

Mappings:

- · u2 maps to p3.
- u1 maps to the LAG that includes p1 and p2.

Link Aggregation Groups (LAGs) use Link Aggregation Control Protocol (LACP) for the transport network.

Map the uplinks on NSX to uplinks on VDS and not to physical NICs directly.


In the example, logical uplink 1 is mapped to a physical LAG (including physical port p1 and p2). Logical uplink 2 is mapped to physical port p3.


3-55 About Uplink Profiles

The uplink profile is a template that defines how VDS connects to the physical network.

The uplink profile specifies the following details:

- Format of the uplinks of VDS
- Default teaming policy applied to those uplinks
- Active and standby uplinks
- Transport VLAN used for overlay traffic
- MTU of the uplinks





An uplink profile is a container of properties or capabilities that you want your network adapters to have. It allows you to consistently configure identical capabilities for network adapters across multiple hosts or nodes.

When an administrator modifies a parameter in the uplink profile, it is automatically updated in all the transport nodes following the uplink profile.

If NSX Edge is installed on bare metal, you can use the default uplink profile. The default uplink profile requires one active uplink and one passive standby uplink.

3-56 Default Uplink Profiles


Uplink profiles enable you to configure consistent capabilities for network adapters across multiple transport nodes.

You can find the default uplink profiles by navigating to System > Fabric > Profiles > Uplink Profiles in the NSX UI.



3-57 About Teaming Policies

As part of the uplink profile configuration, the teaming policy defines the uplink redundancy and failover model. Only one teaming policy can be applied to an entire VDS.




The tearning policy only defines how the NSX virtual switch balances traffic across its uplinks. The uplinks can in turn be individual pNICs or LAGs. A LAG uplink has its own hashing options. However, those hashing options only define how traffic is distributed across the physical members of the LAG uplink, whereas the tearning policy defines how traffic is distributed between NSX virtual switch uplinks.

3-58 Teaming Policy Modes

You can select a teaming policy for the new uplink profile:

- Fallover Order: Uses one active port and a list of standby ports
- Load Balanced Source: Uses a list of active ports as input
- Load Balanced Source Mac: Determines the uplink based on the source VM's MAC address



The image shows that you can specify a type of teaming policy for the uplink profile.

You can select from the following teaming policy modes:

- Failover Order: An active uplink is specified with an optional list of standby uplinks. If the
 active uplink fails, the next uplink in the standby list replaces the active uplink. No actual load
 balancing is performed with this option.
- Load Balanced Source: A list of active uplinks is specified, and each interface on the
 transport node is pinned to one active uplink based on the source port ID. This configuration
 allows use of several active uplinks at the same time.
- Load Balanced Source Mac: This option determines the uplink based on the source VM's MAC address.

The number of VTEPs on transport nodes is determined by the NIC teaming policy:

- The fallover order NIC teaming policy creates a single VTEP on transport nodes.
- The load balance source and source MAC create multiple VTEPs on transport nodes.

3-59 About LLDP

The Link Layer Discovery Protocol (LLDP) is used to advertise a device's identity, the device's abilities, and other devices connected in the same network.

LLDP provides the following benefits:

- Simplifies the use of network management tools in a multivendor environment.
- Provides accurate discovery of physical network topologies, which simplifies troubleshooting in enterprise networks

LLDP is configured as a vCenter Server distributed switch advanced property in the vSphere Client.




3-60 About Transport Node Profiles

A transport node profile captures the configuration required to create a transport node.

The transport node profile can be applied to an existing vSphere cluster to create transport nodes for the member hosts.

A transport node profile defines:

- Transport zones
- VDS switch configuration
- Uplink profile
- IP assignment
- Mapping of physical NICs



Transport node creation begins when a transport node profile is applied to a vSphere cluster.


NSX Manager prepares the hosts in the cluster and installs the NSX components on them.

Transport nodes are created based on the configuration specified in the transport node profile.

3-61 Benefits of Transport Node Profiles

Transport node profiles make deployments easier for customers by using vSphere clusters:

- Speed deployments
- Enable the faster deployment of NSX Infrastructure components
- Drive consistency and avoid manual errors
- Can be reused multiple times




3-62 Prerequisites for Transport Node Profile

Before you configure transport node profiles, several requirements must be met:

- The ESXI hosts planned to configure NSX are part of a vCenter Server Instance.
- The ESXI hosts are added in the vSphere cluster.
- The vCenter Server Instance is added to NSX Manager as a compute manager.
- The transport zone is configured.
- An IP address pool is configured, or the DHCP server is available in the network.

3-63 Creating a Transport Node Profile

You create a transport node profile with the VDS setting and define the transport zone.



When modifying a transport node profile, the clusters using that profile are immediately updated. Also, when an ESXI host is added to a cluster, the host is updated with the transport node profile attached to the cluster.

3-64 Attaching a Transport Node Profile to the vSphere Cluster

You apply a transport node profile to a vSphere cluster.




To configure NSX on ESXI hosts managed by vCenter Server in a vSphere cluster:

- Select the vSphere cluster and click CONFIGURE NSX.
- From the Transport Node Profile drop-down menu, select the transport node profile to attach to the vSphere cluster and click APPLY.

Attaching a transport node profile is required only when configuring ESXI hosts managed by vCenter Server at the cluster level.

3-65 Reviewing the ESXi Transport Node Status (1)


After the ESXI host is prepared, you verify that the Configuration Status appears as Success, the NSX version appears, and the Node Status appears as Up.



The screenshot shows that the ESXI hosts prepared for NSX, sa-esxi-01.vclass.local, and sa-esxi-02.vclass.local are automatically listed as transport nodes in the NSX UI.

3-66 Reviewing the ESXi Transport Node Status (2)

You can view the status of the host transport nodes from the NSX Manager dashboard.



You can check the status of host transport nodes in the System view of the dashboard. Point to the circle and messages appear. These messages provide details about the nodes. For example, in the screenshot, out of four nodes, two nodes are configured as transport nodes and two nodes are not configured for NSX.

The following colors define:

- Green: Indicates a healthy environment where all components are working without issues
- Red: Indicates a critical issue
- Orange: Indicates degraded performance
- Gray: Indicates not configured

3-67 Verifying the ESXi Transport Node by CLI

The NSX kernel modules are packaged in VIB files and downloaded to hosts. The kernel modules provide services such as distributed routing, distributed firewall, and so on.

```
[root@sa-esxi-01:~] esxcli software vib list | grep nsx
nsx-adf
                           4.0.1.1.0-7.0.20598730
                                                              VMware VMwareCertified
nsx-ofdagent
                          4.0.1.1.0-7.0.20598730
                                                              VMware VMwareCertified
                          4.0.1.1.0-7.0.20598730
nsx-dontext-mux
                                                             VMware VMwareCertified
nsx-opp-libs
                          4.0.1.1.0-7.0.20598730
                                                             VMware VMwareCertified
nsx-esx-datapath
                          4.0.1.1.0-7.0.20598730
                                                             VMware VMwareCertified
nsx-exporter
                          4.0.1.1.0-7.0.20598730
                                                             VMware VMwareCertified
                           4.0.1.1.0-7.0.20598730
                                                             VMware VMwareCertified
nsx-host
nax-ids
                          4.0.1.1.0-7.0.20598730
                                                             VMware VMwareCertified
nsx-monitoring
                                                             VMware VMwareCertified
                          4.0.1.1.0-7.0.26598730
DAX-DDS
                          4.0.1.1.0-7.0.20598730
                                                             VMware VMwareCertified
                          4.0.1.1.0-7.0.20598730
                                                             VMware VMwareCertified
nax-nestdb
                          4.0.1.1.0-7.0.20598730
                                                             VMware VMwareCertified
nax-netopa
                          4.0.1.1.0-7.0.20598730
                                                             VMware VMwareCertified
nsx-opsagent
nsx-platform-client
                       4.0.1.1.0-7.0.20598730
                                                             VMware VMwareCertified
                           4.0.1.1.0-7.0.20598730
                                                             VHware VMwareCertified
nsx-proto2-libs
                           4.0.1.1.0-7.0.20598730
                                                             VHware VMwareCertified
nax-proxy
nax-python-gevent
                                                             VHware VMwareCertified
                          1.3.5.pv35-19972216
nax-python-greenler
                                                                     VMwareCertified
                                                              VMware
                           0.4.14.py35-19345965
                           4.0.1.1.0-7.0.20598730
                                                              VMware VMwareCertified
nsx-python-logging
nax-python-protobuf
                          2.6.1-19195979
                                                              VHware
                                                                      VMwareCertified
                           4.0.1.1.0-7.0.20598730
                                                                      VMwareCertified
nsx-python-utils
                                                              VMware
                                                              VHware
                           4.0.1.1.0-7.0.20598736
                                                                      VMwareCertified
nsx-shared-libs
                           4.0.1.1.0-7.0.20598730
                                                              VMware VMwareCertified
                           4.0.1.1.0~7.0.20598730
nax-vdp1
                                                              VMware VMwareCertified
naxc11
                           4.0.1.1.0-7.0.26598730
                                                              VMwareCertified
```

After an ESXI host is prepared for NSX, VIBs are installed for the host to participate in networking and security operations.

The functions of the VIBs are defined as follows:

- nsx-esx-datapath: Provides NSX data plane packet-processing functionality.
- nsx-exporter: Provides host agents that report runtime state to the aggregation service.
- nsx-host: Provides metadata for the VIB bundle that is installed on the host.
- nsx-mpa: Provides communication between NSX Manager and hypervisor hosts.
- nsx-python-protobuf: Provides Python bindings for protocol buffers.
- nsx-sfhc: Service fabric host component (SFHC) provides a host agent for managing the life cycle of the hypervisor as a fabric host.
- nsxcli: Provides the NSX CLI on hypervisor hosts.

3-68 Lab 4: Preparing the NSX Infrastructure

Deploy transport zones, create IP pools, and prepare hosts for use by NSX:

- Prepare for the Lab
- Create Transport Zones
- Create IP Pools
- Prepare the ESXI Hosts

3-69 Review of Learner Objectives

- Describe the functions of transport zones, transport nodes, N-VDS, and VDS.
- Explain the relationships between transport nodes, transport zones, N-VDS, and VDS.
- Create IP pools for the node address assignment
- Describe and configure uplink profiles
- Prepare ESXI hosts to participate in NSX networking
- Verify the status of ESXI transport nodes and VIB installation

3-70 Lesson 4: DPU-Based Acceleration for VMware NSX

3-71 Learner Objectives

- Define a DPU
- Describe the use cases and benefits of DPU-based acceleration
- Explain the architectural components of a DPU
- Describe the hardware and networking configurations supported with DPUs
- Define the NSX features supported by DPUs
- Install NSX using DPUs

3-72 Traditional Infrastructure Challenges

Some of the challenges in today's environments include:

- Meeting new application requirements
- Shared hypervisor computing resources between applications and infrastructure (network, storage, and hypervisor management tasks)
- Effect of network and security service performance on applications
- Inadequate isolation between provider and tenant in multitenant environments based on bare-metal clouds



The following challenges are faced in the traditional infrastructure:

- Modern applications require speed of delivery, application availability, security of applications, and better performance
- Hypervisor computing resources are not dedicated for modern applications and create CPU and I/O bottlenecks.
- Need for computing resources increases exponentially as you move to higher speeds such as 25, 100, and 200 Gbps.
- Network and security services, such as line rate speeds or encryption, have a cost on performance because the hypervisor resources are shared with other applications and infrastructure services.
- Bare-metal clouds interest customers with demands for high-performance workloads, direct hardware access, and custom multitenant environments. They do not provide the isolation between the provider network and tenant network in SDDC environments.

3-73 Next-Generation Infrastructure with DPUs

DPUs offer the following advantages:

- Dedicated computing resources and hardware acceleration
- Full datapath offloading to achieve high throughput and low latency
- Security services enabled in the DPU without performance impact
- Enhanced observability and operations capabilities for monitoring, troubleshooting, logging, and compliance
- Isolation between tenant and provider both on ESXI and bare-metal platforms (available in future releases)




DPUs are also called SmartNICs.

Next-generation infrastructure runs NSX services directly in the SmartNiC. This architecture is known as DPU-based acceleration for NSX.

3-74 DPU-Based Acceleration Use Cases

DPU-based acceleration has the following use cases:

- Applications with high network bandwidth demand and low latency
- Security services offloading for a better performance
- Enhanced observability requirements



These are the benefits of DPU-based acceleration:

- Networking offloads for high-bandwidth applications such as video streaming applications
- Security offloads for a better performance for applications with high security or encryption demands
- Enhanced observability such as duplication and encapsulation tasks for remote L3SPAN, packet captures, and IPFIX collection that are offloaded to the DPU or SmartNIC

3-75 SmartNICs Ecosystem

These SmartNIC or DPU vendors are currently supported.




SmartNIC vendors that currently Integrate with NSX are Nvidia and Pensando (AMD).

As a part of the new generation infrastructure with SmartNICs, leading system OEMs such as Dell Technologies, HPE, and Lenovo use SmartNICs in their integrated platforms.

3-76 Architecture Changes with DPUs

DPU-based architecture includes the following changes:

- NSX and Infrastructure services, such as storage and I/O control, are offloaded to the DPU's processor.
- ESXI and NSX instances run directly in the DPUs.



The new architecture runs the infrastructure services on the SmartNIC or DPU providing the necessary separation between application workloads that run on the x86 computing platform and the infrastructure services.


SmartNIC hardware is independent from the x86 server and has a programmable pipeline with dedicated CPUs. This architecture allows network and storage virtualization to run directly on the NIC. It saves the x86 CPU cycles and improves the performance through offload.

These instances are required to run the ESXI and NSX functionalities on the new architecture with DPUs:

- Lightweight ESXI that runs on the ARM processors in the DPU
- NSX components (transport node components) that run on the DPU

3-77 SmartNIC Hardware Components


A SmartNIC or DPU has main components.



SmartNIC or DPU components:

- Flash memory to store the boot image and local storage.
- Ethernet port for management purposes.
- Two SFP ports for high-speed Ethernet (more than two ports are not supported).
- Programmable hardware accelerator packet processing pipeline for datapath offloading.
- Embedded processors with their own memory to save server x86 CPU cycles and improve network performance.
- High-speed interconnect to provide very low latency and high bandwidth communication between the programmable hardware accelerator and the SmartNIC's embedded CPUs.
- Virtualized device functions to show virtual device functions to the x86 server. Several physical and virtual devices are included (NVME, VMXNET3, and PCIE).

3-78 Network Offloading with SmartNICs



Network offloading on SmartNiCs is achieved with the following components:

- SmartNIC contains a control plane with dedicated CPUs for slow path and software fast path forwarding.
- SmartNIC hardware accelerator is a dedicated hardware for fast path forwarding, providing line-rate speeds. This high speed can be achieved in one of the following ways:
 - NSX packet processing, currently implemented in software, can use the packet processing hardware pipelines designed in SmartNlCs.
 - NSX software packet processing can be moved from the server to the SmartNIC. This
 method reduces the server CPU consumption and the associated cache and memory
 resources, which are currently shared with the VM and container workloads.
 - VMs can use passthrough VFs and use the NSX functionalities. The hardware packet processing pipeline and embedded processors can implement the NSX datapath functionalities for that traffic.

SmartNIC architecture provides the benefits of both passthrough and current NSX Enhanced Data Path with the new VMXNET3 UPTv2 drivers.

A dedicated UPTv2 VF (Virtual Functions) module implements the new UPT architecture in the ESXI.

UPTv2 Virtual Functions present the virtualized instances of the physical network adapter.

SmartNiCs traffic for emulated devices, such as standard VMXNET3 VNiCs, passes through a shared MUX switch implemented in the ESXI to forward the traffic.

SmartNIC support for VMXNET3 UPTv2 has the following advantages:

- VNICs can work in passthrough mode while preserving the functionalities of emulated VNICs such as vSphere vMotion.
- VMs can use networking hardware from multiple vendors without using vendor-specific guest drivers.


3-79 Supported NSX Features

NSX 4.0.1 supports the following features with DPUs:

- L2 and L3 overlay networking
- L2 VLAN networking
- Observability features such as packet capture, IPFIX, TraceFlow, and port mirroring
- East-west security features such as DFW, IDS, and IPS (Tech Preview)

Technology preview (Tech Preview) software is an unreleased, concept version of a VMware software, in object form only, excluding any open-source software provided with such software. The media and documentation are provided by VMware to the licensee for which the licensee is granted a use license pursuant to an agreement.

3-80 NSX Traffic Flow with DPUs




Traffic flow has the following process:

- User enables network offloading to the DPUs.
- When the first packet arrives, it is considered as a flow miss and processed at software level.
- 3. The new packet is forwarded for software slow path processing:
 - If a packet is not allowed by a rule, the packet is dropped and a flow's entry is not created.
 - If the packet is allowed, a flow entry is created.
- When the software processing successfully inserts a flow entry, it programs the flows (fastpath flows) in the DPU hardware for faster processing.

3-81 Networking Configurations: SmartNIC Only

You can configure an ESXI with a SmartNIC with the following considerations:

- Only one SmartNIC is supported per host in NSX 4.0.1.
- SmartNIC uplinks can only be attached to vSphere Distributed Switch (VDS) compatible with network offloads.
- A SmartNIC has two uplinks so that it can be associated with a maximum of two virtual switches.




SmartNIC (DPU) Only

Current uplink profiles, such as active-active, active-standby, and LAG, are supported with SmartNICs.

3-82 Networking Configurations: SmartNIC with Standard NICs

You can configure ESXI with a SmartNIC and standard NICs with the following considerations:


- Uplinks from SmartNICs and standard NICs must not be part of the same VDS.
- Uplink profiles on SmartNICs must use the uplinks from the same SmartNIC.



SmartNiC (DPU) with Standard NICs

3-83 Installing NSX with DPUs (1)

Create a vSphere 8.0 VDS version with the network offloads compatibility enabled.




Select Configure settings > Network Offloads compatibility and select a supported SmartNIC from the drop-down menu, NVIDIA BlueField and Pensando are the available options.

A vSphere 8.0 distributed switch with network offloads compatibility enabled is required to connect hosts with DPUs.

Installation is only supported in greenfield deployments.

3-84 Installing NSX with DPUs (2)

Add a host with a DPU installed to the VDS.



You can filter the hosts that are compatible by turning on the Compatible toggle.

3-85 Installing NSX with DPUs (3)


Select **Physical adapters** by navigating to the host menu **Configure** > **Networking** to validate that the DPU adapters (SmartNIC) are detected.



In the screenshot, vmnic2 and vmnic3 are DPU-backed adapters (SmartNIC).

3-86 Installing NSX with DPUs (4)

Select the cluster and click Install NSX to Install NSX in a cluster with the DPU adapters.



Select the cluster that is connected to the virtual distributed switch configured with network offloads compatibility.

3-87 Installing NSX with DPUs (5)


Verify that the DPU uplinks exist in the cluster by reviewing the NSX installation details.



Verify that the VDS uplinks are detected as DPUs.

3-88 Installing NSX with DPUs (6)

Select Host Transport Nodes from the Configuration menu and verify that the host transport nodes are DPU supported.



In the screenshot, vmnic2 and vmnic3 are DPU supported.

3-89 Review of Learner Objectives

- Define a DPU
- Describe the use cases and benefits of DPU-based acceleration.
- Explain the architectural components of a DPU
- Describe the hardware and networking configurations supported with DPUs
- Define the NSX features supported by DPUs
- Install NSX using DPUs

3-90 Key Points (1)

- The management plane, control plane, and policy functions are deployed in NSX Manager.
- You can deploy the NSX Manager nodes on ESXI hosts.
- The NSX UI has two sections called Policy and Manager.
- The ESXI hosts that are managed by vCenter Server can only use VDS during the transport node preparation.
- A transport zone defines a collection of hosts that can communicate with each other across a physical network infrastructure. Overlay and VLAN are the available types of transport zone.

3-91 Key Points (2)

- Uplink profiles enable you to configure consistent identical capabilities for network adapters across multiple hosts or nodes.
- A teaming policy applies to each VDS uplink and defines how VDS uses its uplinks for redundancy and load balancing.
- VIBs install kernel modules that run in the hypervisor kernel and provide services such as distributed routing, distributed firewall, and other capabilities.
- DPUs accelerate the network performance in the SDDC and allow customers to free computing resources from the hypervisors.
- The NSX Instance runs directly in the DPU.

Questions?

Module 4

NSX Logical Switching

4-2 Importance

To build and run layer 2 switching in NSX, you must understand the overall architecture and components that interact during logical switching. You must deploy, configure, and manage layer 2 features that are provided by NSX, such as segments, segment profiles, and Generic Network Virtualization Encapsulation (Geneve).

4-3 Module Lessons

- Overview of Logical Switching Architecture
- Configuring Segments
- Configuring Segment Profiles
- 4. Logical Switching Packet Forwarding

4-4 Lesson 1: Overview of Logical Switching Architecture

4-5 Learner Objectives

- · Describe the functions of NSX segments
- Recognize different types of segments
- Explain tunneling and the Geneve encapsulation protocol
- Describe the interaction between components in logical switching


4-6 Use Cases for Logical Switching

Traditional data center switching challenges:

- Multitenant or application segmentation
- VM mobility requirement of layer 2 everywhere
- Large layer 2 physical network sprawl because of STP issues
- Hardware memory table (MAC and FIB) limits

NSX logical switching benefits:

- Scalable multitenancy across the data center
- Ability to use the existing physical topology
- Enabling layer 2 over layer 3 infrastructure by using overlay networks
- Segments that span across physical hosts and network switches



4-7 Prerequisites for Logical Switching


Before configuring logical switching, ensure that the following conditions are met:

- The NSX management cluster must be formed, stable, and ready to use.
- The transport nodes preparation must be complete.

Transport nodes are ESXI hosts, bare-metal servers, and NSX Edge Instances that participate in NSX.

4-8 Logical Switching Terminology

Logical switching involves several concepts.



A segment, also called a logical switch, reproduces switching functionality in an NSX virtual environment. Segments are similar to VLANs. Segments segregate networks and provide network connections to which you can attach VMs. The VMs can then communicate with each other over tunnels between ESXI hosts if the VMs are connected to the same segment. Each segment has a virtual network identifier (VNI), similar to a VLAN ID. However, unlike VLANs, VNIs scale beyond the limits of VLAN IDs.

A segment contains multiple segment ports. Entities such as routers, VMs, or containers are connected to a segment through the segment ports.

Segment profiles include layer 2 networking configuration details for logical switches and logical ports. NSX Manager supports several types of switching profiles and maintains one or more system-defined default switching profiles for each profile type.

Segment profiles contain different configurations of the logical ports. These profiles can be applied at a port level or at a segment level. Profiles applied on a segment are applicable on all ports of the segment unless they are explicitly overwritten at the port level. Multiple segment profiles are supported, including IP Discovery, MAC Discovery, SpoofGuard, Segment Security, and Quality of Service (QoS).


The virtual distributed switch (VDS for ESXI transport nodes and N-VDS on bare-metal servers) is configured on each transport node, which provides layer 2 functionality within the transport zone.

4-9 About Segments (1)

A segment is a representation of a layer 2 broadcast domain across transport nodes.

VMs attached to the same segment can communicate with each other, even across transport nodes.

Each segment is assigned a virtual network identifier (VNI), which is similar to a VLAN ID.



One or more VMs can be attached to a segment. The VMs connected to a segment can communicate with each other through tunnels between hosts.

Segments are similar to VLANs. Segments separate networks from each other. Each segment has a virtual network identifier (VNI), similar to a VLAN ID.

4-10 About Segments (2)

The type of segment created on a host depends on the transport zone to which it is attached.

A transport zone defines the span of a segment.

A segment is created either in an overlay or VLAN-based transport zone.

Each overlay segment is created as NSX distributed port groups in vCenter Server.

Segment configuration changes are allowed only from the NSX UI and API.


Workloads, such as VMs and containers, are connected to the segment ports.

4-11 About Tunneling

Tunneling encapsulates the virtual network traffic data and carries it over the physical network. NSX uses Geneve to encapsulate the data traffic.

Tunnels are set up between tunnel endpoints (TEPs).

VM frames are encapsulated with Geneve tunnel headers and sent across the tunnel.



The NSX overlay network implementation is based on tunneling. It provides isolation between the underlay network (physical network) and the overlay network (virtual network). This isolation is achieved by encapsulating the overlay frame with a Geneve header.

The underlying transport network can be another layer 2 network, or it can cross layer 3 boundaries.

The transport node endpoints in an NSX overlay network are called the tunnel endpoints (TEPs):


- TEPs are the source and destination IP addresses used in the external IP header to identify the transport nodes.
- TEPs typically carry two types of traffic: VM traffic and control (health check) traffic.

4-12 About Geneve

Geneve is an IETF overlay tunneling mechanism providing L2 over L3 encapsulation of data plane packets.

The Geneve-encapsulated packets are communicated in the following ways:

- The source TEP encapsulates the VM's frame in the Geneve header.
- The encapsulated UDP packet is transmitted to the destination TEP over port 6081.
- The destination TEP decapsulates the Geneve header and delivers the source frame to the destination VM.



NSX uses a tunneling encapsulation mechanism called Geneve.

The Geneve protocol is comparable to other tunneling protocols (such as VXLAN, NVGRE, and STT) and is more flexible.

The Geneve protocol encapsulates only data plane packets.

The Geneve-encapsulated packets are communicated over standard back planes, switches, and routers:

- Packets are sent from one tunnel endpoint to one or more tunnel endpoints using unicast addressing.
- The Geneve protocol does not modify the end-user application and the VMs in which the application runs.
- The tunnel endpoint encapsulates the end-user Ethernet frame in the Geneve header.
- The completed Geneve packet is transmitted to the destination endpoint in a standard User Datagram Protocol (UDP) packet. Both IPv4 and IPv6 are supported.


- The receiving tunnel endpoint strips the Geneve header, interprets any included options, and directs the end-user frame to its destination in the virtual network.
- The Geneve specification offers recommendations on ways to achieve efficient operation by avoiding fragmentation and taking advantage of equal-cost multipath (ECMP) routing and NIC hardware offload facilities.

4-13 Geneve Header Format

The Geneve protocol offers a new approach to encapsulation and offers control plane independence between tunnel endpoints.

The Geneve header has the following features:

- Runs on UDP
- Adds an 8-byte UDP header
- Uses the port number 6081
- Uses a 24-bit VNI to identify the segment
- Is supported by tcpdump and Wireshark tools



To support the needs of network virtualization, the tunneling protocol draws on the evolving capabilities of each type of device in both the underlay and overlay networks.

This process imposes a few requirements on the data plane tunneling protocol:

- The data plane is generic and extensible enough to support current and future control planes.
- Tunnel components are efficiently implemented in both hardware and software without restricting capabilities to the lowest common denominator.
- High performance over existing IP addresses is required.

The Geneve packet format includes a compact tunnel header encapsulated in UDP over either IPv4 or IPv6. A small, fixed tunnel header provides control information, as well as a base level of functionality and interoperability with a focus on simplicity. This header is followed by a set of variable options for future development. The payload consists of a protocol data unit of the indicated type, such as an Ethernet frame.

The following fields are in a Geneve header:

- Version (2 bits): The current version number is 0.
- Options Length (6 bits): This variable results in a minimum total Geneve header size of 8 bytes and a maximum of 260 bytes.
- O (1 bit): Operations, Administration, and Maintenance (OAM) packet. This packet contains a control message instead of a data payload.
- C (1 bit): This field indicates that critical options are present.
- Rsvd. (6 bits): The Reserved field must be zero on transmission and ignored on receipt.
- Protocol Type (16 bits): The field indicates the type of protocol data unit appearing after the Geneve header.
- Reserved (8 bits): The Reserved field must be zero on transmission and ignored on receipt.
- Virtual Network Identifier: A unique VNI identifies each logical network. The VNI uniquely
 identifies the segment that the inner Ethernet frame belongs to. It is a 24-bit number that is
 added to the Geneve frame, allowing a theoretical limit of 16 million separate networks. The
 NSX VNI range is 5,000 through 16,777,216.


The base Geneve header is followed by zero or more options in type-length-value format. Each option includes a 4-byte option header and a variable amount of option data interpreted according to the type. Geneve provides NSX-T Data Center with the complete flexibility of inserting metadata in the type, length, and value fields that can be used for new features. One of the examples of this metadata is the VNI. You must use an MTU of 1600 to account for the encapsulation header.

The Geneve protocol offers the following benefits:

- Supports proprietary type, length, and value fields
- Can add new metadata to the encapsulation without revising the Geneve standard
- Allows VMware to develop software-based features without hardware dependencies
- Provides the same kind of NIC offloads as VXLAN (check compatibility list)
- Is open so that third-party tools, such as Wireshark, can decode it

4-14 Logical Switching: End-to-End Communication

Two VMs connected to the same segment communicate end to end.



The example diagram shows the following details:

- The ESXI host sa-esxi-01 is configured as a transport node using vSphere Distributed Switch (dvs-SA-Datacenter) as a host switch with TEP IP: 172.20.11.51. The VMkernel interface VMK10 is created on the ESXI host.
- The ESXI host sa-esxi-02 is configured as a transport node using vSphere Distributed Switch (dvs-SA-Datacenter) as a host switch with TEP IP: 172.20.11.52. The VMkernel Interface VMK10 is created on the ESXI host.
- The sa-esxi-01 and sa-esxi-02 host transport nodes are configured in the transport zone named PROD-OVERLAY-TZ.
- Transport node A (sa-esxi-01) is running VM-1 with IP address 10.1.10.11 and MAC address ABC.

- Transport node B (sa-esxi-02) is running VM-2 with IP address 10.1.10.12 and MAC address DEF.
- VM-1 and VM-2 are connected to the segment ports on Web-Segment 69632. This web segment is an overlay-based segment configured in the transport zone named PROD-OVERLAY-TZ
- When VM-1 communicates with VM-2, the source hypervisor encapsulates the packet with the Geneve header and sends it to the destination transport node, which decapsulates the packet and forwards it to the destination VM.

During VM-1 to VM-2 communication:

- VM-1 sends the traffic to Web-Segment.
- The source hypervisor encapsulates the packet with the Geneve header.
- The source transport node forwards the packet to the physical network.
- The destination transport node receives the packet and performs the decapsulation.
- The destination TEP forwards the L2 frame to the destination VM.

4-15 Review of Learner Objectives

- Describe the functions of NSX segments
- Recognize different types of segments
- Explain tunneling and the Geneve encapsulation protocol
- Describe the interaction between components in logical switching.

4-16 Lesson 2: Configuring Segments

4-17 Learner Objectives

- Use the NSX UI to create segments
- Attach VMs to segments
- Describe the workflow for segments creation
- Review the switching configuration using the Network Topology


4-18 Segment Configuration Tasks

To create segments and attach VMs to segments:

- Create a segment.
- Attach a VM to a segment.
- Verify the segment port status.
- Verify the layer 2 VM-to-VM connectivity.

4-19 Creating Segments

You use the NSX UI to create segments.



A segment connects to gateways and VMs. A segment performs the functions of a logical switch.


To create segments:

- From your browser, log in to NSX Manager and select Networking > Segments > NSX > ADD SEGMENT.
- In the Segment Name text box, enter the segment name.
- From the Connected Gateway drop-down menu, select one option.
 - None
 - Tier-0 Gateway
 - Tier-1 Gateway
- From the Transport Zone drop-down menu, select the transport zone.
- Enter the segment subnet in the Subnets text box.
- Click Save to save the segment configuration.

4-20 Creating Segments Workflow


To create segments:

- A user creates a segment through the NSX UI.
- The policy role pushes the configuration to the manager role.
- NSX Manager realizes the segment information as logical switches in the Corfu database.
- The manager role forwards the segment information to the CCP.
- The CCP sends the information to the APH.
- The APH service sends the switching configuration to the local control plane (nsx-proxy) over port 1235.
- The nsx-proxy agent forwards the switching configuration to the nsxt-vdl2 kernel module, which creates and configures the segments in the datapath.



4-21 Viewing Configured Segments


You can connect to vCenter Server with the vSphere Client to view the configured segments.



Segments are represented as NSX distributed port groups in vCenter Server.

4-22 Attaching a VM to a Segment

You can attach a VM to a segment.




You attach the VM to the desired segment by editing its settings in the vSphere Client.

4-23 Workflow: Attaching a VM to a Segment (1)


To attach a VM managed by vCenter Server to a segment:

- You attach a VM to a segment (Web-Segment in this example) from vCenter Server.
- vCenter Server sends the request to the ESXI host.
- The ESXI host sends the request to the management plane.




4-24 Workflow: Attaching a VM to a Segment(2)

- NSX Manager configures logical interface 1 (LIF 1) on the segment with a virtual interface (VIF 1) attachment.
- 5. NSX Manager advertises the attachment configuration to the CCP.
- The CCP sends the request to the ESXI host on which the VM resides.
- LIF 1 is created on Web-Segment with VIF 1 attached.



4-25 Verifying the Segment Port Status

You verify that the status of the segment and the status of the port (to which the VM is connected) appear as Success. VMs attached to the same segment should be able to ping each other.



After you successfully set up the segment and attach VMs to it, you can test the connectivity between VMs on the same segment. In the example, you can test the connectivity in the following way:


- Using SSH or the VM console, log in to the sa-web-01 (172.16.10.11) VM, which is attached to Web-Segment.
- Ping the sa-web-02 (172.16.10.12) VM, which resides on another ESXI host. This VM is also attached to Web-Segment.

4-26 About Network Topology

The Network Topology feature enables users to understand how the different NSX networking components are configured and interconnected.


This feature displays:

- Tier-0 and Tier-1 gateways with their attached segments and workloads
- Details for each VM, segment, and Tier-1 and Tier-0 gateways
- IP addresses configured in the network
- Underlying fabric infrastructure details of the network topology



4-27 Using Network Topology to Validate the Segment Configuration

The Network Topology tool shows the segment configuration and the VMs that are attached to segments.



If you click the individual segment, a panel opens with details about the segment, such as the subnet, traffic type, and the attached segment profiles as shown in the screenshot.

4-28 Review of Learner Objectives

- Use the NSX UI to create segments
- Attach VMs to segments
- · Describe the workflow for segments creation
- Review the switching configuration using the Network Topology

4-29 Lesson 3: Configuring Segment Profiles

4-30 Learner Objectives

- Describe the purpose of segment profiles
- Identify the functions of the segment profiles in NSX
- Create segment profiles and attach them to segments and ports

4-31 About Segment Profiles (1)

Segment profiles provide layer 2 networking configuration details for segments and ports. You can create various types of segment profiles from the NSX UL



4-32 About Segment Profiles (2)

Each type of segment profile has a different function:


Segment Profile	Function
SpoofGuard	Helps prevent NIC spoofing by authenticating the IP and MAC address of the virtual NIC
IP Discovery	Learns the VM MAC and IP addresses
MAC Discovery	Supports MAC learning and MAC address change
Segment Security	Provides stateless layer 2 and layer 3 security
Quality of Service (QoS)	Provides high-quality and dedicated network performance for preferred traffic

NSX supports several types of segment profiles and maintains one or more system-defined default segment profiles:

- SpoofGuard prevents traffic with incorrect source IP and MAC addresses from being transmitted.
- The IP Discovery profile uses DHCP snooping, Address Resolution Protocol (ARP) snooping, or VMware Tools to learn the VM MAC and IP addresses.
- The MAC Discovery profile supports two functionalities: MAC learning and MAC address change.
- Segment Security provides stateless layer 2 and layer 3 security by checking the ingress traffic to the segment and matching the IP address, MAC address, and protocols to a set of allowed addresses and protocols. Unauthorized packets are dropped.
- QoS provides high-quality and dedicated network performance for preferred traffic.

4-33 Default Segment Profiles


The system default segment profiles are not editable.



You cannot edit or delete the default segment profiles, but you can create custom segment profiles.


4-34 Applying Segment Profiles to Segments

You can apply default or user-created profiles to a segment.



4-35 Applying Segment Profiles to Segment Ports

You can apply default or custom profiles to segment ports. A segment or segment port can be associated with only one segment profile of each type.



For example, two QoS segment profiles cannot be associated with a segment or segment port.

When the segment profile is associated or disassociated from a segment, the segment profile for the child segment ports is applied based on the following criteria:

- If the parent segment has a profile associated with it, the child segment port inherits the segment profile from the parent.
- If the parent segment does not have a segment profile associated with it, a default segment profile is assigned to the segment, and the segment port inherits that default segment profile.
- If you explicitly associate a custom profile with a segment port, this custom profile overrides
 the existing segment profile.

You can associate a custom segment profile with a segment and retain the default segment profile for one of the child segment ports. You must make a copy of the default segment profile and associate it with the specific segment port.

4-36 SpoofGuard Segment Profile

The purpose of the SpoofGuard profile is to prevent traffic with incorrect source IP and MAC addresses from being transmitted. SpoofGuard provides the following functions:

- Prevents a rogue VM from assuming the IP address of an existing VM
- Ensures that the IP addresses of VMs are not altered without intervention
- Ensures that distributed firewall rules are not inadvertently or deliberately bypassed

The SpoofGuard profile includes several features.

SpoofGuard Feature	SpoofGuard Function
MAC SpoofGuard	Authenticates the MAC address of a packet.
IP SpoofGuard	Authenticates the MAC and IP addresses of a packet.
Dynamic ARP Inspection	ARP and GARP SpoofGuard and ND SpoofGuard validation are checked against the MAC source, IP source, and IP-MAC source mapping in the ARP, GARP, and ND payload.

SpoofGuard provides protection against spoofing with MAC+IP+VLAN bindings. If a VM's IP address does not match the IP address on the corresponding logical port and switch address binding in SpoofGuard, the VM's vNIC is prevented from accessing the network entirely. SpoofGuard can be configured at the port or switch level.

SpoofGuard might be used in your environment for the following reasons:

- Preventing a rogue VM from assuming the IP address of an existing VM.
- Ensuring that the IP addresses of VMs cannot be altered without intervention. You might
 not want VMs to alter their IP addresses without proper change control review. You can use
 SpoofGuard to ensure that the VM owner cannot alter the IP address and continue working
 unimpeded.
- Ensuring that the distributed firewall rules are not inadvertently (or deliberately) bypassed.
 For distributed firewall rules created using IP sets as sources or destinations, a VM could have its IP address forged in the packet header, thereby bypassing the rules in question.

4-37 Creating a SpoofGuard Segment Profile

SpoofGuard ensures that address bindings are enforced in the datapath.

If the IP address of a VM changes, traffic from the VM might be blocked by SpoofGuard until the corresponding configured port-segment address bindings are updated with the new IP address.



A SpoofGuard profile applied to a segment or a port blocks traffic determined to be spoofed. When SpoofGuard is configured, if the IP address of a VM changes, traffic from the VM might be blocked until the corresponding configured port or segment address bindings are updated with the new IP address.

You can enable SpoofGuard for the port groups containing the guests. When enabled for each network adapter, SpoofGuard Inspects packets for the prescribed MAC and its corresponding IP address.

A SpoofGuard profile can be applied to a segment or a port:

- At the port level, the allowed MAC, VLAN, or IP allowlist is provided through the Address
 Bindings property of the port. When the VM sends traffic, it is dropped if its MAC, VLAN, or
 IP address does not match the MAC, VLAN, or IP properties of the port. The port-level
 SpoofGuard deals with traffic authentication, that is, the traffic consistent with VIF
 configuration.
- At the segment level, the allowed MAC, VLAN, or IP allowlist is provided through the Address Bindings property of the segment. This property is typically an allowed IP range or subnet for the segment, and the segment-level SpoofGuard deals with traffic authorization.

Traffic must be permitted by the port and the segment levels by SpoofGuard before it is allowed into a segment. Enabling or disabling port-level and segment-level SpoofGuard can be controlled using the SpoofGuard segment profile.

4-38 IP Discovery Segment Profile

The IP Discovery segment profile uses the following mechanisms to learn the MAC and IP addresses of VMs:

- DHCP/DHCPv6 snooping
- ARP snooping
- VMware Tools
- ND snooping
- Duplicate address detection

Learned addresses are shared with the CCP to achieve ARP/ND suppression.


In NSX, the IP Discovery profile works in the following ways:

- DHCP/DHCPv6 Snooping inspects the packets exchanged between the VM DHCP/DHCPv6 client and the DHCP/DHCPv6 server to learn the VM's IP and MAC addresses
- ARP Snooping inspects a VM's outgoing ARPs and GARPs to learn the IP and MAC addresses of the VM.
- The VMware Tools software runs on a VM hosted on ESXI and can provide the VM's configuration information.
- ND Snooping is the IPv6 equivalent of ARP snooping. It inspects neighbor solicitation (NS)
 and neighbor advertisement (NA) messages to learn the IP and MAC addresses.
- Duplicate address detection checks whether a newly discovered IP address is already present on the realized binding list for a different port.

The VMware Tools IP Discovery method can also provide the VM's configuration information and is available for only VMs hosted by ESXI.

The IP Discovery profile might be used in the following scenario: The distributed firewall depends on the IP-to-port mapping to create firewall rules. Without IP Discovery, the distributed firewall must find the IP of a logical port through SpoofGuard and manual address bindings, which is a cumbersome and error-prone process.

4-39 Creating an IP Discovery Segment Profile





In the screenshot, a custom IP Discovery profile named Custom- IP-Discovery-Profile is created. After the profile is created, it can be applied to various segments.

By default, the discovery methods, ARP snooping and ND snooping, operate in a mode called trust on first use (TOFU). In the TOFU mode, when an address is discovered and added to the realized bindings list, that binding remains in the realized list forever. TOFU applies to the first n unique <IP, MAC, VLAN> bindings discovered using ARP/ND snooping, where n is the binding limit that you can configure. You can disable TOFU for ARP/ND snooping. The methods then operate in trust on every use (TOEU) mode. In the TOEU mode, when an address is discovered, it is added to the realized bindings list, and when it is deleted or expires, it is removed from the realized bindings list. DHCP snooping and VMware Tools always operate in the TOEU mode.

4-40 MAC Discovery Segment Profile

The MAC Discovery profile supports MAC learning, MAC address change, unknown unicast flooding, MAC limit, and MAC limit policy functions.



The MAC Discovery profile supports source MAC address learning:

- Source MAC address-based learning is a common feature in the physical world for learning the MAC address of a machine. The MAC Learning feature provides network connectivity to deployments where multiple MAC addresses are configured behind one vNIC. For example, in a nested hypervisor deployment, an ESXI VM runs on an ESXI host and multiple VMs run in the ESXI VM.
- Without MAC Learning, when the ESXI VM's vNIC connects to a segment port, its MAC
 address is static. VMs running inside the ESXI VM do not have network connectivity because
 their packets have different source MAC addresses. With MAC Learning, the source MAC
 address of every packet coming from the vNIC is inspected, the MAC address is learned,
 and the packet is allowed to go through. If a MAC address that is learned is not used for 10
 minutes, it is removed. This aging property is not configurable.
- MAC Learning also supports Unknown Unicast Flooding. When a unicast packet is received by a port that has an unknown destination MAC address, the packet is flooded out on all segment ports that have MAC Learning and Unknown Unicast Flooding enabled. This property is enabled by default, but only if MAC Learning is enabled.

The MAC Discovery profile also supports a VM's ability to change its MAC address:

- A VM connected to a port with MAC Change enabled can run an administrative command to change the MAC address of its vNiC and still send and receive traffic on that vNiC.
- This feature (disabled by default) is used when a VM needs the ability to change its MAC address and not lose network connectivity.

The number of MAC addresses that can be learned is configurable. The maximum value is 4,096, which is the default. You can also set the policy for when the limit is reached. The options are:


- Drop: Packets from an unknown source MAC address are dropped. Packets inbound to this MAC address are treated as unknown unicast. The port receives the packets only if it has unknown unicast flooding enabled.
- Allow: Packets from an unknown source MAC address are forwarded although the address is not learned. Packets inbound to this MAC address are treated as unknown unicast. The port receives the packets only if it has unknown unicast flooding enabled.

If you enable both MAC Learning and MAC Change, you should also enable SpoofGuard to improve security.

For information about creating a MAC Discovery profile and associating the profile with a segment or a port, see NSX Administration Guide at https://docs.vmware.com/en/VMware-NSX/4.0/administration/GUID-C24D09FC-E213-4920-BD12-60063AC3B08E.html.

4-41 Segment Security Profile

The Segment Security profile provides stateless layer 2 and layer 3 security. It protects segment integrity by filtering malicious attacks from VMs in the network.



The Segment Security profile provides stateless layer 2 and layer 3 security by checking the ingress traffic to the segment and dropping unauthorized packets sent from VMs. The profile matches the IP address, MAC address, and protocols to a set of allowed addresses and protocols.

You can configure the Bridge Protocol Data Unit (BPDU) filter, DHCP snooping, DHCP server block, and rate limiting options:

- BPDU Filter: When the BPDU filter is enabled, all BPDU traffic is blocked for each port on the segment.
- BPDU Filter Allow List: You click the destination MAC address from the BPDU destination MAC addresses list to allow traffic to the permitted destination.
- To enable DHCP filtering, you turn on the Server Block and Client Block toggles. DHCP Server Block blocks traffic from a DHCP server to a DHCP client. It does not block traffic from a DHCP server to a DHCP relay agent.
- DHCP filtering can also be configured for IPv6 traffic by using the Server Block-IPv6 and Client Block-IPv6 options.
- Turning on the Non-IP Traffic Block toggle allows only IPv4, IPv6, ARP, GARP, and BPDU traffic. The rest of the non-IP traffic is blocked. The permitted IPv4, IPv6, ARP, GARP, and BPDU traffic is based on other policies set in address binding and SpoofGuard configurations. By default, this option is disabled to allow non-IP traffic to be handled as regular traffic.

- Turn on the RA Guard toggle to filter out ingress IPv6 router advertisements. The ICMPv6type 134 packets are filtered out. This option is enabled by default.
- You can configure rate limits for the ingress or egress broadcast and multicast traffic. Rate limits are configured to protect the segment or the VM from threats such as broadcast storms. To avoid any connectivity problems, the minimum rate limit value must be greater than or equal to 10 PPS.

4-42 QoS Segment Profile

The QoS profile provides high-quality network performance for preferred traffic that requires high bandwidth.

This profile supports the following methods:

- Class of Service (CoS)
- Differentiated Services Code Point (DSCP)



QoS provides high-quality and dedicated network performance for preferred traffic that requires high bandwidth. The QoS mechanism achieves this performance by providing sufficient bandwidth, controlling latency and jitter, and reducing data loss for preferred packets even with network congestion. This level of network service is provided by using the existing network resources efficiently.

The QoS profile supports the following methods:

- Class of Service (CoS): Marks the packet's layer 2 header to specify its priority
- Differentiated Services Code Point (DSCP): Inserts a code value into the packet's layer 3 header for prioritization

The layer 2 CoS allows you to specify priority for data packets when traffic is buffered in the segment because of congestion. The layer 3 DSCP detects packets based on their DSCP values. CoS is always applied to the data packet regardless of the trusted mode.

NSX trusts the DSCP setting applied by a VM or modifies and sets the DSCP value at the segment level. In each case, the DSCP value is propagated to the outer IP header of encapsulated frames. In this way, the external physical network can prioritize the traffic based on the DSCP setting on the external header. When DSCP is in the trusted mode, the DSCP value is copied from the inner header. When in the untrusted mode, the DSCP value is not preserved for the inner header. DSCP settings work only on tunneled traffic. These settings do not apply to traffic inside the same hypervisor.

You can use the QoS segment profile to configure the average ingress and egress bandwidth values to set the transmit limit rate. To prevent congestion on the northbound network links, you can use the peak bandwidth rate to specify the upper limit that traffic on a segment is allowed to burst. The settings in a QoS segment profile do not guarantee the bandwidth but help limit the use of network bandwidth. The actual bandwidth you observe is determined by the link speed of the port or the values in the segment profile, whichever is lower.

For information about the QoS segment profile, see NSX Administration Guide at https://docs.vmware.com/en/VMware-NSX/4.0/administration/GUID-62BB7145-EDD7-4611-A50D-17F4A0EAE57C.html.

4-43 Review of Learner Objectives

- Describe the purpose of segment profiles
- Identify the functions of the seament profiles in NSX
- Create segment profiles and attach them to segments and ports

4-44 Lesson 4: Logical Switching Packet Forwarding

4-45 Learner Objectives

- · Describe the functions of each table used in packet forwarding
- · Describe how BUM traffic is managed in switching
- Explain how ARP suppression is achieved

4-46 NSX Controller Tables

NSX Controller maintains a set of tables that are used for identifying the data plane component associations and for traffic forwarding.

To achieve network virtualization, a network controller must configure the hypervisor virtual segment with network flow tables that form the logical broadcast domains.


Several network flow tables are available:

- TEP Table: Maintains VNI-to-TEP IP bindings
- ARP Table: Maintains the VM MAC-to-VM IP mapping
- MAC Table: Maintains the VM MAC address-to-TEP IP mapping

4-47 TEP Table Update (1)


When a powered-on VM is connected to a segment:

The VNI-to-TEP mapping is registered on the transport nodes in its local TEP table.




4-48 TEP Table Update (2)

2. Each transport node updates the CCP about the learned VNI-to-TEP IP mapping.




4-49 TEP Table Update (3)

The CCP maintains the consolidated entries of VNI-to-TEP IP mappings.



4-50 TEP Table Update (4)


The CCP sends the updated TEP table to all the transport nodes where the VNI is realized.



4-51 MAC Table Update (1)


When a powered-on VM is connected to a segment:

The VM MAC-to-TEP IP mapping is registered on the transport nodes in its local MAC table.




4-52 MAC Table Update (2)

2. Each transport node updates the CCP about the learned VM MAC-to-TEP IP mapping.




4-53 MAC Table Update (3)

The CCP maintains the consolidated entries of VM MAC-to-TEP IP mappings.



4-54 MAC Table Update (4)

4. The CCP sends the updated MAC table to all the transport nodes where the VNI is realized.




4-55 About the ARP Table

The ARP suppression technique reduces the amount of ARP flooding in segments:

- NSX uses the ARP table maintained in the CCP to provide ARP suppression.
- Transport nodes learn the MAC-to-IP association by snooping the ARP and DHCP traffic.
- The mapping is recorded each time a VM initiates a communication (sends or receives traffic).
- The learned information is pushed from each transport node to the control plane.


4-56 ARP Table Update (1)

Each transport node records the local VM IP-to-MAC mapping in its local table.




4-57 ARP Table Update (2)

2. Each transport node sends known VM IP-to-MAC mappings to the CCP.




4-58 ARP Table Update (3)

The CCP updates its ARP table based on the VM IP-to-MAC mappings received from transport nodes.



4-59 ARP Table Update (4)

The CCP sends the updated ARP table to all the transport nodes.




The ARP table values in both the CCP and the transport nodes are flushed after 10 minutes.

4-60 Unicast Packet Forwarding Across Hosts(1)


VM1 assumes that the ARP is resolved:

VM1 starts sending traffic to VM2.




4-61 Unicast Packet Forwarding Across Hosts(2)

The original packet is encapsulated in the Geneve header by the ESXI-A source transport node.




4-62 Unicast Packet Forwarding Across Hosts(3)

The packet is sent to the ESXI-B destination transport node.



4-63 Unicast Packet Forwarding Across Hosts(4)

 The destination transport node decapsulates the Geneve header and delivers the original source VM frame to VM2




4-64 Overview of BUM Traffic

A VM's broadcast, unknown unicast, and multicast (BUM) traffic must be flooded to all other VMs that belong to the same segment.

The BUM traffic originated by a VM on a transport node must be replicated to remote transport nodes (running the VMs connected to the same segment).

The following BUM traffic replication modes are supported:

- Hierarchical Two-Tier Replication
- Head Replication



All broadcast, unicast, and multicast (BUM) traffic is treated the same: flooded to all participating ESXI hosts in the segment. The replication is performed in software.

Each host transport node has a tunnel endpoint. Each TEP has an IP address. These IP addresses can be in the same subnet or in different subnets, depending on your configuration of IP pools or DHCP for your transport nodes.

When two VMs on different hosts communicate directly and ARP is resolved, unicastencapsulated traffic is exchanged between the two TEP IP addresses without any need for flooding. However, as with any layer 2 network, sometimes traffic that is originated by a VM, such as an ARP request, needs to be flooded. For the layer 2 BUM traffic, the packet must be sent to all the other VMs belonging to the same segment.

In the diagram, Source VM residing on transport node 1 (TN1) must send traffic to Destination VM residing on TN9. The destination's VM MAC address is unknown to TN1 or the control plane. Therefore, the Source VM sends an ARP request (broadcast frame) seeking the destination VM's MAC address. TN1 floods this ARP request frame to all other transport nodes within VNI 73728. The destination VM on TN9 receives the ARP request and responds with an ARP reply. ARP tables on hosts are updated to reduce future flooding.


To enable flooding, the NSX segment supports the following types of replication modes:

- Head Replication mode: This mode is also called Source Mode or Headend Replication. The source host duplicates each BUM frame and sends a copy to each TEP (on a particular VNI) that it knows.
- Hierarchical Two-Tier Replication: This mode is also called the MTEP mode. It involves a host
 in another L2 domain that performs replication of BUM traffic to other hosts within the same
 VNI.

The TEP only replicates traffic in which the replication option TLV is set in the Geneve header.

4-65 Managing BUM Traffic: Head Replication

The Head replication mode performs source-based replication. The BUM packet is replicated by the source transport node to all other transport nodes participating in that VNI.



In this example, a BUM packet arrives on TN1:


- TN1 replicates to TN2 and TN3 because they are in the same L2 domain.
 TN1 replicates because the control plane does not have the desired information.
- Meanwhile, TN1 also needs to replicate the packet to the remote transport nodes (TN4 and TN5 in one L2 domain and TN7, TN8, TN9 in another L2 domain).
- Because TN6 does not participate in VNI 73728, the packet is not replicated to TN6.

4-66 Managing BUM Traffic: Hierarchical Two-Tier Replication

The source transport node replicates the BUM packet locally within its L2 domain.

The source transport node elects a proxy TEP (MTEP) from each remote L2 domain and sends the BUM packet to each proxy TEP.

The proxy TEP replicates the BUM packet to the transport nodes within its L2 domain.



Hierarchical two-tier mode is also known as the MTEP replication mode.

In the diagram, a BUM packet arrives on TN1:

- An MTEP is elected for each L2 domain (segment). TN1 elects an MTEP for each remote L2 domain.
- TN1 replicates the BUM traffic locally to TN2 and TN3.

TN1 sends a copy of the BUM packet to each remote MTEP with the replication option TLV
embedded in the Geneve header.

The role of MTEP is to replicate the received BUM packet locally and forward it to other TNs within the same L2 domain.

- MTEP TN7 forwards the BUM packet to TN8 and TN9.
- MTEP TN5 forwards the packet to TN4.
- Because TN6 does not participate in VNI 73728, the packet is not sent to TN6.

4-67 Lab 5: Configuring Segments

Create segments for VMs residing on the ESXI hosts:

- Prepare for the Lab
- Create Segments
- Attach VMs to Segments
- Use Network Topology to Validate the Logical Switching Configuration
- Test Layer 2 Connectivity and Verify the Configuration of Segments

4-68 Review of Learner Objectives

- Describe the functions of each table used in packet forwarding
- Describe how BUM traffic is managed in switching
- Explain how ARP suppression is achieved

4-69 Key Points

- A segment is a representation of the L2 broadcast domain across transport nodes.
- A VNI is assigned to each segment.
- Types of segments are overlay segments and VLAN segments.
- Geneve is an IETF overlay tunneling mechanism that provides L2 over L3 encapsulation of data plane packets.
- Segment profiles provide L2 networking configuration details for segments and ports.
- Five types of segment profiles are available: SpoofGuard, IP Discovery, MAC Discovery, Segment Security, and QoS.
- You can apply default or custom segment profiles to segments or ports.
- Network flow tables used in packet forwarding include TEP, ARP, and MAC table.
- BUM traffic replication supports head mode and hierarchical two-tier mode.

Questions?

Module 5

NSX Logical Routing

5-2 Importance

In NSX, logical routing provides an optimized and scalable way to manage east-west and northsouth traffic. You must understand the NSX logical routing architecture, routing components, and routing features to build an efficient and secure layer 3 network infrastructure.

5-3 Module Lessons

- Overview of Logical Routing
- NSX Edge and Edge Clusters
- Configuring Tier-0 and Tier-1 Gateways
- Configuring Static and Dynamic Routing
- ECMP and High Availability
- Logical Routing Packet Walk
- VRF Lite

5-4 Lesson 1: Overview of Logical Routing

5-5 Learner Objectives

- Explain the function and features of logical routing
- Describe the architecture of NSX two-tier routing
- Differentiate between north-south and east-west routing
- Describe the gateway components
- Recognize the various types of gateway interfaces

5-6 Use Cases for Logical Routing

In NSX, logical routing is used in many ways:

- Support for single or multitenant deployment models
- Separation of tenants and networks
- Solution for cloud environments with containerized workloads
- Optimized routing path and simplified routing in virtual networks
- Distributed routing and centralized services in data centers
- Ability to extend logical networks to physical environments

The NSX logical routing has many use cases:

- The logical routing functionality focuses on multitenant environments. Gateways can support
 multiple instances where a separation of tenants and networks is required.
- Logical routing is optimized for cloud environments, it suits containerized workloads and multicloud data centers.
- The distributed routing architecture provides optimal routing paths. Routing is done closest
 to the source. For example, traffic from two VMs on different subnets residing on the same
 host can be routed in the kernel. The traffic does not need to leave the host to be routed.
 This method helps avoid hairpinning.
- NSX Edge transport nodes that host gateways provide network services that cannot be distributed to hosts.
- Gateways exist where east-west routing, north-south routing, and centralized services (such as NAT or VPN) are required.

- A dynamic routing protocol is not needed between the two-tiered gateways, simplifying data center routing.
- Logical routing makes it easy to extend logical networks to physical environments.

5-7 Prerequisites for Logical Routing

For logical routing to work, certain requirements must be met:


- The NSX management cluster must be formed and available.
- Transport zones and N-VDS/VDS should be created.
- Hypervisors must be prepared as NSX transport nodes and added to the management plane.
- Transport nodes must be attached to the appropriate transport zones.
- An N-VDS/VDS Instance must be created on each transport node.
- The NSX Edge nodes must be deployed and preconfigured according to the requirements.

N-VDS virtual switch is not supported on ESXI transport nodes from NSX 4,0.0.1. N-VDS is still supported on NSX Edge transport nodes.

5-8 Logical Routing in NSX

The NSX gateways provide the following features:

- North-south and east-west routing
- Static and dynamic routing
- Multitenancy
- High availability
- IPv6 and multicast
- Centralized services such as gateway firewall, Network Address Translation (NAT), or VPN



Logical Routing in NSX

An NSX gateway reproduces routing functionality in a virtual environment:

- Logical routing is distributed and decoupled from the underlying hardware. Basic forwarding decisions are made locally on the prepared transport nodes.
- The services running on the NSX Edge nodes provide layer 3 functionalities, such as Network Address Translation (NAT).
- When multiple gateway instances are installed, multitenancy and network separation are supported on a single gateway. Logical routing is enhanced for most cloud use cases that involve multiple service providers and tenants.

The NSX gateways provide north-south and east-west connectivity:

- North-south routing enables tenants to access public networks. Traffic leaves or enters a
 tenant administrative domain. Connections to and from the entities outside the tenant's
 premises are considered north-south traffic.
- East-west traffic flows between various networks in the same tenant. Traffic is sent between logical networks (segments) under the same administrative domain.

5-9 Tier-0 and Tier-1 Gateways

NSX provides Tier-0 and Tier-1 gateways. Each gateway has distinct characteristics.

Tier-O Gateway	Tier-1 Gateway
Owned and configured by the provider or infrastructure administrator	Owned and configured by tenants
Supports static or dynamic routing	Does not use dynamic routing protocols
Supports equal-cost multipath (ECMP) routing to upstream physical gateways	Does not support ECMP routing
Forwards the traffic between logical and physical networks (north-south)	Enables routing between segments (east-west) and must be connected to a Tier-O gateway to provide external connectivity
Requires an NSX Edge cluster	Only requires an NSX Edge cluster if centralized services are configured

Gateways are distributed across the kernel of each host. A gateway can be deployed as either a Tier-0 or a Tier-1 gateway:

- Tier-0 gateways provide north-south connectivity.
- Tier-1 gateways provide east-west connectivity.

The Tier-1 gateway must connect to the Tier-0 gateway to access external networks. The Tier-0 gateway is directly connected to upstream physical gateways.

The Tier-1 gateway does not require an NSX Edge node if no services are used. It has connections (preprogrammed by the management plane) toward its upstream Tier-0 gateway


Both Tier-0 and Tier-1 gateways support stateful services, such as NAT. Stateful services are centralized on NSX Edge nodes.

Dynamic routing processes are centralized services while packet forwarding is stateless.

5-10 Single-Tier Topology

In a single-tier topology:

- Only Tier-O gateways are included.
- Segments are connected directly to the Tier-0 gateway.




In a single-tier deployment, only Tier-O gateways are used. Tier-1 gateways are not used. The segments are directly connected to the Tier-O layer. The service provider provides upstream connectivity. The tenant performs southbound connectivity.

5-11 Multitier Topology

In a multitier topology:

- Tier-0 and Tier-1 gateways are included.
- Tier-1 gateways are connected to the Tier-0 gateways.
- Segments are connected to the Tier-1 gateways.



In the diagram, segments A, B, C, and D are connected to the Tier-1 gateways.

The two-tier routing topology is not mandatory. If the provider and the tenant do not need to be separated, a single-tier topology can be used.

In most use cases, the provider owns and configures the Tier-O gateway. The tenants own and configure the Tier-1 gateway. Cloud management platforms typically provision Tier-1 gateways.


5-12 Edge Nodes and Edge Clusters

NSX Edge nodes have the following functions:

- Provide connectivity to external networks.
- Regulred to host the Tier-0 gateways.
- Run gateways with centralized and stateful services such as NAT or VPN.

An NSX Edge cluster is a group of edge nodes:

- NSX Edge clusters provide redundancy and scalability.
- NSX Edge nodes must join an edge cluster.



Edge Cluster


NSX Edge nodes provide computational resources to deliver dynamic routing and services for NSX gateways.

An NSX Edge node can only host one Tier-0 gateway.

An NSX Edge cluster can have up to 10 edge nodes.

5-13 Tier-O Gateway Uplink Connections

Each Tier-0 gateway can have one or more uplinks to physical networks per NSX Edge node.



The diagram shows two different configurations for the uplinks of the NSX Edge nodes:

- On the left, the Tier-0 gateway has one uplink per NSX Edge node mapped to one VLAN to connect to external networks.
- On the right, the Tier-O gateway has two uplinks per NSX Edge node mapped to different VLANs.

In both scenarios, the NSX Edge cluster contains two NSX Edge nodes.


The Tier-O deployment can be active-active or active-standby. When using a dynamic routing protocol, ECMP can be enabled for multiple northbound uplinks.

5-14 Gateway Components: Distributed Router and Service Router (1)

A distributed router (DR) has the following features:

- Provides basic packet-forwarding functionalities
- Spans all transport nodes (host and edge transport nodes)
- Runs as a kernel module in the ESXI hypervisor
- Provides distributed routing functionality
- Provides first-hop routing for workloads

Gateway



A service router (SR) has the following features:


- Provides north-south routing
- Provides centralized services, such as NAT and VPN
- Regulred for uplinks to external networks
- Deployed in edge transport nodes

5-15 Gateway Components: Distributed Router and Service Router (2)



SRs implement centralized services and are required for uplinks.

DRs implement distributed forwarding.



A gateway can be either a Tier-0 or a Tier-1 gateway, depending on the design requirements:

- A Tier-O gateway provides north-south connectivity:
 - In a single-tier topology, the Tier-0 gateway also provides east-west connectivity.
- A Tier-1 gateway provides east-west connectivity.


A Tier-1 and a Tier-0 gateway can have DR and SR components:

- The DR component is distributed among all hypervisors and provides basic packet forwarding:
 - A DR is always created when creating a gateway.
- The SR component is only located in the NSX Edge nodes and provides services:
 - An SR is automatically created on the edge node when you configure the gateway with an edge cluster.

5-16 Realization of Distributed Routers and Service Routers

Distributed routers and service routers are realized in the following manner:

- Distributed router instances can be realized on host and edge transport nodes.
- Service router instances are realized only on edge transport nodes.
- Distributed routers and service routers are automatically interconnected.




The diagram shows that SR and DR instances can be distributed across the Compute and Management clusters.

SR instances are only realized on the edge transport nodes that are running in the Management cluster while DR instances span across both clusters as they run in host transport nodes and in edge transport nodes. The service and distributed routers are interconnected through the TransitLink ports that are automatically created at the time of deployment.

5-17 Gateway Components in a Single-Tier Topology

The diagram shows a logical and physical view of a single-tier configuration.




The diagram represents a single-tier topology where the Tier-O gateway (TO-GW) has two uplinks configured to physical networks and each uplink is connected to a different SR to provide redundancy.

In the physical view, the DR component of the Tier-O gateway is distributed across all transport nodes (ESXI-1, ESXI-2, Edge Node-1, and Edge Node-2), whereas the SR components are located only on the NSX Edge nodes.

In active-active mode, up to eight SRs can be used to connect the uplinks.

5-18 Gateway Components in a Multitier Topology (1)

The diagram shows a logical and physical view of a multitier configuration in which the Tier-1 gateways are not configured with an NSX Edge cluster.




The diagram represents a multitier topology in which the Tier-0 gateway (T0-GW) has two uplinks configured to physical networks. The two Tier-1 gateways have no configured services and therefore were not configured with an NSX Edge cluster. The Tier-1 gateways have no SR components.

In the physical view, the DR component of the Tier-O gateway is visible. The DR component of Tier-1 Gateway A and the DR component of Tier-1 Gateway B are distributed across all transport nodes (ESXI-1, ESXI-2, Edge Node-1, and Edge Node-2).

The Tier-O gateway is configured as active-active. SR1 is in Edge Node-1, and SR2 is in Edge Node-2.

5-19 Gateway Components in a Multitier Topology (2)

The diagram shows the logical and physical views of a multitler configuration with services configured on both Tier-1 gateways.



Logical View

Physical View

The diagram represents a multitier topology in which the Tier-O gateway (TO-GW) has two uplinks configured to physical networks.

Some services are configured on the two Tier-1 gateways. Each gateway has an SR component.

The DR component of Tier-1 Gateway A and the DR component of Tier-1 Gateway B are distributed across all transport nodes (ESXI-1, ESXI-2, Edge Node-1, and Edge Node-2).

The DR component of the Tier-0 gateway is distributed across the NSX Edge transport nodes (Edge Node-1 and Edge Node-2). It is not distributed across the host transport nodes (ESXI-1 and ESXI-2) when the Tier-1 gateway SR components are deployed.

The Tier-O gateway has each uplink connected to a different SR. SR1 is in Edge Node-1, and SR2 is in Edge Node-2.

Tier-1 gateways are automatically configured in Active Standby mode so that the SRs are deployed on the two NSX Edge nodes. You can select the preferred active node in each gateway:


- Edge Node-1 is the active node for Tier-1 Gateway A
- Edge Node-2 is the active node for Tier-1 Gateway B.

A dedicated NSX Edge cluster can be used to deploy Tier-1 gateways with services to achieve better performance in larger environments.

5-20 Gateway Interfaces

The following types of interfaces are used by gateways:

- Uplink Interfaces connect the Tier-O gateways to upstream physical devices.
- Downlink interfaces connect segments (logical switches) to gateways.
- RouterLink ports connect the Tier-0 and Tier-1 gateways.
- An intratier TransitLink is an internal link between the distributed and service routers on a gateway.
- The service interface is a special interface for VLAN-based services and partner service redirection.



In a logical router deployment in NSX, different types of connections require different types of interfaces:

- The uplink interface provides connections to the external physical infrastructure. VLAN and
 overlay interface types are supported depending on the use case. The uplink interface is
 where the external BGP peerings and OSPF adjacencies can be established. External
 service connections, such as IPSec VPN, can also be used through the uplink interface.
- The downlink interface connects workload networks (where endpoint VMs are running) to the routing infrastructure. A downlink interface is configured to connect to a logical switch (corresponding to the segment defined at the policy). This interface provides the default gateway for the VMs in that subnet.
- RouterLink is a type of interface that connects Tier-0 and Tier-1 gateways. The interface is created automatically when Tier-0 and Tier-1 gateways are connected through an internal logical switch also created automatically. It uses a subnet assigned from the 100.64.0.0/16 IPv4 address space by default.
- The Intratier TransitLink connection is also created when a service router is created. It is an
 internal logical switch between the distributed and service routers on a gateway. By default,
 the Intratier TransitLink has an IP address in the 169.254.0.0/24 subnet range.
- The service interface is a special-purpose port to enable services for VLAN-based networks. North-south service insertion is another use case that requires a service interface to connect a partner appliance and redirect north-south traffic for partner services. Service interfaces are supported on both active-standby Tier-O logical routers and Tier-1 routers. Firewall, NAT, and VPNs are supported on this interface. The service interface is also a downlink.

5-21 Review of Learner Objectives

- Explain the function and features of logical routing
- Describe the architecture of NSX two-tier routing
- Differentiate between north-south and east-west routing
- Describe the gateway components
- Recognize the various types of gateway interfaces

5-22 Lesson 2: NSX Edge and Edge Clusters


5-23 Learner Objectives

- Explain the main functions and features of the NSX Edge node
- Describe the functions of the NSX Edge cluster
- Identify the NSX Edge node form factors and sizing options
- Describe the different NSX Edge node deployment methods

5-24 About the NSX Edge Node

The NSX Edge node has several functions:

- Provides connectivity to external networks
- Hosts the Service Router components of Tier-0 and Tier-1 gateways
- Runs the dynamic routing processes and services such as DHCP, NAT, or VPN
- Establishes Geneve tunnels for the overlay networks
- Enables Service Insertion with third-party vendors



NSX Edge is a component of NSX transport zones.

NSX Edge nodes support Data Plane Development Kit (DPDK) for faster packet forwarding in high-performance environments.


NSX Edge nodes use:

- Container-based architecture for most services
- Separate routing tables for management and overlay traffic

5-25 About the NSX Edge Cluster

An NSX Edge cluster is formed by a group of edge nodes and has the following characteristics:

- Provides extra resources to scale out.
- Provides high availability.
- Supports up to 10 edge nodes, and a maximum of 160 clusters can be configured.
- Failure domains can be defined within an edge cluster.



Edge nodes must join an edge cluster to be used as a transport node

An edge cluster can be formed with edge nodes of different form factor types.


Failure domains can be configured with NSX APIs and are used to automatically place the Tier-1 gateway active and standby instances. Failure domains guarantee service availability, for example, if a rack failure occurs. Active and standby Tier-1 gateway services always run in different failure domains.


The NSX edge cluster scaling and maximums are available at https://configmax.vmware.com.

5-26 NSX Edge Node Form Factors

The NSX Edge node supports the following form factors:

- VM on an ESXI host
- Bare-metal node





5-27 NSX Edge VM Sizing Options

For NSX Edge nodes deployed as VMs on hypervisors, several deployment sizes are available.

Size	Memory	VCPU	Disk Space	VM Hardware Version		
Small	4 GB	2	200 GB	ESXI 6.0 or later (VM version 11 or later)		
Medium	8 GB	4	200 GB	ESXI 6.0 or later (VM version 11 or later)		
Large	32 GB	8	200 GB ESXI 6.0 or later (VM version 11 or			
Extra Large	64 GB	16	200 GB	ESXI 6.0 or later (VM version 11 or later)		
	-					

An NSX Edge node that is deployed as a VM runs on ESXI host hypervisors.

For an NSX Edge node VM deployment, the following sizes are available:

- The small appliance is for proof-of-concept deployments.
- The medium size is suitable when only L2 through L4 features, such as NAT, routing, and L4 firewall, are required and the total throughput requirement is less than 2 Gbps.
- The large size is suitable when only L2 through L4 features, such as NAT, routing, and L4 firewall, are required and the total throughput is 2 through 10 Gbps.
- The extra large size is suitable when the total throughput required is multiple Gbps for VPN and north-south Malware Detection.

For additional information, see NSX 4.0 Installation Guide at https://docs.vmware.com/en/VMware-NSX/4.0/installation/GUID-3E0C4CEC-D593-4395-84C4-150CD6285963 html.

5-28 Prerequisites for Deploying the NSX Edge Node VM

For deploying an NSX Edge node in the VM form factor, the following prerequisites must be satisfied:

- The supported deployment media are OVA, OVF, ISO, and preboot execution environment (PXE).
- You can only deploy the NSX Edge node VM on an ESXI hypervisor.
- If using PXE, the password for root and admin users must be encrypted with SHA-512.
- The host name must not contain invalid characters.
- You cannot remove or replace VMware Tools on the NSX Edge node VM.
- The required ports and protocols must be open.
- All the edge nodes in an edge cluster should use the same NTP service.

For DPDK support, the underlaying platform must meet the following requirements:

- CPU must have AESNI capability.
- CPU must have 1 GB Huge Page support.


For more information, see NSX 4.0 Installation Guide at https://docs.vmware.com/en/VMware-NSX/4.0/Installation/GUID-3E0C4CEC-D593-4395-84C4-150CD6285963.html.

For information about the required ports and protocols, see VMware Ports and Protocols at https://ports.esp.vmware.com/home/NSX.

5-29 Deployment Considerations for NSX Edge Node VM Interfaces

An edge node deployment requires various interface types and assignments:

- In the vSphere virtual switch, you must allocate at least two ports for the NSX Edge node.
- The first interface must be defined for management access (eth0) by using one NSX Edge VM vNiC.
- The other four interfaces are datapath interfaces (fp-ethX) and are dedicated for overlay tunneling and uplink connections by using the remaining vNICs.



Edge node VM deployment requires specific interface assignments:


- The first interface of the deployment must be assigned to a management interface.
- Other interfaces must be assigned to the datapath process that creates the overlay or VLAN-based N-VDS.

From NSX-T Data Center 3.2.1, you can use up to four datapath interfaces (fp-ethx) for greenfield deployments. For brownfield deployments upgraded to NSX-T Data Center 3.2.1, you can redeploy the NSX Edge nodes if you need four interfaces for datapath as stated in the VMware NSX-T Data Center 3.2.1 Release Notes at https://docs.vmware.com/en/VMware-NSX/3.2.1/rn/vmware-nsxt-data-center-321-release-notes/index.html.

5-30 Deploying the NSX Edge Node VM with Multiple N-VDS

An edge node VM with multiple N-VDS has the following characteristics:

- Five vNICs are available on the VM.
- The first internal interface is dedicated for management (eth0).
- The remaining interfaces are allocated for the datapath module (fp-ethX).
- Connectivity to VSS and VDS is supported in the hypervisors.
- Multiple N-VDS exist on the edge for overlay and VLAN uplink traffic.




The fp-ethx Interfaces can be assigned to overlay or external VLAN traffic.

Each N-VDS in the edge node can have its own teaming policy.

5-31 Deploying the NSX Edge Node VM with a Single N-VDS

An edge node VM with single N-VDS has the following characteristics:

- Single N-VDS in the VM edge node.
- It carries both overlay and VLAN uplink traffic.
- Two TEPs are configured to provide load balancing for the overlay traffic.



Use the predefined uplink profile, nsx-edge-multiple-vteps-uplink-profile, to configure the edge node VM with two TEPs. TEPs are mapped to each of the uplinks configured in the uplink profile for the overlay traffic.


The ESXI TEP and the edge node TEP IP addresses are in the same subnet.

A named teaming policy can be used for each VLAN uplink traffic for better load balancing of the traffic across the uplinks and override the default teaming policy.

This architecture aligns with the existing support for a single N-VDS in bare-metal edge nodes.

5-32 NSX Edge Node VM Network Offloading with SmartNICs

Like any other VM, you can achieve network offloading of NSX Edge traffic using the new SmartNIC framework. To use this feature, enable Uniform Passthrough (UPT) mode in NSX Edge.



NSX 4.0.1 Edge, using VMXNET3 UPTv2 vNICs, can work in passthrough mode while preserving the functionalities of emulated vNICs such as vSphere vMotion.

5-33 Requirements for the NSX Edge Bare-Metal Node

The NSX Edge node can be installed on bare-metal hardware.

Form Factor	Memory	CPU Cores	Disk	DPDK CPU Requirements
Bare metal (minimum	32 GB	8	200 GB	AES-NI
requirements)				1 GB Huge Page support
Bare metal (recommended	256 GB	24	200 GB	AES-NI
requirements)				1 GB Huge Page support

The NSX Edge bare metal supports only specific CPU types and has some NIC requirements.

For a list of requirements, see NSX 4.0 Installation Guide at https://docs.vmware.com/en/VMware-NSX/4.0/installation/GUID-3E0C4CEC-D593-4395-84C4-150CD6285963.html.

If the hardware is not listed, the storage, video adapter, or motherboard components might not work on the NSX Edge node.

5-34 Prerequisites for Deploying the NSX Edge Bare-Metal Node

For deploying an NSX Edge node in the bare-metal form factor, the following prerequisites must be satisfied:

- The only supported deployment media is ISO either with or without the preboot execution environment (PXE).
- The bare-metal form factor has specific hardware regulrements.
- If using PXE, the password for root and admin users must be encrypted with SHA-512.
- The host name must not contain invalid characters.
- The required ports and protocols must be open.
- All the edge nodes in an edge cluster should use the same NTP service.

For more information, see NSX 4.0 Installation Guide at https://docs.vmware.com/en/VMware-NSX/4.0/installation/GUID-3E0C4CEC-D593-4395-84C4-150CD6285963.html.

For information about the required ports and protocols, see VMware Ports and Protocols at https://ports.esp.vmware.com/home/NSX.

5-35 Deployment Methods for NSX Edge Nodes

The following ways are available to deploy an edge node in the VM form factor:


- Use the NSX UI.
- Deploy an OVF template in vCenter.
- Use the OVF tool command-line utility.
- Use an ISO file and a PXE server to automate the network configuration.

For the bare-metal form factor, an ISO file is used for the installation:

You can use a PXE server to automate the network configurations.

5-36 Deploying NSX Edge Nodes from the NSX UI (1)

You can deploy edge transport nodes directly from the NSX UI by navigating to **System >** Fabric > Nodes > Edge Transport Nodes.



In the NSX UI, select **System > Fabric > Nodes** and click the **Edge Transport Nodes** tab to add an edge VM.

On the Name and Description page of the Add Edge VM wizard, configure the following settings:

- Name
- Host name/FQDN
- Form factor

You can also select the resource reservation (CPU, memory, and shares) during the NSX Edge deployment.

5-37 Deploying NSX Edge Nodes from the NSX UI (2)

The datapath interfaces are defined when adding the edge transport node:

- The number of TEP Interfaces is based on the Uplink Profile selection.
- An NSX Edge node can belong to one overlay network and multiple VLAN transport zones.
- An NSX Edge node must belong to at least one VLAN transport zone to provide uplink access.
- The uplink profile determines the number of uplinks.



All nonmanagement links on the edge node are used for the uplinks and tunnels.

In the example, one uplink is used for the tunnel endpoint (overlay network), and another uplink is used for the external physical network (VLAN).

During the N-VDS creation, the uplinks can be individually assigned per N-VDS. The uplink profiles (single-nic-uplink-profile in the diagram) determine the number of uplink interfaces.

You can modify the datapath interfaces later by editing the edge transport nodes.

5-38 Deploying NSX Edge Nodes from vCenter

You can deploy NSX Edge nodes in the vSphere Client from an OVF template.



The NSX Edge nodes can be installed or deployed using various methods. If you prefer an interactive edge installation, you can use a UI-based VM management tool, such as the vSphere Client connected to vCenter.


The image shows the option to deploy through vCenter or the vSphere Client. A wizard guides you through the steps so that you can provide the required details.

This process does not register the NSX Edge node with the management plane. Additional command-line operations are required.

5-39 Using PXE to Deploy NSX Edge Nodes from an ISO File

By using PXE, the networking settings, such as IP address, gateway, network mask, NTP, and DNS, are automatically configured.

The PXE boot process includes several components, including DHCP, HTTP, and TFTP servers.



This operation automates the installation process. You can preconfigure the deployment with all the required network settings for the appliance.

The password for root and admin users must be encrypted with SHA-512.

The preboot execution environment (PXE) boot can also be used to install NSX Edge nodes on a bare-metal platform.

The PXE method supports only the NSX Edge node deployment. It does not support NSX Manager deployments.

5-40 Installing NSX Edge on Bare Metal

To Install NSX Edge on bare metal using an ISO file:

- Verify that the system BIOS mode is set to Legacy BIOS.
- Create a bootable disk with the NSX Edge ISO file on It.
- Boot the physical machine from the disk.
- Select Automated Install.



Manual installation is also available when you install NSX Edge nodes on a bare-metal server.

After the listed requirements are verified, the installation process should start automatically from the installation media.

After the bootup and power-on processes are complete, the system requests an IP address (manual or DHCP).

By default, the following credentials are used:

- Root login password: vmware
- Password: default

Further setup procedures include enabling the interfaces and joining the edge node to the management plane.

You can use PXE to automate the network configuration installation.

5-41 Joining NSX Edge Bare Metal with the Management Plane

Installing the NSX Edge node by any method other than the NSX UI does not automatically join NSX Edge to the management plane.

To Join an NSX Edge node with the management plane:

- Open an SSH session to the NSX Manager appliance and retrieve the SSL thumbprint by entering get certificate api thumbprint at the command prompt.
- Open an SSH session to the edge node and run the join management-plane command.



The manual installation of NSX Edge nodes does not include an automated procedure to ensure that the management plane sees edge nodes as available resources.

You must Join NSX Edge with the management plane so that they can communicate with each other.

Joining NSX Edge nodes to the management plane ensures that the edge nodes are available from the management plane as managed nodes.

First, you must verify that you have administration privileges to access NSX Edge nodes and the NSX UI. Then you can use the CLI to join the NSX Edge nodes to the management plane.

5-42 Verifying the Edge Transport Node Status

In the NSX UI, navigate to **System > Fabric > Nodes > Edge Transport Nodes** to verify the nodes status and configuration state.



In the NSX UI, select **System > Fabric > Nodes** and click the **Edge Transport Nodes** tab to view the status of the edge nodes known by NSX Manager or the management plane.

The Edge Transport Nodes tab lists the following categories:

- Configuration state
- Node status
- N-VDS
- NSX version


Clicking the information icon next to the node status provides additional information about the reasons for a given status.

Clicking the number of N-VDS gives information about the attached transport zones.

If you want to verify the datapath interfaces, click Edit on the given edge transport node.

5-43 Changing the NSX Edge VM Resource Reservations

You can change the VM resource reservation for the NSX Edge VMs deployed by using NSX Manager. Navigate to **Actions** > **Change Edge VM Resource Reservations**.




To change the NSX Edge VM resource reservations:

- Click Change Edge VM Resource Reservations to access the reservations for the selected NSX Edge node.
- Select the reservations to be changed:
 - CPU Reservation Priority
 - Memory Reservation (%)
 - CPU Reservation (MHz)

5-44 Changing Node Settings






To change the NSX Edge node settings:

- Click Change Node Settings to access the settings for the selected NSX Edge node.
- Select the settings to be changed:
 - Host name/FQDN
 - Search Domain Names
 - Allow SSH Login
 - DNS Servers
 - NTP Servers
 - UPT Mode

5-45 Enabling UPT Mode on NSX Edge Node VMs

You enable the Uniform Passthrough (UPT) mode on an NSX Edge VM during Edge node setting configuration.



You enable the Uniform Passthrough (UPT) mode on an NSX Edge node to improve overall system performance and reduce latency.

Changing the NSX Edge UPT mode on a production NSX Edge node affects services.

The NSX Edge UPT mode setting will only be applicable when:

- At least one of the host's datapath interfaces is SmartNiC-enabled.
- The vSphere supported hardware is version 20 or later.


5-46 Postdeployment Verification Checklist

After deployment, you can verify the connectivity of the NSX Edge nodes in several ways:

- √ If you enabled SSH, verify that you can use SSH to access the newly deployed edge nodes.
- ✓ Verify that you can ping your NSX Edge node.
- ✓ Verify that the NSX Edge nodes can ping their corresponding default gateway.
- Verify that the NSX Edge nodes can ping the hypervisor hosts that are in the same network as the NSX Edge nodes.
- Verify that the NSX Edge nodes can reach their configured DNS server and NTP server.

5-47 Creating an NSX Edge Cluster

You can deploy an NSX Edge cluster from the NSX UI by navigating to System > Fabric > Nodes > Edge Clusters.



You might want to create an NSX Edge cluster for the following reasons:

- Having a multinode cluster of NSX Edge nodes ensures that at least one NSX Edge node is always available.
- An NSX Edge cluster is required to configure Tier-0 gateways uplinks and enable stateful services such as NAT, load balancer, and so on.

To configure an NSX Edge cluster:

- Click +ADD EDGE CLUSTER to start the process for creating an NSX Edge cluster.
- Select a predefined NSX Edge cluster profile or select nsx-default-edge-high-availabilityprofile, which is the default profile.
- Include NSX Edge node members in the NSX Edge cluster.

An NSX Edge transport node can be added to only one NSX Edge cluster.

After creating the NSX Edge cluster, you can later edit it to add NSX Edge nodes.

5-48 Lab 6: Deploying and Configuring NSX Edge Nodes

Deploy NSX Edge nodes and configure them as transport nodes:

- Prepare for the Lab
- Deploy Two NSX Edge Nodes
- Configure an Edge Cluster

5-49 Review of Learner Objectives

- Explain the main functions and features of the NSX Edge node
- Describe the functions of the NSX Edge cluster
- Identify the NSX Edge node form factors and sizing options
- Describe the different NSX Edge node deployment methods

5-50 Lesson 3: Configuring Tier-0 and Tier-1 Gateways


5-51 Learner Objectives

- Describe how to configure a Tier-1 gateway
- Explain how to configure a Tier-O gateway
- Test end-to-end connectivity provided by Tier-0 and Tier-1 gateways

5-52 Gateway Configuration Tasks

To achieve full network connectivity, you must configure the following components:

- Create the Tier-1 gateway and its segments.
- Connect the segments to the Tier-1 gateway.
- Create uplink segments.
- 4. Create the Tier-O gateway and define the uplink connections.
- Configure static or dynamic routing on the Tier-0 gateway.
- Configure the connectivity between the Tier-0 and Tier-1 gateways.
- Enable route advertisement and redistribution.



Depending on the environment, the order of the configuration tasks can vary. Sometimes you might want to create the Tier-0 gateway before the Tier-1 gateway.

Before configuring the Tier-1 and Tier-0 gateways, verify the following settings:


- Your NSX management cluster is stable.
- At least one NSX Edge node is installed.
- An NSX Edge cluster is configured.

The gateways are not automatically connected to each other during the creation process. The management plane cannot determine which Tier-1 instance should connect to which Tier-0 instance. You must manually connect the gateways after their creation.

After you manually connect these instances, the management plane programs the routes in these instances to establish connectivity between tiers.

5-53 Creating the Tier-1 Gateway

Create a Tier-1 gateway by navigating to Networking > Connectivity > Tier-1 Gateways.



Select the **Distributed Only** HA mode if you are not planning to use any stateful services on the Tier-1 gateway.

In the Distributed Only HA mode, no edge cluster is required for the Tier-1 gateway. No SR is created for the Tier-1 gateway (only the DR component is created). This method saves resources and protects from unintended hairpinning of traffic over the edge nodes.


In the Active Standby HA mode, traffic is processed by the active Tier-1 gateway member. The standby member only takes over the traffic if the active member fails.

In the Active Active HA mode, if the Tier-1 gateway is connected to a Tier-0 gateway with stateful services, the Tier-1 gateway also supports stateful services.

For Active Standby and Active Active HA modes, an edge cluster is required to provide the resources for stateful services.


5-54 Connecting Segments to the Tier-1 Gateway

Connect segments to the Tier-1 gateway by navigating to **Networking > Connectivity > Segments > NSX**.



5-55 Using Network Topology to Validate the Tier-1 Gateway Configuration

The topology diagram shows the segments connected to T1-GW-01 and its subnets.



Pointing to an entity highlights the path to the root, for example, VM or segment or Tier-1 gateway.

If the magnification level is less than 1, then the entity names are not displayed in the topology diagram.

Clicking an entity opens the side panel with more details per entity type. An entity might be a VM, a segment, or gateways,

5-56 Testing East-West Connectivity


VMs on various subnets (segments) attached to the Tier-1 gateway can reach each other.

```
Tier-1
                                                                                                                                                                        Gateway
                                                                                                                                                                        72.16.30 1
2.16.20.1
64 hytes from 172.16.20.11: |cop_meq=1 ttl=63 time=3.76 mm
64 bytes from 172.16.20.11: [cmp_seq=2 tt1=63 time=1.78 mm
64 bytes from 172.16.20.11: [cmp_seq=3 tt1=63 time=1.65 mm
64 bytes from 172.16.20.11: [cmp_seq=4 tt1=63 time=1.71 mm
64 bytes from 172.16.20.11: [cmp_seq=5 tt1=63 time=1.74 mm
64 bytes from 172.16.20.11: [cmp_seq=5 tt1=63 time=1.74 mm
                                                                                                                               WALE-Singerneett
                                                                                                                                                             App Segment
                                                                                                                                                                                         DB-5egment
                                                                                                                           VIII
                                                                                                                                                         VIII
                                                                                                                                                                                      vm
- 172.16.20.11 plng statistics -- 5 packets transmitted, 5 received, 0% packet loss, time 4006ms rtt min/aug/max/ndev = 1.575/2.088/3.762/0.838 ms
                                                                                                                                  Segments (Tenant Networks)
                    -01 I - Ja
```

The Tier-1 gateway is created and the interfaces for various logical networks are configured. You can verify the east-west connectivity in the tenant environment.

5-57 Creating the Uplink Segments






Each Tier-0 gateway can have multiple uplink connections, depending on the requirements and the actual configuration.

In the example, two different segments are configured to connect the Tier-O gateway uplink interfaces.

5-58 Creating the Tier-O Gateway (1)






Enable stateful services for Active Active HA mode if you need the Tier-O gateway to support services that need connection state tracking. Turn off stateful services if you want to run only stateless services.

5-59 Creating the Tier-O Gateway (2)


Configure the Tier-O gateway interfaces to associate with the previously created uplink segments.



After Tier-O is created, you can set up the interfaces.

5-60 Configuring Routing

Configure static or dynamic routing to the remote networks by editing the Tier-0 gateway.




You can configure a static or dynamic route to remote networks. BGP is the dynamic routing protocol enabled by default.

BGP and OSPF routing protocols can be active at the same time but for simplicity, you must disable BGP before enabling the OSPF dynamic routing protocol.

5-61 Connecting the Tier-1 and Tier-0 Gateways

Connect the Tier-1 gateway to the Tier-0 gateway by navigating to **Networking** > **Connectivity** > **Tier-1 Gateways** and editing the Tier-1 gateway.



Edit the Tier-1 gateway to connect it to the desired Tier-0 gateway to provide north-south routing and access to external networks.

5-62 Enabling Route Advertisement in the Tier-1 Gateway

Enable route advertisement on the Tier-1 gateway for tenant networks to be propagated to the Tier-0 gateway.



Using route advertisement ensures that the networks defined for tenant segments are available for the connected Tier-O gateway, which can advertise them with the preferred dynamic routing protocol.

5-63 Configuring Route Redistribution on the Tier-O Gateway

Configure route redistribution on the Tier-O gateway to redistribute learned routes to the upstream routers.



Navigate to **Networking** > **Connectivity** > **Tier-0 Gateway** and edit the Tier-0 gateway to configure route redistribution.

BGP and OSPF are the supported destination protocols.


If only static routes are used in your environment, you do not have to configure route redistribution in Tier-0 gateways.

Redistribution into OSPF happens with the following considerations:

- Routes are redistributed as OSPF E2 route type (N2 in NSSA areas).
- Redistribution of OSPF E1 routes (N1 in NSSA areas) is not supported.
- The OSPF cost of the redistributed routes is always 20.

5-64 Using Network Topology to Validate the Tier-O Gateway Configuration

Network Topology shows the Tier-0 gateway connected to the Tier-1 gateway and their network subnets.




The topology diagram shows the IP addresses, such as uplink IPs, router link IPs, and interface IPs, between the NSX objects based on the magnification level.

If the magnification level is less than 1, then the entity names do not appear in the topology diagram.

5-65 Testing North-South Connectivity

VMs on the tenant networks can communicate with external workloads.



In the diagram and the command output, the sa-web-01 (172.16.10.11) VM can ping the Tier-0 gateway (192.168.100.2) and the upstream physical router (192.168.100.1), assuming that routing is configured on the physical router.

sa-web-01 can also ping a remote VM 172.20.10.80.

Complete north-south connectivity is now established.

5-66 Lab 7: Configuring the Tier-1 Gateway

Create and configure a Tier-1 gateway for east-west L3 connectivity:

- Prepare for the Lab
- Create a Tier-1 Gateway
- Connect Segments to the Tier-1 Gateway
- Use Network Topology to Validate the Tier-1 Gateway Configuration
- Test East-West L3 Connectivity

5-67 Review of Learner Objectives

- Describe how to configure a Tier-1 gateway
- Explain how to configure a Tier-0 gateway
- Test end-to-end connectivity provided by Tier-0 and Tier-1 gateways

5-68 Lesson 4: Configuring Static and Dynamic Routing

5-69 Learner Objectives

- Distinguish between static and dynamic routing
- Configure static routes on the Tier-0 gateway
- Configure BGP on the Tier-0 gateway
- Configure OSPF on the Tier-0 gateway

5-70 Static and Dynamic Routing

Static routing:

- Static route configuration is a manual procedure performed by administrators.
- The configuration process enables fine-tuning of route selection.
- Route changes cannot be made dynamically.
- Limited scalability is because of administrative overhead.
- Fallover planning is possible:
 - Network administrators must design and account for all network failure scenarios.
 - Route redundancy must be configured manually.

Dynamic routing:

- Dynamic route configuration enables gateways to exchange information about the network.
- Routing protocols are used to dynamically obtain routes to access the networks.
- Routers Inform neighbor gateways when a network change occurs.

Dynamic routing protocol categories:


- Interior Gateway Protocols (IGPs): These protocols are used for routing in a single routing domain under the administration of a single organization. Some IGP routing protocols are RIP, EIGRP, OSPF, and IS-IS.
- Exterior Gateway Protocols (EGPs): These protocols are used to establish network connectivity between autonomous systems (AS) run by different organizations. BGP protocol is an example of an EGP.

NSX Implements Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF).

5-71 Tier-O Gateway Routing Configurations (1)

The Tier-O gateway supports the following routing configurations:

- Static routing toward upstream physical gateways
- Dynamic routing using Border Gateway Protocol (BGP):
 - External BGP (eBGP) sessions with peers in a different autonomous system (AS)
 - Internal BGP (IBGP) sessions with peers in the same AS



In the diagram, external BGP (eBGP) is used to establish neighbor relationships between Tier-0 and upstream physical gateways with different AS network prefixes that are exchanged between the Border Gateway Protocol (BGP) peers.

The BGP dynamic neighbor enables peering to a group of remote neighbors.


BGP supports 4-byte autonomous system number (ASN).

The Tier-O gateway BGP topology should be configured with redundancy and symmetry between the Tier-O gateways and the external peers.

5-72 Tier-O Gateway Routing Configurations (2)

Tier-0 gateways also support dynamic routing configurations using Open Shortest Path First (OSPF):

- OSPF over point-to-point networks
- OSPF over broadcast networks



Open Shortest Path First (OSPF) is available since NSX-T Data Center 3.1.1.

In the diagram, OSPF establishes adjacencies between Tier-0 and upstream physical gateways. OSPF adjacencies have the following characteristics:

- OSPF is a link state routing protocol, and OSPF establishes and maintains neighbor relationships for exchanging routing updates with other routers.
- Two OSPF routers are neighbors if they are members of the same subnet and share the same area ID, subnet mask, timers, and authentication.
- Setting a password is optional. Authentication methods can be MD5 hashing or clear text.


These adjacencies can be established over broadcast networks or point-to-point networks.

Broadcast and point-to-point networks are defined as follows:

- Broadcast networks support multiple routers connected to the same network. A single broadcast packet can reach all the attached routers. The Ethernet protocol is an example of a broadcast network.
- Point-to-point networks are networks that only join a single pair of routers. They are typically seen on WAN links.


5-73 Configuring Static Routes on a Tier-0 Gateway (1)

You can configure static routes in the ROUTING section of the Tier-O gateway.




5-74 Configuring Static Routes on a Tier-0 Gateway (2)

You can add one or multiple static routes and configure the next hops.



5-75 Configuring Dynamic Routing with BGP on Tier-0 Gateways (1)

To configure dynamic routing, you can configure BGP in the BGP section of the Tier-O gateway.



BGP is enabled by default on Tier-0 gateways. You must set the local AS and configure the BGP neighbors.

You can also configure the following advanced BGP settings:

- Inter-SR routing so Service Routers (SRs) components exchange routing information through iBGP in the same Tier-O gateway.
- Multipath Relax to enable ECMP across different neighboring ASNs if all other BGP attributes are equal.
- Route advertisement filtering using:
 - IP prefix lists to define the networks with subnet masks that are permitted or denied based on a match condition.
 - Community lists to specify the BGP communities allowed. A community is a BGP attribute that can be used to tag a specific set of routes that share common properties.
 - Route maps include a sequence of IP prefix lists or community lists with an associated action to filter or modify the routes advertised. When a match occurs, the gateway performs the action and stops scanning the rest of the route map.

5-76 Configuring Dynamic Routing with BGP on Tier-0 Gateways (2)

You can configure BGP neighbors by adding their AS number, IP addresses, and source addresses.




You can also configure the following advanced BGP settings:

- Bidirectional Forwarding Detection (BFD) is an end-to-end protocol that can detect forwarding path failures.
- Enable Allowas-In to prevent BGP process from dropping the routes received that contain the same AS as the one defined in the Tier-O gateway. Do not enable unless required because the default BGP configuration designed to avoid loops might break.
- Graceful Restart can eliminate or reduce the disruption of traffic associated with routes learned from a BGP neighbor when a control plane fallover occurs. The default mode is Helper Only.

5-77 Verifying the BGP Configuration of the Tier-0 Gateways

You verify the BGP Connectivity Status for each neighbor.




You can also use the get bgp neighbor summary nsxcli command to verify that the BGP neighbor state is established.

5-78 BGP Route Aggregation

Route aggregation is a BGP feature that allows the aggregation of specific routes into one route:


- Reduces the size of the routing tables
- Reduces the number of advertised routes
- Accelerates the best path calculation



The aggregated routes can be received from different AS.


5-79 Configuring Route Aggregation with BGP

Route aggregation can be configured in the BGP section of the Tier-O gateway.



5-80 Configuring Dynamic Routing with OSPF on Tier-0 Gateways (1)

OSPF is not enabled by default in the Tier-O gateway configuration. It must be enabled before setting any parameter.



The OSPF protocol is not enabled by default when creating a Tier-0 gateway.

Turn on the OSPF toggle to enable OSPF.

OSPF Router-ID cannot be manually configured in the UI. It is automatically populated.

Enable OSPF Graceful Restart to keep sending traffic if a control plane fallover occurs in a ToR router.

5-81 Configuring Dynamic Routing with OSPF on Tier-0 Gateways (2)

You configure the area for the Tier-O gateway.



An OSPF network is divided into areas that are the logical groupings of hosts and networks:

- OSPF routers in an area have the same detailed topology for only their own area.
- An area border router (ABR) is the OSPF boundary between two areas.
- In a single-area OSPF, any area can be used. Area 0 is not required.
- The Tier-O gateways do not have ABR functionality currently.
- All inter-area traffic traverses area 0, preventing routing loops.

NSX supports the following OSPF area types:

- Standard Area: Nonbackbone area that must be connected to a backbone area using an Area Border Router (ABR).
- Backbone Area: Must be designed while considering redundancy and cannot be partitioned.
 This area has knowledge of the entire topology. Inter-area traffic must flow through this area.
- Not-so-Stubby Area (NSSA): Blocks external routes from other areas (Inter-area) but can import external routes type-2 from other AS.

Stub Areas, Totally Stubby Areas, and Virtual links to connect areas to the Backbone Area through nonbackbone areas are not supported in NSX.

Follow these criteria when configuring the areas in Tier-O gateways:

- Area Definition supports only one area per Tier-0 gateway.
- Area ID must be either a single number (0) or use dotted format (0.0.0.0).
- Area type can be Normal (Standard or Backbone) or NSSA.
- Authentication is optional and can be configured using MD5 (hashing) or Password (plain text).

5-82 Configuring Dynamic Routing with OSPF on Tier-0 Gateways (3)

Configure the Tier-O gateway Interfaces that will form OSPF adjacencies.



You must follow these criteria when configuring OSPF in the interfaces:

- You can configure a maximum of two uplink interfaces in OSPF per edge node.
- The two interfaces on the NSX Edge node must be in the same area.
- You can configure BFD on the OSPF-enabled interfaces.
- The BFD Helio interval supports a minimum value of 500 milliseconds for an interface configured in an edge VM and 50 milliseconds for an interface configured on bare-metal edges. Dead interval minimum values are 1,500 milliseconds for edge VM and 150 milliseconds for bare-metal edges.

5-83 Verifying OSPF Configuration of the Tier-0 Gateways

You verify the OSPF Neighbors State.




You can also use the get ospf neighbor nsxcli command to verify that the OSPF adjacencies are established.

5-84 OSPF Route Summarization

You must use route summarization in large-scale environments for the following reasons:

- Reduce the Link State Advertisement (LSA) flooding
- Preserve CPU and memory resources
- Ease troubleshooting



Link State Advertisements (LSAs) are the messaging system used in OSPF routing protocol:


- LSAs are exchanged between routers.
- LSAs contain information regarding OSPF links and their status.
- Different types of LSA exist depending on the type of information exchanged.
- LSAs are stored in a local Link State Database (LSDB) in each OSPF router.

The diagram displays:

- Network 10.1.1.0/24 can be advertised as a summary route for both /25 Tier-1 gateway segments.
- The LSA for the summary route is advertised as a type 5 LSA. These LSA types are used to inform about redistributed routes from other routing protocols including static routes. The summarized route is advertised as an OSPF external route of type-2 (N E2).

5-85 Configuring Route Summarization with OSPF

Set Route Summarization in the OSPF section to define the summarized networks.



5-86 Lab 8: Creating and Configuring a Tier-0 Gateway with OSPF

Create a Tier-0 gateway and use OSPF to configure the north-south end-to-end connectivity:

- Prepare for the Lab
- Create an Uplink Segment
- Create a Tier-O Gateway
- Connect the Tier-O and Tier-1 Gateways
- Use Network Topology to Validate the Tier-O Gateway Configuration
- Test the End-to-End Connectivity

5-87 Lab 9: Configuring the Tier-O Gateway with BGP

Create a Tier-0 gateway and use BGP to configure the north-south end-to-end connectivity:

- Prepare for the Lab
- Create an Uplink Segment
- Create a Tier-O Gateway
- Connect the Tier-0 and Tier-1 Gateways
- Use Network Topology to Validate the Tier-O Gateway Configuration
- Test the End-to-End Connectivity

5-88 Review of Learner Objectives

- Distinguish between static and dynamic routing
- Configure static routes on the Tier-O gateway
- Configure BGP on the Tier-0 gateway
- Configure OSPF on the Tier-0 gateway

5-89 Lesson 5: ECMP and High Availability


5-90 Learner Objectives

- Explain the purpose of ECMP routing
- Use the NSX UI to configure ECMP routing
- Identify the active-active and active-standby modes for high availability
- Recognize failure conditions and explain the failover process

5-91 About Equal-Cost Multipath Routing

Equal-cost multipath (ECMP) routing has several features and functions:

- ECMP routing increases the north-south communication bandwidth by combining multiple
 uplinks.
- ECMP routing performs traffic load balancing.
- ECMP routing provides fault tolerance for falled paths.
- A maximum of eight ECMP paths are supported.
- ECMP hashing is based on a 5-tuple algorithm.
- ECMP routing is available for Tier-0 gateway uplinks.



ECMP hashing is based on a 5-tuple algorithm that uses source IP address, destination IP address, source port, destination port, and IP protocol. This method allows a better distribution of the traffic across all the available paths.

5-92 Enabling ECMP in BGP

ECMP is enabled by default on Tier-O gateways when Border Gateway Protocol (BGP) is enabled. ECMP can be disabled in the BGP section on the Tier-O Gateway configuration page.



5-93 Enabling ECMP in OSPF

ECMP is enabled by default on Tier-0 gateways when Open Shortest Path Protocol (OSPF) is enabled. ECMP can be disabled in the OSPF section on the Tier-0 Gateway configuration page.



When configuring ECMP in OSPF, a maximum of two uplink interfaces can be enabled per edge node.

5-94 About High Availability

You can configure high availability on the gateways for redundancy.

High availability can be configured in the following modes:

- Active-active:
 - All the edge nodes are active and run the gateway services simultaneously.
 - The workload is distributed between all nodes to prevent overloading one single node.
- Active-standby:
 - One edge node is active, and one edge node remains on standby.
 - The standby node takes over when the active node becomes unavailable.

Grouping edge nodes offers the benefits of high availability for edge node services. The service router runs on an edge node and has two modes of operation: active-active or active-standby.

The active-active mode is offered on NSX Edge and has the following characteristics:

- Logical routing is active on more than one NSX Edge node at a time.
- Active-active mode is stateless by default. In stateless mode, no tracking tables are kept with the state of connections and services.
- From NSX 4.0.1, when creating a Tier-0 or a Tier-1 the high availability mode can be configured to stateful active-active if you are planning to host stateful services like NAT.


The active-standby mode is also offered on NSX Edge and has the following characteristics:

- Logical routing is active on only one NSX Edge node at a time.
- The standby node takes over if the active member has a failure.
- Active-standby mode can be configured on both Tier-0 and Tier-1 gateways.

5-95 Active-Active HA Mode

The active-active mode is the default high availability mode for Tier-O gateways and provides:

- ECMP to load balance traffic across the different edge nodes.
- Logical routing services are active on more than one edge node at a time.
- NSX version 4.0.1 adds support for stateful services such as NAT in active-active HA mode.



Active-active is a high availability mode where a gateway is hosted on more than one edge node at a time:

- In the active-active mode, traffic is load-balanced across all members.
- For northbound traffic, the DR component sends traffic across the different active SR components.

- When one node falls, traffic is not disrupted but bandwidth is constrained.
- A gateway can span up to eight edge nodes to provide load balancing and redundancy.


Active-active mode is stateless by default. So, services such as NAT and stateful firewall cannot be configured. Only routing and stateless services, such as reflexive NAT, are enabled.

NSX version 4.0.1 supports stateful services such as NAT in stateful active-active mode by pinning specific flows to an individual NSX Edge node.

5-96 Active-Active Topology with BGP

Active-active HA mode topologies with BGP have the following characteristics:

- BGP peering with physical routers is established in all SRs.
- The DRs load balance traffic across all SRs.



In the active-active mode, all the SRs process the northbound and southbound traffic.

Traffic flow considerations:


- In the diagram, DR sends traffic to both active SRs with IPs 169.254.0.2 and 169.254.0.3 in the Transit Segment.
- Routing decision is not influenced on BGP peers.

5-97 Active-Active Topology with OSPF

Active-active HA mode topologies with OSPF have the following characteristics:

- OSPF adjacencies with physical routers are established in all SRs using the same cost.
- The DRs load balance traffic across all SRs.

Physical Topology



For the Tier-O gateway in an active-active configuration with OSPF:


- ECMP can be leveraged with the physical routers but a maximum of two uplink interfaces can be enabled for OSPF per edge node.
- An OSPF cost of 20 is announced by all Tier-0 gateways, influencing the routing decision on physical routers with equal cost paths.

5-98 Active-Standby HA Mode

Active-standby is a high availability mode where a gateway is operational on only a single edge node at a time.

The following centralized stateful services are provided in the active-standby mode:

- Stateful Gateway Firewall
- VPN



Active-standby is a high availability mode where a gateway is operational on only a single edge node at a time.


In the active-standby mode, an elected active member processes all traffic. If the active member fails, a new member is elected to be active:

- Tier-0:
 - The active-standby SRs have different northbound IP addresses and have dynamic routing sessions established on both links.
 - Gateway state is synchronized but does not actively forward traffic. Both SRs maintain dynamic routing peering with the physical gateway.
- Tier-1:
 - The active-standby SRs have the same northbound IP addresses.

5-99 Active-Standby Topology with BGP

Active-standby HA mode topologies with BGP have the following characteristics:

- BGP peering is still established on both SRs.
- The standby SR performs AS path prepending and does not forward traffic to the physical routers.
- The DR sends traffic to the active SR only.




Traffic flow considerations in the active-standby mode with BGP:

- In the diagram, 169.254.0.2/25 is used on both active and standby SRs.
- The standby SR transit segment interface is down for datapath traffic.
- AS path prepending influences BGP peer path selection so the standby SR is less preferred to receive any traffic.
- BGP peering over the standby path ensures optimal BGP route convergence time during fallover.

5-100 Active-Standby Topology with OSPF

The standby Tier-O uses a high OSPF cost to influence route selection on physical routers.

Physical Topology




The Tier-O gateway is configured with two uplinks in an active-standby configuration:

- From an OSPF standpoint, the standby Tier-0 is active.
- The standby Tier-0 uses the OSPF cost to influence the routing decisions on the physical routers.
- The OSPF cost sent by the standby Tier-0 is always 65534, a hard-coded value that cannot be adjusted.
- The route with the lowest value for cost is chosen as the best southbound route.
- The standby SR transit segment interface is down for datapath traffic.
- The DR sends traffic to the active SR only, using the active SR as the northbound route.

5-101 Failover Detection Mechanisms

The fallover process uses the following mechanisms to check the connectivity between tiers:

- Bidirectional Forwarding Detection (BFD): On the management and overlay network
- Dynamic Routing Protocol (BGP or OSPF): On the uplinks




5-102 About BFD

High availability uses BFD to detect forwarding path failures.

BFD provides a low-overhead detection of faults even on physical media that do not support failure detection of any kind, such as Ethernet.

BFD keepalives are sent on both management and tunnel Interfaces.




Management Network

BFD is a network protocol used to detect faults between two forwarding engines connected by a link. Fallures are detected per logical router. The conditions used to declare an edge node as down are the same in active-active and active-standby high availability modes.

To ensure uninterrupted routing of network traffic, the NSX Edge nodes exchange keepalive messages, which are BFD sessions running between the nodes. The edge nodes in an edge cluster exchange BFD keepalive on the management and tunnel interfaces. When the standby Tier-0 gateway falls to receive keepalives on both management and tunnel interfaces, it announces itself as active.

5-103 Failover Scenario with BFD


If a standby gateway fails to receive BFD keepalives on both management and tunnel interfaces, the gateway becomes active.



The BFD protocol provides fast detection of failure for forwarding paths or forwarding engines, improving convergence. Edge VMs support BFD with a minimum BFD timer of 500 milliseconds with three retries, providing 1.5 seconds failure detection time. Bare-metal edges support BFD with a minimum BFD timer of 50 milliseconds with three retries, which implies a 150 milliseconds failure-detection time.

5-104 Failover Scenario with Dynamic Routing

Dynamic routing peering sessions are established on the uplinks with physical routers. If an active gateway loses all its routing neighbors and a standby gateway is available, the active gateway steps down and becomes the standby gateway and the standby gateway is promoted to the new active gateway.



If an active gateway loses all its dynamic routing peerings and a standby gateway is configured, failover occurs. An active SR on an edge node is declared down when all the dynamic routing sessions on the peer SR are down.

This scenario is only applicable on Tier-0 with dynamic routing.

BGP or OSPF is configured on the uplink between each NSX Edge node and the exterior physical gateways.

Status is monitored during the dynamic routing keepalive exchanges.

The default BGP timers are a keepalive interval of 60 seconds and the minimum time between advertisements is 30 seconds. The default OSPF timers are a Helio interval of 10 seconds and a Dead interval of 40 seconds.

If all overlay tunnels to the compute hypervisors are down, the active edge node does not receive tunnel traffic from compute hypervisors. Then the standby edge node takes over.

5-105 Failover Modes

You can select different fallover modes in active-standby HA mode:

- Preemptive: If the preferred node falls and recovers, it takes over its peer and becomes the
 active node. The peer changes its state to standby.
- Non Preemptive: If the preferred node falls and recovers, it checks whether its peer is
 active. If the peer is active, the preferred node stays in standby mode.



Preemptive and non-preemptive modes are used in a fallback scenario after a fallover occurs.

During the fallover, the standby node becomes active.

The fallback happens when the node that falled becomes available again:

- If non-preemptive mode is configured, nothing happens.
- If preemptive is configured, the original active (preferred) node takes over again.

5-106 Review of Learner Objectives

- Explain the purpose of ECMP routing
- Use the NSX UI to configure ECMP routing
- Identify the active-active and active-standby modes for high availability
- Recognize failure conditions and explain the failover process

5-107 Lesson 6: Logical Routing Packet Walk


5-108 Learner Objectives

- Describe the datapath of single-tier routing
- Explain the datapath of multitier routing

5-109 Single-Tier Routing: Egress to Physical Network (1)

A packet is sent from the source VM 10.1.1.10 to the destination VM 192.168.10.1:

The packet is forwarded to its default 10.1.1.1 gateway.




The default gateway 10.10.10.1/24 is on the Tier-0 DR (T0 DR) component of the hypervisor where the VM resides.

The Internal transit subnet 169.254.0.0/24 in the diagram, which connects the Tier-0 DR (T0 DR) in the hypervisor with the Tier-0 SR (T0 SR) in the NSX Edge node, is the default subnet, but it can be configured with a different range.

5-110 Single-Tier Routing: Egress to Physical Network (2)


The gateway (TO DR) checks its forwarding table. Because a specific route does not exist
for the 192.168.10.0/24 network, the packet is sent to the default 169.254.0.2 gateway,
which is the TO SR component on the edge node.



The packet is sent to the default 169.254.0.2 gateway over the Internal Transit Subnet. 169.254.0.2 is an interface of the TO SR component on the edge node.

5-111 Single-Tier Routing: Egress to Physical Network (3)


To send the packet from the hypervisor to the edge node, the packet is encapsulated with a Geneve header.



The source host (TEP 172.16.215.67) encapsulates the packet with a Geneve header to send it to the remote host (TEP 172.16.215.124). The original packet is intact.


5-112 Single-Tier Routing: Egress to Physical Network (4)

4. The encapsulated packet is sent to the edge node across the overlay tunnel.



5-113 Single-Tier Routing: Egress to Physical Network (5)


The edge node decapsulates the packet and sends it to its SR component. The gateway (TO SR) routing table shows a route for the 192.168.10.0/24 network over the uplink segment.



To reach the destination 192.168.10.0/24 network, the next-hop 172.16.215.98 is used.


5-114 Single-Tier Routing: Egress to Physical Network (6)

The edge node sends the packet to its upstream physical gateway, which routes the packet to its destination 192.168.10.1.




5-115 Single-Tier Routing: Ingress from Physical Network (7)

For the return packet, the source VM 192.168.10.1 sends the packet to its default gateway, which routes the packet to the edge node.



5-116 Single-Tier Routing: Ingress from Physical Network (8)

 The SR and the DR components on an edge node share their routing table. A route is directly connected to the 10.1.1.0/24 network over Segment 1. The packet is sent to the remote host by using the TO DR Interface.




In the edge node, the SR and DR components share their routing table. This method removes the extra step of using the Internal Transit Subnet to route from SR to DR.

The Internal Transit Subnet is used when routing from a DR component from a hypervisor to an SR component in an edge node.

Because the routing table is shared, when the TO SR component in the edge receives the packet, it sends the packet to the remote host (Hypervisor with TEP 172.16.215.67) through the TO DR interface.


5-117 Single-Tier Routing: Ingress from Physical Network (9)

To send the packet from the edge node to the hypervisor, the packet is encapsulated with a Geneve header.




5-118 Single-Tier Routing: Ingress from Physical Network (10)

The encapsulated packet is sent across the overlay tunnel.



5-119 Single-Tier Routing: Ingress from Physical Network (11)


The receiving host decapsulates the packet and routes it to its destination (VM 10.1.1.10).



5-120 Multitier Routing: Egress to Physical Network (1)

A packet needs to be sent from the source VM 10.1.1.10 to the destination VM 192.168.10.1:

The packet is forwarded to its default 10.1.1.1 gateway.




In the example, the Tier-1 gateway is a distributed router only and has no services configured, and therefore has no SR components.

The default gateway 10.10.10.1/24 is on the Tier-1 DR component of the hypervisor where the VM resides.

The Tier-0 Tier-1 (T0-T1) transit subnet 100.64.16.0/31 in the diagram, which connects the Tier-1 DRs (T1 DR) with the Tier-0 DRs (T0 DR) in the hypervisor and the NSX Edge node, is the default subnet, but it can be configured with a different range.


5-121 Multitier Routing: Egress to Physical Network (2)

 The gateway (T1 DR) checks its forwarding table to make a routing decision. Because no specific route exists for the 192.168.10.0/24 network, the packet is sent to the default 100.64.16.0 gateway, which is the DR instance of Tier-0 on the same hypervisor.




5-122 Multitier Routing: Egress to Physical Network (3)

The packet is sent to the TO DR instance on the same hypervisor through TO-T1 Transit Subnet.



5-123 Multitier Routing: Egress to Physical Network (4)


 The gateway (TO DR) checks its forwarding table to make a routing decision. The packet is sent to the default 169.254.0.2 gateway, which is the TO SR component on the edge node.



The packet is sent to the default 169.254.0.2 gateway over the Transit segment. 169.254.0.2 is an interface of the Tier-O SR component that attaches to the Internal Transit network.

5-124 Multitier Routing: Egress to Physical Network (5)


To send the packet from the hypervisor to the edge node, the packet is encapsulated with a Geneve header.



The source host (TEP 172.16.215.67) encapsulates the packet with a Geneve header to send it to the edge node (TEP 172.16.215.124). The original packet is intact.


5-125 Multitier Routing: Egress to Physical Network (6)

The encapsulated packet is sent to the edge node across the overlay tunnel.




5-126 Multitier Routing: Egress to Physical Network (7)

7. The edge node decapsulates the packet and sends It to Its TO SR Instance.



5-127 Multitier Routing: Egress to Physical Network (8)


 The gateway (T0 SR) routing table shows a route for the 192.168.10.0/24 network over the uplink segment.



To reach the destination network 192.168.10.0/24, the next-hop 172.16.215.98, is used.


5-128 Multitier Routing: Egress to Physical Network (9)

The edge node sends the packet to its upstream physical gateway, which routes the packet to its destination, 192.168.10.1.




5-129 Multitier Routing: Ingress from Physical Network (10)

 For the return packet, the source VM 192.168.10.1 sends the packet to its default gateway, which routes the packet to the edge node.



5-130 Multitier Routing: Ingress from Physical Network (11)

 The SR and the DR components of the Tier-0 gateway share their routing table because they are both on the edge node. The routing decision is made to send the packet to the Tier-1 DR instance in the same edge node.




In the edge node, the SR, and DR components share their routing table. This method removes the extra step of using the Internal Transit Subnet to route from SR to DR.

The Internal Transit Subnet is used when routing from a DR component in a hypervisor to a SR component in an edge node.

As the routing table is shared, when the Tier-O SR component receives the packet, the packet is sent to the Tier-1 DR component of the edge node through the DR interface.


5-131 Multitier Routing: Ingress from Physical Network (12)

12. The packet is sent to the T1 DR instance on the edge node through T0-T1 Transit Subnet.




5-132 Multitier Routing: Ingress from Physical Network (13)

 The gateway (T1 DR) checks its forwarding table to make a routing decision. A route is directly connected to the 10.1.1.0/24 network over Segment 1. The packet is sent to the remote host.



5-133 Multitier Routing: Ingress from Physical Network (14)


 To send the packet from the edge node to the hypervisor, the packet is encapsulated with a Geneve header.



The source host (TEP 172.16.215.124) encapsulates the packet with a Geneve header to send it to the remote host (TEP 172.16.215.67). The original packet is intact.


5-134 Multitier Routing: Ingress from Physical Network (15)

15. The encapsulated packet is sent to the edge node across the overlay tunnel.



5-135 Multitier Routing: Ingress from Physical Network (16)

The receiving host decapsulates the packet and routes it to its destination (VM 10.1.1.10).



5-136 Review of Learner Objectives

- Describe the datapath of single-tier routing
- Explain the datapath of multitier routing

5-137 Lesson 7: VRF Lite


5-138 Learner Objectives

- Describe VRF Lite
- Explain the benefits of VRF Lite
- Configure and validate VRF Lite deployments

5-139 About VRF Lite

VRF Lite has the following characteristics:

- Multiple routing instances can be configured without deploying additional Tier-0 gateways and NSX Edge nodes.
- Logical routing isolation is provided in NSX and to external peers that are compatible with the VRF Lite technology.
- MPLS/MP-BGP protocols are not used.



Virtual Routing and Forwarding (VRF) allows the coexistence of multiple routing instances in one routing device. Independent routing and forwarding tables are maintained for each instance.

Separation between tenants and applications does not require additional Tier-0 gateways and NSX Edge nodes with VRF Lite.

VRF Lite provides logical routing isolation in NSX and spans it to external peer devices that support this technology.

VRF Lite differs from other VRF implementations because it does not rely on MPLS and MP-BGP protocols running in the physical network.

- Multiprotocol Label Switching (MPLS): This layer 2 protocol is used to forward traffic based on labels. MPLS does not use network addresses like IP protocol. These labels identify the paths between the endpoints in VRFs.
- Multiprotocol Border Gateway Protocol (MP-BGP): This BGP protocol extension is used to propagate the VRF routing information across MPLS network devices.

5-140 VRF Lite Requirements and Limitations

A VRF Lite deployment has the following requirements:

- A deployed Tier-0 gateway
- External connectivity with a layer 3 peer
- Peer device that supports 802.1Q protocol (VLAN tagging)

The following services cannot be configured in a VRF gateway:

- VPN
- OSPF routing

A Tier-O gateway must be used to deploy VRF gateways. It is the default Tier-O gateway and is the parent gateway of the VRF gateways.

The Tier-O gateway, used as the default Tier-O gateway, can be an existing Tier-O gateway with connected Tier-1 gateways.

You can have more than one Tier-0 gateway with VRF gateways.


VLAN tagging is used to separate the VRFs in the uplink segment that connects with the external devices.

These limitations apply only to the VRF gateway. You can connect a Tier-1 gateway configured with a VPN to a VRF gateway and that is fully supported.

5-141 Use Cases for VRF Lite

VRF Lite can be used to enable the following features:

- Allow the same network address to coexist in different routing domains.
- Provide feature compatibility with existing network installations.
- Run multiple routing instances in the same gateway to optimize existing resources.




The only solution for customers that require separate network routing instances for each tenant is to deploy multiple Tier-O gateways. However, these deployments can create scalability issues because only a single Tier-O gateway can be deployed per NSX Edge node, specifically, for deployments based on bare-metal edges.


VRF Lite helps network administrators to deal with the overlapping of network ranges in the same routing domain between business units or after a merger.

It also allows existing VRF Lite deployments in the physical network infrastructure to be extended to NSX.

5-142 VRF Lite Topologies

VRF Lite can be deployed in single-tier and multitier topologies.






The following topology requirements must be considered:

- VRF gateways can only be deployed as Tier-0 gateways.
- A trunk is used to interconnect the different VRFs with the physical gateway.
- The physical gateway and the underlying infrastructure like vSphere Distributed Port Groups have to support trunking.
- Tenants can be connected to Tier-0 segments and Tier-1 segments.

5-143 VRF Lite Gateway Interfaces

The following types of interfaces are used with VRF gateways:

- The Logical Router (LR) trunk port connects the parent Tier-0 gateway to upstream physical devices.
- The VRF Uplink Interface is internally connected to the LR trunk port of the parent Tier-O gateway.
- The Intratier Transit Link is the Internal link between the service router (SR) and distributed router (DR) of a VRF gateway.
- The Downlink Interfaces connect VRF gateways to segments with attached workloads.
- The RouterLink ports connect VRF gateways with Tier-1 gateways.



The LR trunk port is a network interface while the VRF uplink port can be seen as a subinterface with a specific VLAN ID.


The LR trunk port is internally created in the parent Tier-0 gateway and is the only port connected to the uplink trunk segment.

The other interfaces are the same type as the interfaces used in the standard Tier-0 and Tier-1 gateways.

5-144 VRF Lite: Control and Data Planes

VLAN tagging (802.1Q) in the uplink trunk segment provides isolation for each VRF:

- VLAN is the channel for the data plane.
- BGP protocol instance in each VRF provides the control plane functionality.



A dedicated BGP instance runs in every VRF. You do not need to use the extensions in the MP-BGP protocol to exchange the VRF routing information.

BGP is the control plane because it dynamically propagates and updates routing information to all VRF peers.

Each VLAN is mapped to a VRF and only transports traffic for that VRF.

5-145 Configuring VRF Lite

Follow these steps to configure VRF Lite.



The deployment of a Tier-0 gateway is optional if an existing Tier-0 gateway is used instead as the default Tier-0 for the VRF gateway.

VRF gateways inherit the following configuration options from the default Tier-O gateway:


- HA mode
- Edge cluster
- BGP AS number
- Graceful restart settings
- BGP multipath relax

You do not need to connect a Tier-1 gateway to the VRF gateways. Tenants can be directly connected to VRF gateways.

5-146 Deploying the Default Tier-O Gateway

To deploy and configure the default Tier-O gateway as a standard Tier-O gateway:

- Navigate to Networking > Connectivity > Tier-O Gateways in the NSX Ut.
- Select ADD GATEWAY > Tier-0.




You configure the following parameters to deploy the default Tier-0 gateway:

- HA mode
- Edge cluster
- Uplink Interfaces

5-147 Adding Uplink Interfaces to the Default Tier-O Gateway

Connect the default Tier-O uplink interfaces to the uplink segments in the Set Interfaces window.



A VLAN ID is not configured in the uplink interface.

Uplink interfaces are required to deploy the default Tier-0 gateway in the NSX Edge nodes.

5-148 Configuring BGP for the Default Tier-O Gateway

Configure BGP parameters to use dynamic routing with external routers in the BGP configuration section.



VRF gateways inherit the global BGP configuration:

- Local AS
- Graceful restart
- Graceful restart timer
- Graceful restart stale timer
- Multipath relax

These parameters can only be changed in the default Tier-O gateway.

5-149 Adding the Uplink Trunk Segment for the VRF Gateway

To configure the trunk segment for connecting the VRF gateway uplinks:

- Navigate to Networking > Connectivity > Segments > NSX in the NSX UI.
- Select ADD SEGMENT.



A segment is configured as a trunk when more than one VLAN is configured. A range of VLANs can also be specified (VLAN X-Y).

Uplink trunk segments specify which VLANs are allowed but do not add 802.1Q VLAN tagging. Tags are added in the uplink interface of VRF gateways.

You can configure a dedicated uplink trunk segment for each VRF uplink if the trunk is configured as a single VLAN range (X-X). As a best practice, you should connect the VRF uplinks to the same uplink trunk segment. This method reduces the number of resources required (segments, logical switch ports, and logical router ports).

5-150 Deploying the VRF Gateway

To deploy and configure the VRF gateway:

- Navigate to Networking > Connectivity > Tier-O Gateways in the NSX Ut.
- Select ADD GATEWAY > VRF.




Edge cluster and HA mode configuration values are automatically inherited from the default Tier-0 gateway.

You do not need to configure VRF settings for VRF Lite. These settings are used for Ethernet VPN (EVPN).

5-151 Adding Uplink Interfaces to the VRF Gateway


In the Set Interfaces window, connect the VRF gateway uplink interfaces to the uplink trunk segment.




VRF uplink interface configuration options:

- 802.1Q VLAN tagging is added at the uplink level.
- VRF gateway uplinks must be connected to a trunk segment.
- Access VLAN ID is required and must belong to the range specified for the trunk segment.

5-152 Configuring the BGP for the VRF Gateway





The following parameters are inherited from the default Tier-0 gateway and cannot be modified at the VRF level:

- Local AS
- Graceful restart
- Graceful restart timer
- Graceful restart stale timer
- Multipath relax

Inter-SR IBGP is not supported in VRF gateways.


BGP can be enabled or disabled per VRF gateway.

Route aggregation and BGP neighbors are local configurations per VRF.

5-153 Connecting a Tier-1 Gateway to the VRF Gateway

To connect the Tier-1 gateway to the VRF gateway:


- Navigate to Networking > Connectivity > Tier-1 Gateways in the NSX UI.
- Select a Tier-1 Gateway and click Edit from the Actions menu next to >.
- From the Linked Tier-O Gateway drop-down menu, select the VRF gateway.



All Tier-0 gateways and VRF gateways are listed in the **Linked Tier-0 Gateway** drop-down menu.

5-154 VRF Lite Validation

Navigate to **Networking** > **Connectivity** > **Tier-0 Gateways** to obtain the list of VRF gateways with its status and associated errors.



VRF gateways are marked with the VRF tag in the name field.

5-155 Lab 10: Configuring VRF Lite

Configure and verify the VRF Lite functionality to isolate routing domains:

- Prepare for the Lab
- Create the Uplink Trunk Segment
- Deploy and Configure the VRF Gateways
- Deploy and Connect the Tier-1 Gateways to the VRF Gateways
- Create and Connect Segments to the Tier-1 Gateways
- Attach VMs to Segments on Each VRF
- Test the VRF End-to-End Connectivity
- Review the Routing Tables in Each VRF
- Verify the Routing Isolation Between VRFs

5-156 Review of Learner Objectives

- Describe VRF Lite
- Explain the benefits of VRF Lite
- Configure and validate VRF Lite deployments

5-157 Key Points (1)

- The NSX routing function meets the needs of service providers and tenants.
- An administrator manually performs static route configuration.
- Dynamic route configuration enables gateways to exchange information about the network.
- NSX logical routing commonly implements a two-tiered topology.
- Tier-1 gateways have downlink ports to connect to NSX segments and transit ports to connect to Tier-0 gateways.
- A gateway has two components: a distributed router and a service router.
- You can deploy an NSX Edge node through the NSX UI, the OVF tool, and an ISO file in a PXE environment.
- Joining NSX Edge nodes with the management plane ensures that NSX Manager and the NSX Edge nodes can communicate with one another.

5-158 Key Points (2)

- A multinode NSX Edge cluster helps ensure that at least one NSX Edge node is always available.
- NSX Implements Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF) dynamic routing protocols.
- External BGP (eBGP) is used to interchange autonomous system iP addresses with another autonomous system.
- OSPF is a link state routing protocol that maintains adjacencies with neighbor routers over Broadcast and Point-to-Point networks.
- ECMP routing increases the north-south communication bandwidth by adding an uplink to the Tier-0 gateway and configuring it for each NSX Edge node in an NSX Edge cluster.
- Multiple NSX Edge nodes can be pooled in a cluster for scale-out and redundancy.
- High availability supports two modes; active-active and active-standby.
- Stateful active-active high availability mode supports stateful services like NAT.
- VRF Lite enables you to configure multiple routing instances without deploying additional Tier-0 gateways and NSX Edge nodes.

Questions?

Module 6

NSX Logical Bridging

6-2 Importance

Logical bridging enables layer 2 communication between devices on NSX overlay-backed virtual networks and VLAN-backed physical networks. Logical bridging is also useful when you must split a subnet across physical and virtual workloads during a physical-to-virtual migration.

6-3 Lesson 1: NSX Logical Bridging


6-4 Learner Objectives

- Describe the purpose and function of logical bridging
- Distinguish between routing and bridging
- Create bridge profiles
- Create a bridge-backed segment and bridge traffic between virtual and physical environments

6-5 Overview of Logical Bridging

The bridging function provides layer 2 connectivity between the overlay segments and VLANbacked physical networks:

- The layer 2 bridge feature is provided by the NSX Edge node.
- Traffic is bridged in and out of the NSX domain.
- The NSX Edge firewall provides granular control over the bridged traffic.



The bridge feature is available in bare-metal NSX Edge nodes and in NSX Edge VMs.

6-6 Logical Bridging Use Cases

Layer 2 bridging is used in the following common use cases:

- Performing a physical-to-virtual migration
- Connecting non-virtualized compute platforms

Layer 2 bridging is useful when you must split a subnet across physical and virtual workloads during a physical-to-virtual migration.


Layer 2 bridging also connects non-virtualized compute platforms.

6-7 Routing and Bridging for Physical-to-Virtual Communication

You can achieve physical-to-virtual communication through routing or bridging.


Routing:

- Standard layer 3 routing process between physical and virtual workloads
- Equal-cost multipath (ECMP) for scale-out and failure isolation



Bridging:

- The layer 2 flat broadcast domain is used for both physical and virtual workloads, resulting in limited domain size and lack of scalability.
- A single active bridge is needed for an overlay-to-VLAN pair.



Route when you can. Bridge when you have to.


When connecting your physical workloads on traditional physical networks to a virtualized environment, you can use routers running standard routing protocols to route traffic between workloads in the two environments.

If you do not want to use routing and you must place your physical and virtual devices on a single layer 2 subnet, then you can enable bridging.

6-8 Example of Virtual-to-Physical Routing

Routing occurs between VMs in the NSX virtual environment and a server in the physical environment:


- The web tier and application tier belong to the NSX overlay.
- The physical server resides on a different subnet.
- The Tier-O gateway provides north-south routing between the physical server and the application-tier servers.



6-9 Example of Virtual-to-Physical Bridging

Bridging occurs between a VM in the NSX virtual environment and a server in the physical environment on the same subnet:

- The application tier is in the NSX overlay.
- The physical server resides in the physical environment.
- The physical server is on the same subnet as the application tier servers.
- The communication between the physical server and the application tier occurs through the NSX Edge node.



The communication between the App1 VM and the physical server is realized by the gateway that performs bridging.


The diagram demonstrates how the physical server and the App1 server (virtual) can exist on the same subnet.

In the diagram, the traffic between the physical server and the application tier does not pass through the Tier-1 and Tier-0 gateways.

6-10 Logical Bridging Components

Layer 2 bridging consists of the following NSX components:

- A bridge profile is used to specify which edge nodes are involved in bridging.
- A bridge-backed segment is part of an overlay transport zone with an attached bridge profile.



You configure a bridge-backed segment to provide layer 2 connectivity between VMs in an NSX overlay and devices that are outside NSX.


An NSX segment that is attached to a bridge profile provides the following information:

- VLAN to which traffic is bridged
- Physical port to be used (identified by the VLAN transport zone)

For more information about the additional edge configurations, see the section about configuring edge-based bridging in NSX Administration Guide at https://docs.vmware.com/en/VMware-NSX/Index.htmlhttps://docs.vmware.com/en/VMware-NSX/4.0/administration/GUID-FBFD577B-745C-4658-B713-A3016D18CB9A.html.

6-11 Using Multiple Bridge Profiles

You can configure multiple bridge profiles on an NSX Edge node.



Bridge Profiles on Edge Nodes Sample Topology

Bridge profile 1 maps VLAN 1 with Segment 1.

Bridge profile 2 maps VLAN 2 with Segment 2.

Edge 1 is the primary edge for bridge profile 1 and the backup edge for bridge profile 2.

Edge 2 is the primary edge for bridge profile 2 and the backup edge for bridge profile 1.

6-12 Creating an Edge Bridge Profile

To create an edge bridge profile, select **Networking > Connectivity > Segments > Profiles > Edge Bridge Profiles** in the NSX UI.



When you configure an edge bridge profile, you select one of the following fallover modes:

- Preemptive: The bridge on the primary edge node always becomes the active bridge when
 it becomes available again after a failure.
- Non Preemptive: The bridge on the primary edge node remains at standby if it becomes available after a failure when the bridge on the other edge node is already active.

After creating an edge bridge profile, additional configurations are required.


The following configuration options are available:

- Option 1: Configure the Promiscuous mode
- Option 2: Configure MAC address learning
- Option 3: Configure a sink port

6-13 Creating a Layer 2 Bridge-Backed Segment

The bridge-backed segment provides layer 2 connectivity to overlay VMs outside NSX:

- To create a segment, select Networking > Connectivity > Segments > NSX.
- After the segment creation, set up an edge bridge to attach the edge bridge profile.



Before configuring a bridge-backed logical segment, you must verify the following components:


- A bridge profile must be configured.
- At least one ESXI host must exist to serve as a regular transport node, This node hosts VMs that require connectivity with devices outside an NSX deployment.
- A VM or another physical device must exist outside the NSX deployment. This physical device must be attached to a VLAN port matching the VLAN ID of the bridge-backed logical segment.

A regular transport node and a physical device are not mandatory for creating a bridgebacked logical segment. However, if the transport node and physical device are not available, you cannot use the bridge.

The option to add an edge bridge profile is not visible immediately after you create a segment. You must save the configuration and edit it again.

6-14 Monitoring the Bridged Traffic Statistics

To monitor the bridged traffic statistics, you expand the bridge-backed segment and then click **VIEW STATISTICS**.



6-15 Review of Learner Objectives

- Describe the purpose and function of logical bridging
- Distinguish between routing and bridging
- Create bridge profiles
- Create a bridge-backed segment and bridge traffic between virtual and physical environments

6-16 Key Points

- The NSX Edge layer 2 bridge is responsible for bridging traffic between the NSX overlay and VLAN-backed VMs or physical devices outside the NSX deployment.
- With a bridge profile, an NSX Edge cluster can provide layer 2 bridging to a segment.
- The traffic bridged in and out of the NSX domain is subject to the NSX Edge layer 2 bridge firewall.

Questions?

Module 7

NSX Firewalls

7-2 Importance

NSX includes a distributed firewall and gateway firewall to protect both east-west and northsouth traffic. You must understand the architecture and configuration of the NSX firewalls to ensure that your workloads are protected.

7-3 Module Lessons

- NSX Segmentation
- NSX Distributed Firewall
- 3. Use Case for Security in Distributed Firewall on VDS
- 4. NSX Gateway Firewall


7-4 Lesson 1: NSX Segmentation

7-5 Learner Objectives

- Define NSX segmentation
- Recognize use cases for NSX segmentation
- Identify steps to enforce Zero-Trust with NSX segmentation

7-6 Traditional Security Challenges

The traditional security model assumes that all users and components in an organization's network can be trusted



The foundation of IT security has remained almost the same for the last 30 years.

Traditional IT security is built on static information:

- If the device joined the domain.
- Whether the user has the correct password.
- Trusted users are within the firewall. Threats and attacks originate outside.


Virtual Private Network (VPN) and Multifactor Authentication (MFA) were introduced later to ensure external protection. But these types of authentication are not enough.

This approach assumes that a user's identity is not compromised and that all users act responsibly and can be trusted.

This traditional perimeter-centric security approach has proven inadequate for protecting modern IT environments.

7-7 About Zero-Trust Security

Zero-Trust is a security model that does not automatically trust entities in the security perimeter.



Zero-Trust is a security model that does not automatically trust entities in the security perimeter.

This model emerged to mitigate the increase of network attacks and insider threats that exploit the breaches of a traditional perimeter-centric approach to security.

The rapidly changing work styles and increased use of SaaS applications resulted in Zero-Trust security becoming one of the most important forms of alternative security.

Zero-Trust moves the architecture from a single large DMZ to multiple smaller boundaries around each application and data. If an attacker succeeds in penetrating one of these boundaries, the attacker can only move in that perimeter and be easily contained.

7-8 About NSX Segmentation

Segmentation is the process of dividing data center infrastructure into small zones, allowing finegrain control and inspection of traffic flows.

NSX includes a distributed, scale-out internal firewall that simplifies and automates both macrosegmentation and micro-segmentation:

- Macro-segmentation isolates and secures specific environments.
- Micro-segmentation isolates and secures specific applications in an environment.



Macro-segmentation is also called network segmentation. Macro-segmentation isolates and secures specific environments (virtual zones), such as development and production, from each other. Attackers cannot move laterally between these environments.

Micro-segmentation enables security teams to define and enforce granular controls to the workload level of an application.

7-9 Use Cases for NSX Segmentation

NSX segmentation has the following use cases:

- Enforce a Zero-Trust architecture
- Rapidly deploy network segments
- Isolate and secure applications







With NSX, security teams can deploy network segments easily, enable application isolation, and enforce a Zero-Trust architecture with a single solution.

Use cases:

- NSX segmentation enforces a Zero-Trust architecture by creating granular policies between applications, services, and workloads.
- Network segments, virtual security zones, and partner domains are quickly created and configured as they are entirely defined in software. NSX also removes the need to architect the network again and to deploy discrete appliances.
- Critical applications and shared services are protected from being compromised by two
 mechanisms: discovery of application boundaries using NSX Intelligence and setting up
 segmentation policies at the application level. NSX also ensures that policies stay up to date
 as applications evolve or move.


7-10 NSX Segmentation Benefits

NSX Segmentation offers key business and functional benefits:

- Limits lateral movement within the data center
- Minimizes risks and the effect of security breaches
- Simplifies network traffic flows
- Uses the existing underlay network infrastructure
- Lowers capital expenditure and operating expenses
- Automates IT service delivery
- Securely enables business agility

7-11 Enforcing Zero-Trust with NSX Segmentation

NSX segmentation helps build a Zero-Trust approach to security by defining a security perimeter around each application.



NSX improves the security of today's modern workloads by preventing lateral movement using network segmentation. It is distributed, application-aware, and simple to operate.


Follow this process to secure a data center environment with NSX segmentation:

- Apply macro-segmentation to define virtual security zones.
- In those security zones, identify the assets (understand the application behavior and network flows).
- Implement micro-segmentation to secure every application.
- Implement contextual security policies.

You must always monitor the environment for changes or unexpected behavior and adapt the security policies.

7-12 Step 1: Creating Virtual Security Zones

Protect segments of the network by creating virtual security zones.



Using macro-segmentation to isolate environments improves the security of the data center. It prevents lateral movement between virtual zones.


Depending on their business structure and use cases, a security team typically chooses to segment environments that should not be able to directly communicate with each other. Examples include different business units (such as HR, Finance, and so on), partner environment, and production environments.

Macro-segmentation provides a more flexible solution compared to a traditional appliance-based firewall. It enables organizations to easily expand the number of zones needed and adapts easily to changing network and security requirements as the business evolves.

Macro-segmentation is the first step to the Zero-Trust journey. It starts defining security zones in the data center environment that will be further secured.

7-13 Step 2: Identifying the Application Boundaries

Identify the virtual machines and containers used by an application and the network traffic that is necessary for the application to function.



When the network is macro-segmented into virtual security zones, you can move to microsegmentation and secure applications in a virtual zone.

Before implementing micro-segmentation, assets must be identified:

- Define the application boundaries by identifying the VMs and containers that an application is
 using. Also, define the data, assets, applications, and services that need protection, such as:
 - Data: Credit Card Information, Protected Health Information (PHI), Personally Identifiable Information (PII), Intellectual Property
 - Assets: Point-of-sales terminals, manufacturing assets, IoT devices
 - Applications: Custom or off-the-shelf software
 - Services: DNS, DHCP, Active Directory, LDAP, and so on

 Identify how the traffic moves across the organization in relation to the previously defined boundaries.

You must look for the following types of traffic:


- External application traffic: Which user is connecting to the app, which shared services the application is using, and so on
- Internal application traffic: Communication between the different application components

A good understanding of the application footprint and the network traffic is the only way to determine and enforce policies that secure access to the data.

This identification process is tedious and time-consuming when performed manually. NSX Intelligence or VMware Aria Operations for Networks can be used to automate the discovery of the application boundaries.

7-14 Step 3: Implementing Micro-Segmentation

Use micro-segmentation to allow necessary network traffic.



After the application's composition and necessary network traffic are identified, firewall rules must be configured to allow the necessary network traffic.

NSX Distributed Firewall enables users to configure firewall rules from a single point, which are then pushed to all hosts that participate in the NSX network. The creation of rules can be automated with NSX Intelligence. NSX Intelligence can recommend distributed firewall rules based on the discovered traffic flows in the environment.


Micro-segmenting the network immediately reduces the attack surface of an application. It restricts the application to only communicate with the resources that it absolutely needs.

In this process, security tags and security groups can be used:

- Security tags enable labeling NSX objects.
- Security groups include different objects that are added both statically and dynamically and
 can be used as the source and destination of a firewall rule. Security groups can use security
 tags or other criteria (such as IP sets, MAC sets, segment ports, AD user groups, and so on)
 to group virtual machines together.

7-15 Step 4: Securing Through Context

Set up security policies to establish the behavior of virtual machines and containers.



This step secures the traffic and the context of the application by setting up policies based on the behavior of virtual machines and containers.

Security groups and tags can be used.

In the example, a security administrator wants to create a firewall policy to restrict network access to VMs with an earlier version of Windows:

- Tag VMs with the OS version used.
- Create a security group that gathers all VMs that do not match the OS version threshold.
- Create a security policy to restrict access to the members of that security group.

When a VM is created and does not meet the OS version criteria, it is automatically put in that security group and blocked by the firewall rule. This approach removes the need for checking each VMs individually.

Third-party services can be integrated to create more granular control. For more information about a list of NSX partners, see NSX Data Center Technology Partners at https://www.vmware.com/products/nsx/technology-partners.html.

7-16 Review of Learner Objectives

- Define NSX segmentation
- Recognize use cases for NSX segmentation
- Identify steps to enforce Zero-Trust with NSX segmentation

7-17 Lesson 2: NSX Distributed Firewall

7-18 Learner Objectives

- Identify types of firewalls in NSX
- Describe features of distributed firewalls
- Create firewall policies
- Configure firewall rules
- Configure firewall rule attributes: groups, services, and profiles
- Configure the distributed firewall to block malicious IPs
- Save, roll back, export, and import the distributed firewall configuration
- Describe the distributed firewall architecture

7-19 NSX Firewalls


NSX includes the distributed firewall (east-west) and the gateway firewall (north-south).

The distributed firewall is a hypervisor, kernel-embedded stateful firewall:

- It resides in the kernel of the hypervisor and outside the guest OS of the VM.
- It controls the I/O path to and from the vNIC.

The gateway firewall is used for north-south traffic between the NSX gateways and the physical network:

- It applies to Tier-0 and Tier-1 gateway uplinks and service interfaces.
- It is a centralized stateful service enforced on the NSX Edge node.




The gateway firewall can be used as perimeter, intertenant, or interzone firewall.

7-20 Features of the Distributed Firewall

The distributed firewall provides visibility and control for virtualized workloads and networks.

The distributed firewall has the following main features:


- Centralized configuration through the NSX UI or API
- Layer 2 stateless firewall rules
- Layer 3 stateless and stateful firewall rules
- Context-aware (layer 7) firewall rules
- Identity Firewall for Windows workloads
- FQDN Filtering
- Blocking Malicious IPs
- Time-based policies



7-21 Distributed Firewall: Key Concepts

Several key concepts apply to distributed firewalls:

- Security policy: A collection of firewall rules and service configurations.
- Firewall rule: A set of instructions that determine whether a packet should be allowed or blocked.
- Group: A construct with multiple objects statically or dynamically pooled together.
- Service: Defines a port and protocol combination and is used to specify the type of traffic
 to be blocked or allowed in firewall rules.
- Context profile: Inspects the layer 7 content of the packets to allow or deny them.



7-22 Overview of a Security Policy

A security policy is a collection of firewall rules. You can configure different types of security policies from the NSX UI.



The NSX UI enables you to configure several types of policies:

- Firewall policies: Used for configuring firewall rules to control north-south and east-west traffic
- Endpoint policies: Used for configuring Guest Introspection services and rules.
- IDS/IPS policies: You can use these policies to define Intrusion Detection and Prevention rules for east-west and north-south traffic.
- Malware Prevention policies: You can use these policies to define anti-malware rules for east-west and north-south traffic.
- Network Introspection policies: Used for configuring north-south and east-west traffic redirection rules.
- TLS inspection: Used for configuring TLS inspection rules for the north-south traffic.

7-23 Distributed Firewall Policy Categories

A Distributed Firewall policy is a collection of firewall rules applied to east-west traffic.

The NSX UI enables you to group distributed firewall policies into different categories.



The categories for distributed firewall rules include:

- Ethernet: All layer 2 policies. Layer 2 firewall rules are always evaluated before layer 3 rules.
- Emergency: Temporary firewall policies needed in emergency situations, such as blocking an attacker from attacking a web server.
- Infrastructure: Nonapplication policies specific to infrastructure components, such as AD, DNS, NTP, and so on.
- Environment: High-level policy groupings, for example, the production group cannot communicate with the testing group, or the testing group cannot communicate with the development group.
- Application: Specific and granular application policy rules, such as rules between applications
 or application tiers, or rules between microservices.
 - Application tiers: In a multitier application, the functionality of the application is separated into isolated functional areas, called application tiers.
 - Microservices: It is an architectural style that structures an application as a collection of services that are independently deployable.

Each of these categories has its own policies and rules. Firewall rules are enforced left to right and top to bottom across these categories.

You can reorder policies and rules in a specific category. However, you cannot move policies or rules across different categories.

You can configure policies and rules under relevant categories.

7-24 About Distributed Firewall Policies

A firewall policy includes one or more firewall rules, which contain specific instructions for managing various types of traffic.

A policy can contain stateful or stateless rules for enforcement.

Policies are enforced in a top-to-bottom order.



In a firewall policy, each firewall rule contains instructions that determine the following factors:

- Source and destination for the packet
- Whether a packet should be allowed or blocked
- Protocols that the packet can use
- Ports that the packet can use


Policies are used for multitenancy, such as creating specific rules for sales and engineering departments in separate policies.

A policy can be defined as enforcing stateful or stateless rules. Stateless rules are treated as traditional stateless access-control lists (ACLs).

7-25 Distributed Firewall Rule Processing within a Policy

Firewall rules are processed in a top-to-bottom order:

- You can move rules up and down within a policy.
- The first match triggers rule enforcement.
- Packets that do not match any other rule are matched by the default rule.



Firewall rules are enforced in the following ways:


- Like firewall policies, firewall rules are processed in the top-to-bottom order.
- Each packet is checked against the top rule in the rule table before moving down the subsequent rules in the table.
- The first rule in the table that matches the traffic parameters is enforced. Subsequent rules
 cannot be enforced because the search is terminated for that packet.

Because of this behavior, you must place the most granular policies at the top of the rule table.

Packets not matching other rules are enforced by the default rule. The default rule is originally set to the Allow action. This rule ensures that VM-to-VM communication is not broken during stagling or migration phases. To implement a Zero-Trust model where only the specified traffic is allowed, the action for the default rule should be changed to block, and firewall policies defined to specify all traffic to be allowed.

7-26 Applied To Field for the Policy


When creating a Distributed Firewall policy, you can define the scope of the policy. It can be the full DFW or a specific security group.



The Applied To field configured at the policy level overrides the Applied To field configured at the rules within it. You must configure the Applied To field either at the policy level or at the rule level, but not at both levels. For more granularity, consider configuring the Applied To field at the rule level.

7-27 Configuring Distributed Firewall Policy Settings


When creating a Distributed Firewall policy, you specify settings such as TCP Strict, Stateful, and Locked.



7-28 Configuring Time-Based Firewall Policies

You can configure security policies that are only valid for a specific period. You can specify the following parameters in the Time Window:

- Name
- Time zone (UTC or local to transport node)
- Frequency (weekly or one time)
- Recurring days
- Start and end date
- Start and end time



The From and Till parameters must be configured in 30-minute increments. For example, from 09:00 to 09:30 is a valid configuration. However, if a user configures an interval from 09:15 to 09:45, a configuration error appears in the UI.

Before configuring a time-based rule, you must configure NTP servers for the transport nodes.

7-29 Creating Distributed Firewall Rules

Rules are a set of criteria used to evaluate traffic flows. They contain instructions that determine whether a packet should be allowed, dropped, or rejected.



7-30 Configuring Distributed Firewall Rule Parameters

A distributed firewall rule includes parameters such as source, destination, service, and context profile. This rule defines the scope where the rules should be applied to and the action that should be taken on a rule match. It also provides an option of logging when the traffic matches a rule.



Several parameters can be defined when configuring a distributed firewall rule:

- Name: You can specify a name for the rule.
- Sources: You can use previously defined groups.
- Destinations: You can use previously defined groups.
- Services: You can specify a port and protocol combination.
- Context Profiles: You use context profiles to define context-aware or layer 7 rules.
- Applied To: Defines the scope of the rule. It can be the full DFW or a specific security
 group.
- · Action: You can select from the following firewall rule actions:
 - Allow
 - Drop
 - Reject
- ID: NSX Manager automatically assigns an ID to a rule when it is created. The ID can be used for troubleshooting. A user cannot modify the ID.

The order of firewall rules determines how the traffic is managed. You can drag the rules in the UI to change the order.

7-31 Specifying Sources and Destinations for a Rule

When specifying sources and destinations for a firewall rule, you can use an IP or MAC address or an object (such as a group). If you do not specify these parameters, they match **Any**.



Both IPv4 and IPv6 addresses are supported for sources and destinations options of the firewall rule. Multicast addresses are also supported.

7-32 Creating Groups

A group defines a collection of assets on which security policies and rules can be applied.

A group can contain VMs, VIFs, segments, segment ports, IP and MAC addresses, AD user groups, and physical servers.



Before creating a group that includes AD users, you must add an AD domain to NSX Manager. You add this domain through the NSX UI by navigating to System > Configuration > Identity Firewall AD > ADD ACTIVE DIRECTORY.

The main use case for creating a group that includes AD users is to configure identity-based firewall rules.

7-33 Adding Members and Member Criteria for a Group

Groups can be defined by using dynamic or static membership criteria:

- Dynamic group inclusion for VMs can be based on tags, machine names, OS names, or computer names.
- Static group inclusion criteria apply to VMs, VIFs, segments, segment ports, IP sets, MAC sets, AD user groups, physical servers, and nested groups.




7-34 Creating Groups Based on Tags

Tags are typically used to automate security policy enforcement for new applications being provisioned.

Security administrators can assign one or multiple tags to workloads based on a given criteria.


These tags can then be used to create dynamic security groups for use in firewall rules.



7-35 Specifying Services for a Rule

When configuring distributed firewall rules, you specify one or more services.

Services contain the port and protocol definition for network traffic.




NSX includes an extensive list of predefined services. You cannot modify or delete these services. However, you can create additional services to meet your communication requirements.

You can create a service while configuring a distributed firewall rule. Alternatively, you can create additional services by navigating to **Inventory** > **Services** > **ADD SERVICE**.

7-36 Adding a Context Profile to a Rule





NSX Manager includes a list of predefined context profiles. You can also configure custom context profiles for your firewall rules. Layer 7 firewall rules can be defined only in a stateful firewall policy.

Alternatively, you can create context profiles by navigating to **Inventory** > **Profiles** > **Profiles** > **Context Profiles** > **ADD CONTEXT PROFILE**.

7-37 Configuring Context Profile Attributes

When creating a context profile for a distributed firewall rule, you configure two attributes: domain name and application ID.




A context profile defines context-aware attributes, including application ID, domain name, and subattributes such as application version or cipher set.

Context profiles for distributed firewall rules include the following main attributes:

- DOMAIN_NAME: You can choose from a static list of fully qualified domain names (FQDNs)
 or add your own FQDN.
- APP_ID: You can choose from a list of preconfigured applications. You cannot add
 applications. Examples include FTP, SSH, and SSL. Certain applications allow users to
 specify subattributes. For example, when choosing SSL Application, you can specify the
 TLS_VERSION and the TLS_CIPHER_SUITE. For CIFS, you can specify the
 SMB_VERSION.

7-38 Custom FQDN Filtering


You can add your own fully qualified domain name (FQDN) for custom filtering.



Alternatively, you can create context profiles by navigating to **Inventory > Profiles > Attribute**Types > FQDNs > ACTIONS > Add FQDN.

7-39 Setting the Scope of Rule Enforcement

The Applied To attribute optimizes the resource utilization on the ESXI hosts. It also helps in defining targeted policies at specific zones or tenants without affecting the policy defined on other zones or tenants.




The appropriate use of the Applied To field is paramount to optimize resource utilization on the transport nodes and to avoid scalability issues. You must configure the Applied To field in a distributed firewall rule to match the security groups used as the source and destination.

Zones can be compared to departments in an organization, whereas tenants can be compared to different organizations managed by a service provider. In the NSX UI, you can create groups to represent zones or tenants.

7-40 Specifying the Action for a Rule


You configure the following actions in a distributed firewall rule:

- Allow: Allows all traffic with the specified source, destination, and protocol.
- Drop: Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source and destination systems.
- Reject: Rejects packets with the specified source, destination, and protocol. Rejecting a
 packet is a more graceful way to deny a packet, because it sends a destination unreachable
 message to the sender.



7-41 Distributed Firewall Rule Settings






You can configure several settings for distributed firewall rules:

- Logging: You can turn logging off or on. Logs are stored in the /var/log/dfwpktlogs.log file on ESXI hosts.
- Direction: This setting matches the direction of traffic from the point of view of the destination object.
 - Considering VM as an object, in means that only traffic to the VM is checked, Out
 means that only traffic from the VM is checked, and in-Out, means that traffic in both
 directions is checked.
- IP Protocol: IPv4, IPv6, and IPV4-IPv6 protocols are supported.
- Log Label: Log labels can be used to identify a rule when analyzing the log files.

7-42 Blocking Malicious IPs in the Distributed Firewall

The Block Malicious IPs feature in NSX Distributed Firewall uses the Malicious IPs File or Malicious IPs Feed to filter the traffic in NSX.



The Block Malicious IPs feature in NSX Distributed Firewall was introduced in NSX 4.0.0.1.


Traffic between any Internet resource with a malicious IP address and the internal resources is blocked by the distributed firewall by referring to the Malicious IPs Feed or Malicious IPs File.

The distributed firewall allows only traffic between an internet resource with a legitimate IP address and the internal resources. A legitimate IP address does not have an entry in the Malicious IPs Feed or Malicious IPs File.

7-43 Enable or Disable Malicious IP Feeds

You can enable or disable the download of Malicious IP Feeds from NTICS in the NSX UI.

In the navigation pane on the left, select Security > Policy Management > Distributed Firewall > Category Specific Rules. Then, from the Actions drop-down menu, select General Settings.



In the Settings dialog box, click the Malicious IP Feeds tab and turn on the Auto Update Malicious IPs toggle.

You can download the latest feed in the Settings dialog box. The automatic download frequency is 12 hours.

7-44 Default Malicious IP Group


A group called DefaultMaliciouslpGroup is created by default, it contains all the Malicious IPs obtained from the Malicious IPs Feed as its members.



You cannot delete DefaultMaliciousIPGroup, but you can edit this group.

7-45 Creating a Custom Malicious IP Group

You can create a custom Malicious IP group by either adding exception members or exception IP addresses to the default Malicious IP group.



Next to Group Type, you must select IP Addresses Only.

Next to Add Pre-defined Malicious IPs, turn on the **Yes** toggle to include all the Malicious IPs from the Malicious IPs Feed.

The Exception Members and Exception IP Addresses tabs contain the IP addresses that you want to exclude from the pre-defined malicious IP address list during the distributed firewall rule processing.

7-46 Default Malicious IP Block Rules

A default policy with the name Default Malicious IP Block Rules is created in the infrastructure category of the distributed firewall rules.



The first rule is to drop any traffic initiated by a malicious IP address from an external network to the internal network.

The second rule is to drop any traffic from the internal network to a malicious IP address in an external network.

Both the rules are enabled by default in a greenfield environment and disabled in a brownfield environment.

7-47 Configuring Rules to Block Malicious IPs


You can create a custom malicious IP address block rule and attach your custom Malicious IP group to It.



The configured rule drops any traffic from any internal source going to any destination IP address that is included in the group named Custom Malicious IP Group.

7-48 Saving and Viewing the Distributed Firewall Configuration

You can save and view distributed firewall configurations. Every time you publish a distributed firewall rule, a draft of the configuration is saved automatically.




From this view, you can see a timeline with all saved distributed firewall configurations. The following types of saved items are available:

- Auto saved: Drafts automatically saved by the system Immediately after distributed firewall
 changes are published. This feature is enabled by default but can be disabled if required.
 Rolling back to the previous configuration requires reverting to the previously published
 autosave.
 - You can disable the Auto saved feature in the NSX UI by navigating to Security >
 Policy Management > Distributed Firewall > ACTIONS > General Settings and turn off
 the Auto Save Drafts toggle.
- Saved by others: Distributed firewall configurations saved by other users different from the
 user currently logged in to the system.
- Saved by me: Distributed firewall configurations saved by the user currently logged in to the system.

7-49 Rolling Back to a Saved Distributed Firewall Configuration


You can roll back to a previously saved distributed firewall configuration.



When you click the name of the saved configuration in the histogram, a new wizard displays details about the distributed firewall configuration on the top, including name, description, and creation date. The bottom part of the screen displays the differences between the saved configuration and the last published configuration.

7-50 Distributed Firewall Configuration Export and Import






While exporting a firewall configuration, you must provide a passphrase. This passphrase is needed when importing this configuration into the firewall. After the export is complete, you can download the firewall configuration.

While importing a firewall configuration with the passphrase, you must also provide a name for this configuration. All the imports are saved as drafts in the firewall.

7-51 Distributed Firewall Architecture

The high-level distributed firewall workflow includes the following steps:


- Users configure distributed firewall policies through the NSX UI or API.
- The policy role processes these policies.
- Distributed firewall policies are then pushed to the manager role and persisted.
- The manager role forwards the distributed firewall rule configuration to the central control
 plane (CCP).
- The CCP forwards the configuration to the LCP (nsx-proxy) through the Appliance Proxy Hub (APH).
- The host transport nodes (ESXI) store the distributed firewall configuration and configure the datapath accordingly.
- The transport nodes send rule statistics and status to NSX Manager.



7-52 Distributed Firewall Architecture: ESXi

On an ESXI host, the distributed firewall includes several components:

- nsx-proxy receives the configuration changes from the CCP and configures datapath modules.
- Datapath modules include the following components:
 - VSIP: Receives firewall rules and downloads them on each VM's vNIC
 - VDPI: Performs L7 packet Inspection
- Stats Exporter collects flow records from the distributed firewall data plane kernel modules and generates rules statistics.
- nsx-proxy passes rules statistics and real-time data to the management plane.



The following datapath modules are responsible for distributed firewall rule processing:

- VMware Internetworking Service Insertion Platform (VSIP): This module is the main part of the distributed firewall kernel module that receives the firewall rules and downloads them on each VM's vNIC.
- VMware Deep Packet Inspection (VDPI): This deep packet inspection module daemon in the
 user space is responsible for L7 packet inspection. VDPI can identify application IDs and
 extract context for a traffic flow.

L7 rules, like the remaining DFW rules, are programmed into VSIP. VSIP forwards L7 packets to VDPI, which inspects and extracts the L7 information from the packets and returns them to VSIP.

Stats Exporter collects flow records from the VSIP kernel module and generates rule statistics.


7-53 Distributed Firewall Rule Processing: ESXi

vmware-sfw is a software construct where distributed firewall rules are stored and enforced.

vmware-sfw maintains a rule table and a connection table.

When a distributed firewall rule is configured, its information is stored in the rule table. After the initial rule configuration, the traffic is managed as follows:

- A lookup is performed in the connection table to check whether a connection exists.
- If the connection is not present, the packet is matched against the rule table.
- If the packet is allowed and the traffic type is stateful, a connection entry is created in the connection table.
- All subsequent packets for the same connection are serviced directly from the connection table. Stateless packets are always matched against the rule table.



The vSphere ESXI network IOChain is a framework that enables you to insert functions into the network datapath.

The IOChain framework is used by NSX to host vmware-sfw, which is a software construct where distributed firewall rules are stored and enforced. vmware-sfw maintains the following types of tables:

- Rule Table: Locally stores all applicable firewall rules
- Connection Table: Caches flow entries for stateful rules with an action

7-54 Lab 11: Configuring the NSX Distributed Firewall

Create NSX distributed firewall rules to allow or deny the application traffic:

- Prepare for the Lab
- Test the IP Connectivity
- Create Security Tags
- Create Security Groups based on Tags
- Create Distributed Firewall Rules
- Test the IP Connectivity After the Firewall Rule Creation
- Prepare for the Next Lab

7-55 Review of Learner Objectives

- Identify types of firewalls in NSX
- Describe features of distributed firewalls
- Create firewall policies
- Configure firewall rules
- Configure firewall rule attributes: groups, services, and profiles
- Configure the distributed firewall to block malicious IPs
- Save, roll back, export, and import the distributed firewall configuration
- Describe the distributed firewall architecture

7-56 Lesson 3: Use Case for Security in Distributed Firewall on VDS

7-57 Learner Objectives


- · List the distributed firewall on VDS requirements
- Configure the distributed firewall on VDS
- Validate the distributed firewall on VDS configurations

7-58 About Distributed Firewall on VDS

Distributed firewall on VDS enables NSX security features on existing vSphere distributed port groups (DVPG) without changing the vSphere platform.

Distributed firewall on VDS provides the following benefits:

- Removes the need for migrating workloads when configuring NSX security
- Implements the same security policies and rules for network objects irrespective of whether they are created or owned by NSX Manager or vSphere
- Enables the security administrator to deploy the distributed firewall and other security features without involving the networking administrators



Distributed firewall on VDS enables NSX Security features on workloads attached to a distributed port group (DVPG) managed by vCenter Server. Earlier versions of NSX required VMs to be attached to segments managed by NSX (VLAN or overlay) to take advantage of distributed security functions. This functionality has now been expanded to cover DVPG managed by vCenter Server.

Distributed firewall on VDS is also available in vSphere 6.7.

Distributed firewall on VDS is implemented on the NSX Management Plane based on information reported by the vCenter Server inventory.

7-59 Supported Features

Several features are supported when you prepare your vCenter Server cluster for NSX Security.



For all supported features, see NSX Administration Guide at https://docs.vmware.com/en/VMware-NSX/4.0/administration/GUID-E9FBE567-D136-41AF-B8D6-AE95416F4229.html#GUID-E9FBE567-D136-41AF-B8D6-AE95416F4229.

7-60 Distributed Firewall on VDS Requirements

installing the NSX distributed firewall on an existing VDS has the following requirements:

- The NSX version must be 3.2 or later.
- ESXI hosts are not prepared for NSX.
- All hosts in the cluster must have identical VDS configurations.
- The vCenter Server system must be registered as a Compute Manager in NSX Manager.

Supported VDS and vSphere versions:

VDS Version	vSphere Version
6.6.0	vSphere 6.7 and later
7.0.0	vSphere 7.0 Update 1 and later
7.0.2	vSphere 7.0 Update 2 and later
7.0.3	vSphere 7.0 Update 3 and later
8.0.0	vSphere 8.0.0 and later

This feature is only configurable with the NSX API or by using the Quick Start installation in the NSX UI or through the vCenter NSX plug-in.

VDS can span multiple vSphere clusters. The NSX security mode is enabled only on some clusters.

7-61 Installation Workflow

You follow these steps to install the NSX distributed firewall on a VDS with existing dyportgroups.



7-62 Preparing the Cluster for Security

To prepare clusters for NSX Security:

- Navigate to System > Configuration > Quick Start > Prepare Clusters for Networking and Security > GET STARTED.
- Select the clusters that you want to install.
- Click INSTALL NSX > Security Only.




All hosts in the cluster must have identical VDS configuration to perform an NSX Quick Install.

You can also configure the host for security only with a REST API call.

7-63 Validating the Security Cluster Preparation from the NSX UI

You can validate the Security Cluster Preparation by navigating to System > Configuration > Fabric > Nodes > Host Transport Nodes.



You check the following parameters:

- The node status is UP.
- The NSX Configuration status is Success.
- The NSX version is correct.
- One host switch per VDS was created and the Type is VDS.

Because the cluster is not prepared for NSX networking features, it is not expected to have tunnels configured on the host.

7-64 Transport Node Preparation

All the ESXI hosts of the selected clusters are automatically configured as transport nodes in the back end.

During the transport nodes preparation, the following objects are created:

- Autoconfigured transport node profiles
- VLAN transport zones
- Discovered segments

When the ESXI hosts are prepared for NSX security, vCenter Server performs a full inventory sync with NSX Manager to share its objects.

7-65 Autoconfigured Transport Node Profile

A separate transport node profile is created for each cluster.



You can check the autocreated transport node profile by navigating to System > Fabric > Profile > Transport Node Profile.

A separate transport node profile is created for each cluster even if multiple clusters have identical VDS configuration.

The created transport node profiles are neither configurable nor editable.

7-66 VLAN Transport Zones

A VLAN transport zone is created for every VDS.




You can check the autocreated transport zone by navigating to **System > Fabric > Transport Zones > Transport Zones**.

The created transport node profiles are neither configurable nor editable.

7-67 Discovered Segments (1)

Discovered segments are automatically created by the vCenter Server Inventory module. For each vSphere DVPG, a distributed port group is created in NSX with its own NSX ID.




A single instance of a port has two IDs: a vSphere ID and an NSX ID. The NSX Control plane uses the NSX ID.

Discovered segments can be consumed by the NSX security features such as the distributed firewall.

7-68 Discovered Segments (2)

You can access the discovered segments by navigating to **Networking > Connectivity >** Segments > Distributed Port Groups.



Discovered segments do not appear in the NSX UI as segments, but they appear as distributed port groups.

VLAN tags are read and added from vCenter Server.

The vCenter Server Inventory is monitored. Any change that occurs (DVPG added, host added, and so on) is reflected in NSX.


7-69 Configuring Segment Profiles

You can configure IP discovery, segment security, and Spoofguard from the NSX UI.



7-70 Using VDS Attributes to Define NSX Groups

You can use Distributed Port, Distributed Port Groups, and their corresponding tags to define NSX Group membership.



You can select distributed port group and distributed port as members and membership criteria. You can use these groups when configuring distributed firewall rules.

7-71 Review of Learner Objectives

- List the Distributed Firewall on VDS requirements
- Identify the steps to configure Distributed Firewall on VDS
- Validate the configuration

7-72 Lesson 4: NSX Gateway Firewall


7-73 Learner Objectives

- · Describe the functions of the gateway firewall
- Explain the purpose of a gateway policy
- Create gateway firewall policies and rules
- Describe the gateway firewall architecture

7-74 About the Gateway Firewall

The gateway firewall has the following characteristics:

- Stateful and stateless firewall for north-south traffic
- Independent of the distributed firewall
- Implemented on both Tier-0 and Tier-1 gateways and requiring the SR component of the router
- Enforced on the northbound-facing interface of the gateway



The NSX gateway firewall provides essential perimeter firewall protection that can be used in addition to a physical perimeter firewall. The gateway firewall supports stateless and stateful firewall rules. The gateway firewall works independent of the distributed firewall. A user can consume the gateway firewall using either the UI or REST API framework provided by NSX Manager. The gateway firewall configuration is similar to the distributed firewall policy. This configuration is defined as a set of individual rules in a policy. Like the distributed firewall, the gateway firewall rules can use tagging and groups to build policies.

The gateway firewall is a centralized firewall implemented on the northbound-facing interface of the gateway (Tier-0 uplinks and Tier-1 RouterLinks). The firewall is implemented on a Tier-0 or Tier-1 service router (SR) component that is hosted on the NSX Edge node.


The service router component of a Tier-0 or Tier-1 gateway provides north-south routing functionality and centralized services, such as NAT, and so on.

From NSX 4.0.1, the Tier-0 gateway firewall supports stateful firewall fiftering with both activeactive and active-standby high availability modes.

The gateway firewall service is part of the NSX Edge node for both bare-metal and VM form factors. The gateway firewall is useful in developing PCI zones, multitenant environments, or DevOps-style connectivity without forcing the intertenant or interzone traffic onto the physical network. The gateway firewall datapath uses the Data Plane Development Kit (DPDK) framework supported on NSX Edge to provide better throughput.

7-75 Predefined Gateway Firewall Categories

The gateway firewall includes predefined categories on the **All Shared Rules** tab, where rules across all gateways are visible.



The gateway firewall includes several predefined categories for rules.

- Emergency: Used for quarantine and can also be used for Allow rules.
- System: Automatically generated by NSX and specific to internal control plane traffic, such as BFD rules, VPN rules, and so on.
- Pre Rules: Globally applied across all NSX gateways.
- Local Gateway: Rules specific to a particular NSX gateway.
- Auto Service: Autoplumbed rules applied to the data plane.
- Default: Rules that define the default gateway firewall behavior.

Categories are evaluated from left to right. Each category can have its own rules, which are evaluated top to bottom.

7-76 Gateway Firewall Policy


A Gateway Firewall policy includes one or more individual firewall rules and is applied to northsouth traffic.

Gateway policies can be applied to Tier-0 and Tier-1 gateways and their interfaces.



7-77 Configuring Gateway Firewall Policy Settings

To create a Gateway Firewall policy, you assign a policy name and configure the settings. You can also set a Time Window so that the policy is only applicable during the specified period.



To create a Gateway Firewall policy, you assign a policy name.


You can configure the following settings when creating a Gateway Firewall policy:

- TCP strict: In certain circumstances, the firewall might not see the TCP three-way
 handshake for a particular flow (that is, due to asymmetric traffic). By default, the firewall
 does not enforce the need to see a three-way handshake and will pick up sessions that are
 already established. TCP Strict can be enabled per section to turn off midsession pickup.
 When enabling the TCP Strict mode for a particular firewall policy and using a default ANYANY Block rule, packets that do not complete the three-way handshake connection
 requirements and that match a TCP-based rule in this policy section are dropped.
- Stateful: When this option is enabled, the gateway firewall performs stateful packet
 inspection and tracks the state of network connections. Packets matching a known active
 connection are allowed by the firewall, and packets that do not match are inspected against
 the gateway firewall rules.
- Locked: This setting allows you to lock a policy while making configuration changes so that
 others cannot make modifications at the same time.

You can also set a Time Window so that the policy is only applicable during the specified period. For this feature, the NSX Edge nodes need to have an NTP server configured.

7-78 Configuring Gateway Firewall Rules


You create one or more rules in the policy to allow, drop, or reject traffic.



Context Profiles are only available on Tier-1 gateways.

7-79 Configuring Gateway Firewall Rules Settings

You can specify the logging, direction, and IP protocol for the gateway firewall rule. A firewall rule must be published for it to take effect.



The rules are logged in the Syslog file of the edge node.

You can access these logs from the CLI:


In the rule settings, you can also choose the direction and the IP protocol:

- In: The rule only checks traffic to the gateway Interface.
- Out: The rule only checks traffic from the gateway interface.
- In-Out: The rule checks traffic from both directions.

7-80 Gateway Firewall Architecture

The gateway firewall workflow is as follows:

- Users configure gateway policies from the NSX UI.
- The policy role processes gateway policies.
- Gateway policies are sent to the manager role, which validates and forwards them to the CCP.
- The CCP distributes the firewall configuration through APH to the relevant edge nodes.
- nsx-proxy receives the firewall configuration from the CCP and configures the edge data path.
- The Stats Exporter collects flow records from the datapath and generates rule statistics.
- nsx-proxy reports the firewall rules statistics and status to the management plane.




7-81 Gateway Firewall Rule Processing

On an NSX Edge node:

- Gateway Firewall rules are programmed into rule classifier.
- The flow table tracks established connections for stateful firewall rules.

When a Gateway Firewall rule is configured, its information is stored in a rule table in the Rule Classifier module. After the initial rule configuration, the traffic is managed as follows:

- A lookup is performed in the flow table to check whether a connection exists.
- If the connection is not present, the packet is matched against the rule table in Rule Classifier.
- If the type of traffic is stateful, an entry is created in the flow table.
- All subsequent packets for the same connection are serviced directly from the flow table.
 Stateless packets are always matched against the rule table.



The Rule Classifier maintains stateful and stateless rules for the following features:

- Gateway Firewall
- NAT
- IPSec
- Service Insertion

The flow table is responsible for tracking established connections for stateful firewall rules, and NAT edge services. When a new connection is made, the first packet is matched against the flow table to determine if a session exists.

A rule classifier and flow table are created for each gateway. If two gateways are present in the NSX Edge node, two rule classifier instances and flow tables are created: one for each gateway.

7-82 Lab 12: Configuring the NSX Gateway Firewall

Configure and test the NSX gateway firewall rules to control north-south traffic:

- Prepare for the Lab
- Test SSH Connectivity
- Configure a Gateway Firewall Rule to Block External SSH Requests
- Test the Effect of the Configured Gateway Firewall Rule
- Prepare for the Next Lab

7-83 Review of Learner Objectives

- Describe the functions of the gateway firewall
- Explain the purpose of a gateway policy
- Create gateway firewall policies and rules
- Describe the gateway firewall architecture

7-84 Key Points

- Zero-Trust is a security model that does not automatically trust entities in the security perimeter.
- Macro-segmentation is the process of dividing data center infrastructure into smaller zones.
- Micro-segmentation helps to build a Zero-Trust approach to security by defining a security perimeter around each application.
- The distributed firewall is a hypervisor kernel-embedded stateful firewall.
- The distributed firewall resides outside the VM guest OS and controls the I/O path to and from the vNIC.
- The gateway firewall, also called the perimeter firewall, protects traffic from physical environments.

Questions?

Module 8

NSX Advanced Threat Prevention

8-2 Importance

NSX IDS/IPS, NSX Malware Prevention, NSX Intelligence, and NSX Network Detection and Response provide visibility and protection against advanced threats in your network. As a security administrator, you must learn to properly configure these features to successfully prevent malicious attacks against your environment.

8-3 Module Lessons

- NSX Intrusion Detection and Prevention
- NSX Application Platform
- NSX Malware Prevention
- NSX Intelligence
- NSX Network Detection and Response

8-4 Lesson 1: NSX Intrusion Detection and Prevention


8-5 Learner Objectives

- Explain IDS/IPS and its use cases
- Define the IDS/IPS terminologies
- Describe the IDS/IPS architecture
- Configure IDS/IPS
- Interpret IDS/IPS events

8-6 About NSX Distributed IDS/IPS

NSX Distributed IDS/IPS uses real-time deep packet inspection to identify and prevent attempts at exploiting vulnerabilities in your applications:

- Protects virtual workloads from malicious traffic
 - Prevents lateral threat movement in the east-west traffic
- Uses signatures to identify malicious traffic patterns
- Is implemented as a distributed solution across multiple ESXI hosts



Using real-time deep packet inspection, NSX Distributed IDS/IPS performs the following tasks:


- Protects east-west traffic and prevents lateral threat movement: The objective of many
 malicious attacks is not solely to penetrate the network. After the attackers are in, they
 often pivot through multiple systems and explore the network to find their main target and
 gain access to it. NSX Distributed IDS/IPS recognizes and prevents lateral movement across
 the network when perimeter security is compromised.
- Uses signatures to identify malicious traffic patterns: NSX Distributed IDS/IPS uses an
 external cloud-based signature store to remain up to date with known malicious activity,
 including zero-day vulnerabilities. It helps protect against both L4 and L7 attacks.
- Is implemented as a distributed solution across multiple ESXI hosts: Similar to the distributed firewall architecture, NSX Distributed IDS/IPS is implemented as a kernel module in the ESXI hosts. This distributed architecture significantly reduces hairpinning by processing traffic that is closer to the source.

8-7 Use Cases for NSX Distributed IDS/IPS

NSX Distributed IDS/IPS protects against malicious activity, including the following activities:

- Exploits of known application-level vulnerabilities
- Application denial of service
- Lateral movement
- Client-side and server-side exploits

With NSX Distributed IDS/IPS, security administrators can augment or replace discrete appliances.



Application denial of service and client and server exploits have the following characteristics:

- Application denial of service: A denial-of-service attack floods systems, servers, or networks with traffic to exhaust resources and bandwidth. As a result, the system is unable to fulfill legitimate requests.
- Client-side and server-side exploits; Client-side attacks exploit the trust between users and the website or server that they visit. Common client-side and server-side exploits are as follows:
 - Spoofing: Tricking the user into believing a website or server is legitimate.
 - Cross-site scripting (XSS): Executes code in the user's web browser with the aim of controlling a user's session or stealing login credentials.

With NSX Distributed IDS/IPS, security administrators can replace or augment discrete appliances. By using native IDS/IPS capabilities in NSX, you can replace traditional IDS/IPS appliances, including standalone, firewall-based, or virtual host-based solutions. You might also decide to keep the traditional IDS/IPS appliances, while using NSX Distributed IDS/IPS for additional east-west traffic protection.

8-8 About Behavior-Based IDS/IPS

Behavior-based IDS/IPS helps to detect unusual traffic, malicious attacks, and security breaches in the network when compared to a baseline of normal traffic.

Behavior-based IDS/IPS does not require a separate installation because it is part of the NSX Distributed IDS/IPS implementation.

Behavior-based IDS/IPS helps in the following scenarios:

- Identifying periodic callback behavior
- Identifying a client with a high failure rate in SSH authentication with a server
- Tunneling of network data over anonymous proxies such as TOR

Intrusion detection focuses on the use of knowledge-based signatures. These signatures incorporate specific knowledge about an attack.

Behavior-based IDS/IPS attempts to identify anomalous behavior by pinpointing potentially interesting events that are different or unusual compared to a baseline of normal traffic. Behavior-based IDS/IPS also helps to detect and prevent possible zero-day attacks.

A few examples of such alerts:

- Identifying periodic callback behavior in a given flow: A variety of remote access Trojan (RAT) toolkits expose this type of behavior when checking in with a command-and-control server. RATs are a type of malware threat where an attacker takes control of your computer.
- Identifying a client with a high failure rate in SSH authentication with a server (can indicate a credential enumeration attack).
- Tunneling of network data over anonymous proxies such as TOR (not necessarily malicious, but unusual in an enterprise environment). TORs aim to conceal users' identities and their online activity from surveillance and traffic analysis by separating identification and routing.

8-9 Requirements for NSX Distributed IDS/IPS

Before using NSX Distributed IDS/IPS, administrators must consider the following factors:

- The NSX Distributed IDS/IPS components are installed as part of the host preparation.
- NSX Distributed IDS/IPS can be enabled at the vSphere cluster level and for standalone ESXI hosts.
- You can configure NSX Manager to download intrusion detection signatures from the Internet or manually download and upload the signatures to NSX Manager.
- The NSX environment must be configured with a valid license for NSX Distributed IDS/IPS.

For additional information about the type of licenses that are valid for NSX Distributed IDS/IPS, see the VMware NSX Datasheet at

https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/nsx/vmware-nsx-datasheet.pdf.

8-10 About IDS/IPS Signatures

An IDS/IPS signature contains metadata that is used to identify an attacker's attempt to exploit a known operating system or application vulnerability.

NSX Manager downloads IDS/IPS signatures daily from a cloud-based signature repository.

Signatures are classified into the following severity categories:

- Critical
- High
- Medium
- Low
- Suspicious



An IDS/IPS signature contains metadata that is used to identify an attacker's attempt to exploit a known operating system or application vulnerability. Such metadata provides context about the attempt, such as the affected product, attack target, and so on.

IDS/IPS signatures are matched against traffic headers by using regular expressions.

NSX Manager downloads IDS/IPS signatures daily.

IDS/IPS signatures are classified into severity categories based on their Common Vulnerability Scoring System (CVSS) score.

8-11 About IDS/IPS Profiles

An IDS/IPS profile defines the IDS signatures that are included or excluded from detection.



The default IDS profile is configured to include all signatures that are labeled as critical.

8-12 About IDS/IPS Policies and Rules

An IDS/IPS policy is a collection of IDS/IPS rules. An IDS/IPS rule contains a set of instructions that determine which traffic is analyzed, including values for the following parameters:

- Sources and Destinations
- Services
- Security Profiles (IDS/IPS profile)
- Applied to
- Mode



IDS rules can be applied to the distributed firewall or to specific groups.

NSX includes the following modes for an IDS/IPS rule:

- Detect Only: Detects signatures and does not take any action on the traffic. It records and logs the Intrusion.
- Detect & Prevent: Detects signatures and performs the action specified by the security administrator. Available actions are alert, drop, and reject.


Different rules can be configured with different modes.

You typically start with the Detect Only mode. After tuning for false positives, you change to Detect & Prevent.

8-13 IDS/IPS Signature Curation

The NSX IDS curator engine combines the IDS signatures from Trustwave, Secureworks, and Lastline into a single signature set, which it pushes, as an NSX IDS bundle, to NSX Threat Intelligence Cloud.

NSX Threat Intelligence Cloud forwards the NSX IDS bundle to NSX Manager for consumption in the NSX environment.



The NSX IDS curator engine combines the IDS signatures from Trustwave, Secureworks, and Lastline into a single signature set and pushes it as an NSX IDS bundle to NSX Threat Intelligence Cloud. NSX Threat Intelligence Cloud forwards the NSX IDS bundle to NSX Manager for consumption in the NSX environment.


The NSX IDS curator engine performs the following tasks in the back end:

- Identifies low-quality signatures:
 - Avoids false positives.
 - Identifies noisy signatures.
- Identifies redundant signatures:
 - If two signatures match the same traffic, only one signature is kept.
- Adds signature metadata for the NSX UI fields:
 - MITRE ATT&CK tactics and techniques.
 - Impact, severity, and confidence levels.
- Ensures consistent values for existing signature metadata:
 - CVSS scores, signature severity, attack target, and classtype.

8-14 NSX Distributed IDS/IPS Architecture

NSX Distributed IDS/IPS operates as follows:

- NSX Manager downloads curated IDS/IPS signatures from NSX Threat Intelligence Cloud.
- Users configure IDS/IPS profiles and rules.
- Policy stores the IDS/IPS configuration and passes it to NSX Manager.
- NSX Manager passes the information to the central control plane (CCP).
- CCP pushes the IDS/IPS configuration to hosts through the appliance proxy hub.
- The ESXI hosts store the signature information locally and configure the datapath.
- The ESXI hosts collect traffic data and send events to NSX Manager.
- NSX Manager parses and displays these events in the UI.



IDS signatures are written into the IDS module in the datapath, and IDS rules are stored in the VSIP module.

VSIP evaluates traffic against IDS rules. If a match is found, the packet is sent to IDS.

The IDS module evaluates the packets against IDS signatures.

Distributed firewall rules are always evaluated before distributed IDS/IPS rules. If a distributed firewall rule rejects a traffic flow, this traffic is never evaluated by IDS/IPS.


8-15 Configuring NSX Distributed IDS/IPS

To enable distributed intrusion detection and prevention for standalone hosts or clusters, you select Security > IDS/IPS & Malware Prevention > Settings > Shared.



8-16 Configuring IDS/IPS Signatures


On the IDS/IPS tab, you can configure and manage signatures under Settings.



If the **Auto Update** toggle is On, signatures are automatically applied to the ESXI hosts after they are downloaded from the cloud. If the toggle button is Off, the signatures are stopped at the listed version.

8-17 Global Intrusion Signature Management

With Global Intrusion Signature Management, you can override the default action for a given signature and globally disable signatures that are not relevant to your environment.



If a signature is disabled globally, it is removed from custom profiles. You cannot include the disabled signature in newly created custom profiles.


All signatures are preconfigured with a default action that is recommended by VMware. You can override this action globally or per profile. The following actions are available:

- Alert: This action is typically used in new deployments or for new signatures.
- Drop and Reject actions are commonly used in the following circumstances:
 - High-fidelity signatures with no false positives
 - High-impact exploits
 - After a signature is deployed in alert-only mode

Dropping a packet is a silent action with no notification to the source and destination systems. Rejecting a packet is a more graceful way to deny a packet because it sends a destination unreachable message to the sender. The drawback with rejecting a packet is that it can notify a potential attacker of the defense invoked.

8-18 Configuring Custom IDS/IPS Profiles

Using custom IDS/IPS profiles, you specify the IDS signatures that you want to include or exclude for detection based on their severity, attack type, attack target, CVSS, and affected products.



Creating granular workload-specific profiles reduces noise and false positives in your environment.

You create custom IDS profiles by selecting Security > IDS/IPS & Malware Prevention > Profiles.

You configure the IDS signatures that you want to include or exclude for detection based on their severity and more granular criteria:

- Attack Types: Categorizes signatures by attack techniques such as Trojan activity or attempted denial of service (DoS). The types align with the MITRE ATT&CK framework.
- Attack Targets: Broad category of possible attack targets such as IoT, mobile client, or networking equipment.
- CVSS: Include or exclude signatures based on their CVSS score range (none, low, medium, high, and critical).
- Products Affected: Signatures specific to vulnerable applications or operating systems.

Globally disabled signatures are not available for inclusion in profiles.

You can also click the **Manage signatures for this profile** link to disable individual profile signatures that might not be relevant to your workloads, or to override the global action configured for a given signature.


Suspicious IDS signatures are used to detect traffic anomalies in the network.

In a similar way, you create an IDS/IPS profile to specify the IDS signatures that you want to include or exclude for the detection of your north-south traffic.

8-19 Configuring IDS/IPS Rules

You specify how traffic is managed in your environment by configuring one or more rules with an IDS profile and the mode of operation.

As with distributed firewall rules, IDS/IPS rules are evaluated from top to bottom.



To create IDS policies and rules, you select Security > Policy Management > IDS/IPS & Malware Prevention > Distributed Rules.

The Sources, Destinations, Services, and Applied To values work in the same way as those in a distributed firewall rule.

IDS Profile specifies the group of signatures that the traffic is matched against.

Mode determines the action that is taken when a signature is triggered.

The following modes are available for an IDS/IPS rule:

- Detect Only: Regardless of the global or per-signature action, only alerts are generated, and no preventive action is taken. This mode is equivalent to an intrusion detection system.
- Detect & Prevent: The action that is specified for the given signature either globally or at the
 profile level is taken (alert, drop, reject). The action that is specified at the profile level overrides
 the action configured globally. This mode is equivalent to an inline intrusion prevention system.

When configuring IDS/IPS rules, do not use the drop action in a rule that is configured with a security profile that includes suspicious-level signatures. With this configuration, you guarantee that any abnormal traffic is inspected.

Like distributed firewall rules, IDS/IPS rules are evaluated from top to bottom. You must place the most hit rules at the top to avoid the unnecessary evaluation of subsequent rules.

8-20 Monitoring IDS/IPS Events (1)

To monitor IDS/IPS events, you select Security > Threat Event Monitoring > IDS/IPS.



The Events tab shows all intrusion attempts that are detected by the system:

- Critical severity signature events appear in red.
- High severity signature events appear in orange.
- Medium severity signature events appear in yellow.
- Low severity signature events appear in gray.
- Suspicious signature events appear in purple.

Administrators can filter events based on their severity. Free-form text is also available for further filtering of events.

IDS events are graphically represented by using a histogram. Security administrators can specify the period that they are interested in by adjusting the vertical lines in the diagram.

Each IDS/IPS event type is represented by a dot in the histogram. The size of the dot is proportional to the number of occurrences of an event.

Additional information about each type of event appears in a tabular format.

NSX Manager keeps the last 14 days of data or up to 1.5 million records.

8-21 Monitoring IDS/IPS Events (2)

Each event can be expanded to retrieve details about the intrusion attempt, including the attacker, victim, protocol, attack type, and so on.



Each event can be expanded to retrieve the following details:

- Signature ID
- Severity
- Details or description of the event
- Product affected
- Users affected.
- VMs affected

Additional information appears for the last occurrence of the event:

- Attacker and target IP addresses
- Bytes exchanged
- Attack and traffic flow direction

The Intrusion Activity diagram shows the occurrences for a particular signature event .

8-22 About North-South IDS/IPS

North-South IDS/IPS uses real-time deep packet inspection to identify and prevent attempts at exploiting vulnerabilities in your applications:

- Protects north-south traffic and prevents malicious traffic from entering your internal network
- Uses signatures to identify malicious traffic patterns
- Is enabled on a Tier-1 gateway



8-23 Use Cases for North-South IDS/IPS

You can configure North-South IDS/IPS for the following use cases:

- Detecting and preventing intrusions for north-south traffic based on signatures
- Detecting and preventing intrusion attempts across different zones in the data center
- Detecting and reporting suspicious anomalies in the network

8-24 Requirements for Configuring North-South IDS/IPS

Before using North-South IDS/IPS, security administrators must consider the following factors:


- NSX Manager has Internet connectivity to download the signatures from the cloud-based store. Proxy settings can be configured if necessary.
- North-South IDS/IPS is supported on a Tier-1 gateway that is backed by an NSX Edge cluster.
- The NSX environment is configured with a valid license for North-South IDS/IPS.

Additional gateway security capabilities are available with NSX security add-on licenses. 416

8-25 North-South IDS/IPS Architecture

North-South IDS/IPS works as follows:

- NSX Manager downloads IDS/IPS signatures from a cloud-based store.
- NSX Policy stores the IDS/IPS configuration information in the database.
- NSX Policy passes the Information to NSX Manager.
- NSX Manager passes the Information to the CCP.
- CCP pushes the IDS/IPS configuration through APH to NSX Edge using the NSX-proxy.
- NSX Edge stores the signature information locally and configures the datapath.
- The Event Forwarding module collects traffic data and sends events to NSX Manager.
- IDS/IPS reporting on NSX Manager parses and shows these events in the NSX UI.



8-26 Configuring North-South IDS/IPS

To enable north-south intrusion detection and prevention for Tier-1 gateways, you select. Security > IDS/IPS & Malware Prevention > Settings.



8-27 Configuring North-South IDS/IPS Rules

You specify how traffic is managed in your environment by configuring one or more rules with an IDS profile and the mode of operation.

As with gateway firewall rules, gateway specific IDS/IPS rules are evaluated from top to bottom.



You create IDS polices and rules for a specific Tier-1 gateway by navigating to Security > Policy Management > IDS/IPS & Malware Prevention > Gateway Rules and selecting the relevant Tier-1 gateway

You create IDS policies and rules, which are shared between multiple gateways, by navigating to Security > Policy Management > IDS/IPS & Malware Prevention > All Shared Specific Rules.

The Sources, Destinations, Services, and Applied To fields operate in the same way as those in a gateway firewall rule.

Security Profiles specify the group of signatures that the traffic is matched against.

The Mode field in a North-South IDS/IPS rule determines the action that is taken when a signature is triggered. The following modes are available for a North-South IDS/IPS rule:

- Detect Only: Regardless of the global or per-signature action, only alerts are generated, and no preventive action is taken. This mode is equivalent to an intrusion detection system.
- Detect & Prevent: The action specified for the given signature either globally or at the
 profile level is taken (alert, drop, reject). The action specified at the profile level overrides
 the action configured globally. This mode is equivalent to an inline intrusion prevention
 system.

As with gateway firewall rules, the North-South IDS/IPS rules are evaluated from top to bottom. You must place the most hit rules at the top to avoid the unnecessary evaluation of subsequent rules.

8-28 Monitoring North-South IDS/IPS Events

You monitor North-South IDS/IPS events by navigating to **Security** > **Threat Event Monitoring** > **IDS/IPS**.



The same dashboard is used to monitor the north-south IDS/IPS events as in distributed IDS/IPS.

8-29 Lab 13: Configuring Distributed Intrusion Detection and Prevention

Configure Distributed Intrusion Detection and analyze malicious traffic:

- Prepare for the Lab
- Enable Distributed Intrusion Detection and Prevention
- Download the Intrusion Detection and Prevention Signatures
- Create an Intrusion Detection and Prevention Profile
- Configure Intrusion Detection Rules
- Generate Malicious Traffic
- Create a Segment and Attach a VM
- Generate Suspicious Traffic
- Analyze Intrusion Detection Events
- Modify the IDS/IPS Settings to Prevent Malicious Traffic
- Generate and Analyze Intrusion Prevention Events

8-30 Review of Learner Objectives

- Explain IDS/IPS and its use cases
- Define the IDS/IPS terminologies
- Describe the IDS/IPS architecture
- Configure IDS/IPS
- Interpret IDS/IPS events

8-31 Lesson 2: NSX Application Platform

8-32 Learner Objectives

- Describe NSX Application Platform and its use cases
- Deploy NSX Application Platform
- Explain the NSX Application Platform architecture and services

8-33 About NSX Application Platform

NSX Application Platform is a container-based solution that is deployed on an existing Kubernetes cluster.

You must deploy NSX Application Platform before using the following NSX security features:

- NSX Malware Prevention
- NSX Intelligence
- NSX Network Detection and Response
- NSX Metrics

Before deploying NSX Application Platform, you must understand the concepts presented in the course Kubernetes Fundamentals at

https://learning.customerconnect.vmware.com/oltpublish/site/program.do?dispatch=showCourseSession&id=663f8ec8-4078-11eb-8643-0cc47adeb5f8.

8-34 Prerequisites for NSX Application Platform Deployment

To deploy NSX Application Platform, your environment must comply with the following prerequisites:

- Valid, non-expired NSX license
- Tanzu Kubernetes cluster (TKC) v1.20 through v1.22 or upstream Kubernetes v1.20 through v1.24.
- Network connectivity between your microservices environment and the NSX Management cluster
- Working DNS Instance
- Time synchronized between NSX, vSphere, and your Kubernetes environments
- (Optional) Private Harbor registry with a chart repository service configured
 This final item is only required if your Kubernetes environment does not have access to the Internet or if you have specific security restrictions.

For more information about the NSX Application Platform deployment prerequisites, including the supported Tanzu Kubernetes cluster and upstream Kubernetes versions, see *Deploying and Managing the VMware NSX Application Platform* at https://docs.vmware.com/en/VMware-NSX/4.0/nsx-application-platform/GUID-D54C1B87-8EF3-45B3-AB27-EFE90A154DD3.html

The use of a private Harbor instance is required if your Kubernetes environment does not have access to the Internet or you have specific security restrictions.

However, if your Kubernetes environment has external connectivity, you can use registry and repository hosted by VMware to simplify the NSX Application Platform deployment process. This deployment process uses an outbound connection only and does not retain customer data.

8-35 Setting Up a Private Harbor Registry


If your Kubernetes environment does not have access to the Internet, you must set up a private Harbor registry with a chart repository service before you deploy NSX Application Platform.

You then use this registry to upload the Helm charts and Docker images required to deploy NSX Application Platform.

The Helm charts specify the configuration settings to be used for the deployment.

The Docker Images include the container images to be used for the deployment.

For production environments, the private Harbor instance must be configured using external certificate authority (CA)-signed certificates.




VMware Harbor Registry is an enterprise-class registry server that stores and distributes container images.

See the chapter about uploading the NSX Application Platform Docker images and Helm charts in *Deploying and Managing the VMware NSX Application Platform* at https://docs.vmware.com/en/VMware-NSX/4.0/nsx-application-platform/GUID-FAC9DBE3-A8EE-4891-A723-942D0AB679F6.html#GUID-FAC9DBE3-A8EE-4891-A723-942D0AB679F6

8-36 NSX Application Platform Form Factors

Based on the required features, NSX Application Platform can be deployed in different form factors:

- Standard
- Advanced
- Evaluation



The NSX Application Platform form factors have the following features:

- Standard:
 - Supports NSX Network Detection and Response, NSX Malware Prevention, and Metrics
 - Regulres one controller and three worker nodes in the Kubernetes cluster
- Advanced:
 - Supports NSX Network Detection and Response, NSX Malware Prevention, NSX Intelligence, and Metrics
 - Requires one controller and three worker nodes in the Kubernetes cluster
- Evaluation:
 - Supports all available features
 - Requires one controller and one worker node in the Kubernetes cluster
 This form factor is not supported in production environments. It is intended only for evaluation or proof-of-concept.

8-37 NSX Application Platform Deployment (1)


You deploy NSX Application Platform from the NSX UI by selecting **System > Configuration > NSX Application Platform**.

The Helm Repository and Docker Registry URLs must point to the registry and repository hosted by VMware or to your private Harbor Instance.



8-38 NSX Application Platform Deployment(2)

During deployment, you must specify the configuration file for the underlying Kubernetes infrastructure. You also select the form factor based on your feature requirements.



The configuration file for the Kubernetes cluster must be stored locally in the machine where NSX Application Platform deployment is being initiated. This file typically has a YAML format that must be provided by your Kubernetes administrator.

The exact steps required to obtain the Kubernetes cluster configuration file depend on the platform on which your Kubernetes cluster is running. In vSphere with Tanzu environments, you must work with your infrastructure administrator to create a kubeconfig file with a long-lived token to be used in the NSX Application Platform deployment. See the chapter about generating a TKC configuration file with a non-expiring token in *Deploying and Managing the VMware NSX Application Platform* at https://docs.vmware.com/en/VMware-NSX/4.0/nsx-application-platform/GUID-52A52C0B-9575-43B6-ADE2-E8640E22C29F.html

You can then use a file transfer utility to copy the YAML configuration file to your local system.

Apart from specifying the configuration file for the Kubernetes cluster and the desired form factor for the NSX Application Platform, you must also configure the following parameters during NSX Application Platform deployment:

- Interface Service Name (FQDN): The Interface Service Name is used as the HTTPS endpoint to connect to NSX Application Platform.
 - You must configure the FQDN with a static IP address in the DNS server before the NSX Application Platform deployment. The TKC or upstream Kubernetes cluster infrastructure must be able to assign this static IP address.
- Messaging Service Name: The Messaging Service Name value is an FQDN for the HTTPS endpoint that is used to receive the streamlined data from the NSX data sources.

8-39 NSX Application Platform Predeployment Checks

Before proceeding with the deployment of NSX Application Platform, the wizard checks the connection, resources, and correct configuration of the specified Kubernetes cluster.



8-40 NSX Application Platform Deployment Validation

After the deployment, you can validate the status of NSX Application Platform nodes from the NSX UI.



The example illustrates the nodes available after the deployment of NSX Application Platform with an Evaluation form factor.

This type of deployment includes:

- · One controller node: Used as the control plane
- One worker node: Used to perform data processing and analytics tasks

You can review the resource utilization at the overall cluster level, and you can also view the specific resource utilization for each node.


8-41 NSX Application Platform Services

The following core services are deployed as part of NSX Application Platform in both the Standard and Advanced form factors:

- Messaging
- Configuration Database
- Platform Services
- Metrics

The following core services are deployed as part of the Advanced form factor only:

- Analytics
- Data Storage



The following core services are deployed as part of NSX Application Platform in both the Standard and Advanced form factors:

- Messaging: Includes a distributed messaging system that is responsible for collecting network flows and events from the NSX transport nodes and system configuration from NSX Manager.
- Configuration Database: Database system used to store events and recommendations
- Platform Services: Includes all services related to certificate and cluster management.
- Metrics: Collects key performance indicators from the NSX environment.

The following core services are deployed as part of the Advanced form factor only:

- Analytics:
 - Processes and correlates the network flows
 - Generates recommendations and events
- Data Storage: Includes a distributed database used to persistently store correlated flows

During the deployment of NSX Application Platform with a Standard form factor, only the messaging, configuration database, metrics, and platform services are enabled. Installing additional security features on top of NSX Application Platform with a Standard form factor enables additional services. For example, the installation of Malware Prevention and NSX Network Detection and Response automatically enables the analytics service.


8-42 Objects Created During the NSX Application Platform Deployment


During the deployment of NSX Application Platform, several pods are automatically created in the Kubernetes cluster. The number of created pods depends on the form factor.

Pods run in the worker nodes and can run one or more container processes that provide functionality for multiple NSX security features.

One or more pods are deployed per core service of NSX Application Platform.

These pods run in a Kubernetes namespace called nsxi-platform.





8-43 Basic kubectl Commands

Because NSX Application Platform is a container-based solution, you must be familiar with the key kubectl commands to perform basic troubleshooting:

List all the namespaces available in the Kubernetes cluster:

```
kubectl get namespaces
```

List all pods available for a particular namespace:

```
kubectl get pods -n <namespace>
```

Display the detailed state of all containers within a pod:

```
kubectl describe pod <pod-name> -n <namespace>
```

Retrieve the logs for a pod in a given namespace:

```
kubectl logs <pod-name> -n <namespace>
```

Retrieve the logs for a container running in a pod:

```
kubectl logs <pod-name> -c <container-name> -n <namespace>
```

Get relevant events for a given namespace:

```
kubectl get events -n <namespace>
```

The kubectl logs <pod-name> -n <namespace> command works only if one container is running in a given pod. If multiple containers are running in a pod, you run the kubectl logs <pod-name> -c <container-name> -n <namespace> command instead.

You can add the -f flag to follow the log file. The --timestamps flag can be used to retrieve the date and time of the log events. The -p (previous) flag can be used to see the logs from a previously running container.

8-44 Namespaces Available After NSX Application Platform Deployment

The cert-manager, nsxi-platform, and project contour namespaces are automatically created as part of the NSX Application Platform deployment.

kadmin@SA-K8s-CTRL-	01:~\$ ku	bectl get	namespaces
NAME	STATUS	AGE	
cert-manager	Active	3d1h	
default	Active	234d	
kube-node-lease	Active	234d	
kube-public	Active	234d	
kube-system	Active	234d	
local-path-storage	Active	123d	
metallb-system	Active	234d	
nsxi-platform	Active	3d1h	
projectcontour	Active	3d1h	

8-45 Pods Available After NSX Application Platform Deployment

The example shows an extract of the pods created during NSX Application Platform deployment.

NAME	READY	STATUS	RESTARTS	AGE
authserver-cd77bbd4b-47767	1/1	Running	0	34
cluster-api-745655f4db-lxx42	2/2	Running	4	3d1h
common-agent-59ff698dbd-r6t17	1/1	Running	2	3d1h
common-agent-create-kafka-topic-hlpk9	0/1	Completed	o	3d1h
configure-druid-gqqxr	0/1	Completed	1	3d1h
create-kubeapi-networkpolicy-job-gbn6v	0/1	Completed	0	3d1h
druid-broker-7546b89dc4-vdjm6	1/1	Running	1	3dlh
druid-config-historical-0	1/1	Running	0	3d1h
druid-coordinator-d75498dcd-m7hdw	1/1	Running	1	3d1h
druid-historical-0	1/1	Running	0	3d1h
druid-middle-manager-0	1/1	Running	0	3dlh
druid-overlord-899c4469c-z447b	1/1	Running	2	3d1h
fluentd-0	1/1	Running	0	3dlh
kafka-0	1/1	Running	0	34
<cropped output=""></cropped>				

The example shows an extract of the pods created during NSX Application Platform deployment with an Evaluation form factor in the nsxi-platform namespace. The output is cropped. In a healthy environment, the status of all pods should be either Running or Completed. In a Standard form factor deployment, fewer pods are visible because the analytics and data storage pods are not deployed.

The example shows the following functions of the main pods:

- Druid is a time-series distributed database used to store correlated traffic flows.
- Fluentd provides a unified platform to collect logs from different sources.
- Kafka is a distributed messaging system that receives system configuration data from NSX Manager and traffic flows from the transport nodes.

8-46 Lab 14: (Simulation) Deploying NSX Application Platform

Deploy and validate NSX Application Platform:

- Deploy NSX Application Platform
- Validate the NSX Application Platform Deployment from the NSX UI
- Validate the NSX Application Platform Deployment from the Kubernetes Cluster

8-47 Review of Learner Objectives

- Describe NSX Application Platform and its use cases
- Deploy NSX Application Platform
- Explain the NSX Application Platform architecture and services

8-48 Lesson 3: NSX Malware Prevention

8-49 Learner Objectives


- Identify use cases for malware prevention
- Identify the components in the malware prevention architecture
- Describe the malware prevention packet flows for known and unknown files
- Configure malware prevention and validate the configurations

8-50 About Malware Prevention

Malware prevention detects and prevents malicious file transfers by combining signature-based detection of known malware with static and dynamic analyses of malware samples.

NSX Malware Prevention can be configured in the following locations:

- East-west malware prevention detects and prevents malicious files directly on the guest VMs.
- North-south malware prevention detects malicious files on the gateway firewall.



Malware prevention protects your environment against viruses, worms, trojan horses, spyware, and ransomware. Malware prevention uses the following techniques to detect and prevent malicious file transfers:

- Signature-based detection uses databases of known malware patterns. It scans the file and memory of a system for any data that match the pattern of known malicious software.
- Static file analysis extracts the unique characteristics of a file, such as its structure, and uses
 machine learning algorithms to classify and identify malware indicators.
- Dynamic file analysis performs memory analysis and observes how the file interacts with the system, identifying indicators of malicious activity.

Static document analysis occurs when a file is examined without executing it, whereas dynamic analysis examines the actions of the file and can occur only when the file is executing. Both static and dynamic analysis are behavior-based malware detection mechanisms.

Malware prevention can be configured in the following locations:

- East-west malware prevention is configured directly on the guest VMs to protect malware from spreading laterally in the data center.
- North-south malware prevention is configured on the edge node to prevent malware from entering the perimeter.


8-51 About East-West Malware Prevention

East-west malware prevention detects known malicious files on guest VMs and prevents them from being downloaded.

Guest Introspection agents are installed on the guest VMs and perform the following functions:

- Generate alerts if malicious known file hashes are detected
- Extract unknown files and send for local and cloud-based analysis

Alerts are generated based on the return verdict.



East-west malware prevention protects the data center from the spread of internal malware and from malware that makes it past the network perimeter. To perform these tasks, it monitors files downloaded on the guest VMs for malicious content.


To detect and prevent malware from spreading in the environment, malware prevention combines signature-based detection of known malware with static and dynamic analyses of malware samples.

Local analysis, which is a combination of static analysis and machine learning-based analysis of files, is performed on the ESXI host.

Dynamic analysis is performed in a cloud-based environment and is optional.

8-52 Use Cases for East-West Malware Prevention

East-west malware prevention provides malware detection and prevention capabilities for data center and virtual desktop infrastructure (VDI) end users.



East-west malware prevention uses distributed firewall capabilities to protect end users from downloading malicious content.

East-west malware prevention protects guest VMs and virtual desktop infrastructures (VDIs) from lateral malware spreads. Each ESXI host includes a service virtual machine (SVM), which inspects all files that are downloaded at the guest level.

8-53 Requirements for East-West Malware Prevention

East-west malware prevention has the following requirements:

- The NSX environment must be configured with a valid license for malware prevention.
- At a minimum, NSX Application Platform must be deployed with the Standard form factor in the environment.
- The thin agent must be installed on every guest VM.
- NSX Manager must have Internet access.

Proxy settings can be configured if necessary.

For more information about the types of valid licenses for malware prevention, see "VMware NSX" at

https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/nsx/vmwa re-nsx-datasheet.pdf.

Cloud-based analysis is optional.


In Windows operating systems, the thin agent is installed as part of VMware Tools.

In Linux operating systems, the thin agent is a separate package that can be downloaded from the VMware web page. For more information about installing the Guest Introspection thin agent on Linux virtual machines, see NSX Administration Guide at

https://docs.vmware.com/en/VMware-NSX/4.0/admlnistration/GUID-4871C429-CFE6-41C9-86C9-7FCFE9C95EC8.html.

8-54 East-West Malware Prevention Architecture

East-west malware prevention includes ESXI hosts, NSX Application Platform, NSX Threat. Intelligence Cloud, and NSX Advanced Threat Analyzer Cloud.



NSX Malware Prevention uses the following clouds:


- NSX Threat Intelligence Cloud: Obtains file reputations from VMware Carbon Black Cloud and caches them for use in the NSX environment
- NSX Advanced Threat Analyzer Cloud: Performs machine learning-based detection and dynamic analysis of the files

8-55 NSX Application Platform Components

The NSX Malware Prevention feature requires an NSX Application Platform deployment.

In NSX Application Platform, the main components that are deployed for malware prevention are as follows:

- Security Analyzer:
 - Advanced Signature Distribution Service (ASDS) back end
 - Reputation service
 - Messaging pod
 - Configuration database pod
- NSX Cloud Connector



The Standard form factor is enough to use NSX Malware Prevention, but the Advanced form factor is mandatory to install NSX Intelligence.


In NSX Application Platform, the main components that are deployed for malware prevention are as follows:

- Security Analyzer collects all Information that is received from the messaging pod and fetches reputations and signatures from VMware Carbon Black Cloud through the reputation service. Security Analyzer maintains two databases:
 - The Configuration database pod contains all received file events.
 - The NSX Advanced Signature Distribution Service (ASDS) back end gathers the verdict and reputations for all known files.
- NSX Cloud Connector sends files for dynamic analysis to NSX Advanced Threat Analyzer
 Cloud. NSX Cloud Connector acts as a gateway between on-premises services and NSX
 Advanced Threat Analyzer Cloud. Its purpose is to centralize communication and provide an
 authenticated channel between clients and the cloud.

8-56 ESXi Host Components

The main ESXI host components for east-west malware prevention are as follows:

- Guest VM
- Context Multiplexer (MUX)
- Service VM:
 - Guest Introspection agent
 - Security hub
 - Rapid API for Detection (RAPID)
 - ASDS cache



The main ESXI host components for east-west malware prevention are as follows:

 Guest VM: Runs a Network introspection agent, called the thin agent, which offloads files for scanning to the service virtual machine (SVM)

In Windows operating systems, the thin agent is installed as part of VMware Tools and includes the following components:

- NSX Network Introspection (vnetWFP.sys)
- NSX File Introspection (vsepfit.sys)
- VMCl (vsock.svs)

Beginning with NSX 4.0.1, east-west malware prevention is also supported for Linux VMs. The Guest Introspection thin agent for Linux is available as part of the operating system specific packages (OSPs). The packages are hosted on the VMware packages portal.

- Context Multiplexer (MUX): Relays messages between the guest VMs and the SVM, maintains the SVM configuration, and processes the east-west malware prevention policies.
 The MUX is installed as a VIB during transport node preparation.
- SVM: VM appliance that is deployed on every ESXI host part of a malware preventionenabled cluster


The SVM monitors files that are offloaded from the guest VMs, performs local analysis, and connects to NSX Application Platform for file reputation and cloud-based analysis. The SVM contains the following modules:

- Guest Introspection agent: Relays event and data received from the guest VM to the security hub
- Security hub: Collects file events, obtains verdicts for known files, and sends files for local and cloud-based analysis
- Rapid API for Detection (RAPID): Provides local analysis of the files through a combination of static analysis and machine learning-based analysis
- Advanced Signature Distribution Service (ASDS) cache: Maintains a database of known files verdicts and reputations

The ASDS back end is present in the Security Analyzer and each SVM maintains an ASDS local cache.

8-57 East-West Malware Prevention Packet Flow for Known Files

A packet flow occurs when a transfer of a known file is detected on the guest VM.




The following packet flow occurs when a file transfer is detected on the guest VM:

- The thin agent extracts the file, computes the hash, and provides information to the security hub through the MUX and the Guest Introspection agent.
- The security hub checks whether the file is known in the local ASDS cache by sending the hash.
- The local cache sends the verdict back.
- If the file is not in the local cache, ASDS queries the ASDS back end internally for the file reputation.
- The ASDS back end sends the verdict back to the security hub, and the appropriate action is taken.

8-58 East-West Malware Prevention Packet Flow for Unknown Files

If the security hub cannot retrieve the file reputation from the ASDS back end, the file is considered unknown and is sent for local and cloud-based analysis.



If the security hub cannot retrieve the file reputation from the ASDS back end, the file is considered unknown and is sent for local and cloud-based analysis:

- The security hub sends the file to the RAPID module to perform local analysis.
- Based on the malware prevention policy and local analysis results, RAPID sends the file to the NSX Advanced Threat Analyzer Cloud for analysis through the NSX Cloud Connector.


This step occurs only if the policy is set up for cloud-based analysis. If cloud-based analysis is not set up, only the local analysis verdict is sent to the security hub.

- NSX Advanced Threat Analyzer Cloud sends the combined verdicts of the local and cloudbased analysis to the security hub, and the appropriate action is taken.
 - If the verdict of a file is malicious, and if the file's type is portable executable, the security hub sends the file's hash to the reputation service to cross-check its reputation. This step is performed to reduce false positives. The file reputation service queries VMware Carbon Black Cloud to retrieve the file reputation.
- The security hub collects verdicts and statistics and sends an event to the security analyzer.
- The security analyzer reports the verdict and statistics to NSX Manager.

The security analyzer polls the security hub for the local and cloud-based analysis verdicts and updates the ASDS back end accordingly. This data is used for future downloads of the same file.

8-59 Activating Malware Prevention on NSX Application Platform


To enable NSX Malware Prevention in the NSX UI, you select **System > Configuration > NSX Application Platform** and click the NSX Malware Prevention tile.



8-60 Setting the Cloud Region

You select the cloud region, run prechecks, and then click ACTIVATE.

The NSX Malware Prevention installation deploys both NSX Cloud Connector and the components required for malware prevention.



You specify the NSX Advanced Threat Analyzer Cloud instance that you want your environment to connect to:

The FQDN for the United States cloud is nsx.west.us.lastline.com, and the FQDN for the European cloud is nsx.nl.emea.lastline.com. NSX Malware Prevention uses HTTPS port 443 to access NSX Advanced Threat Analyzer Cloud. Before proceeding with the installation of NSX Malware Prevention, the installation wizard verifies that the license is valid and that NSX Advanced Threat Analyzer Cloud is accessible.

NSX Cloud Connector is a shared component between NSX Network Detection and Response and NSX Malware Prevention. The NSX Cloud Connector deployment is skipped if NSX Network Detection and Response is already configured in the environment.

In environments in which both these features are installed, changing the cloud region requires reinstalling both NSX Network Detection and Response and NSX Malware Prevention. Modifying the cloud region after the installation is not supported.

8-61 Validating the Malware Prevention Installation

You can validate the successful installation of NSX Malware Prevention from the NSX UI.



8-62 Service Registration

Before you can use malware prevention on the transport nodes, you must register the malware prevention service and deploy SVMs on each host.

You register the malware prevention service with NSX Manager with the following API call:

```
POST https://sa-nsxmgr-01.vclass.local/napp/api/vl/malware-
prevention/svm-spec

Body:
{
   "ovf_url": "<OVF_PATH>",
   "deployment_spec_name": "MPS-SVM"
}
```

You must use a REST API call to register malware prevention.


8-63 Service Registration Validation from the NSX UI

You validate the malware prevention service registration in the NSX UI by selecting **System** > **Configuration** > **Service Deployments** > **Catalog**.



8-64 Service Deployments

To deploy the SVMs on each host, you select **System > Configuration > Service Deployments** > **Deployment.**



The deployment of the SVMs is cluster-based:

- All hosts inside the cluster are deployed with an instance of the service.
- If a new host is added to the cluster, it is automatically deployed with a service instance.

After clicking **SAVE**, the deployment of the SVMs on each ESXI host starts. You can monitor the deployment by looking at the tasks in vCenter Server.

8-65 Service Deployment Validation from the NSX UI

You can validate the service instance deployment from the NSX UI by selecting **System > Configuration > Service Deployments > Service Instances**.



In a healthy environment, both the Deployment Status and the Health Status appear as Up.


You can check alarms for the NSX Malware Prevention feature by selecting **Home > Alarms > Alarms**.

The following alarms are available for NSX Malware Prevention:

- Analyst API Services Unreachable
- Database Unreachable
- File Extraction Service Unreachable
- NTICS Reputation Service Unreachable
- Service Status Down

8-66 Service Deployment Validation from vCenter Server

You can validate the service instance deployment from vCenter Server.



The SVM is deployed in vCenter Server and runs with the following IP addresses:

- The Management network is defined by the administrator during the configuration.
- The Control network is defined by the system and gets the IP 169.254.1.22.
- 172.18.0.1 is an internal IP address used to connect the RAPID container. Therefore, it is not
 connected to an external network like the other two IP addresses.

8-67 Creating East-West Malware Prevention Profiles

The NSX Malware Prevention profile defines the types of files to be analyzed.




You create a profile for malware prevention by selecting Security > Policy Management > IDS/IPS & Malware Prevention > Profiles > Malware Prevention.

In earlier versions of NSX-T Data Center, only executable files were supported for east-west malware prevention.

Since NSX 4.0.1, all file types are available when configuring east-west malware prevention. The file types are specified in the Malware Prevention Profile.

8-68 Creating Rules for East-West Malware Prevention

To create a rule, you select Security > Policy Management > IDS/IPS & Malware Prevention > Distributed FW Rules.



A malware prevention policy is a collection of malware prevention rules.

A malware prevention rule contains a set of instructions that determine which file is analyzed, including the source and destination, the services, the malware prevention profile, where to apply the rule, and the detection mode.

Malware prevention rules must be applied to the group of VMs that you want to protect. The rules do not work if they are applied to the distributed firewall.

NSX 4.0.1 Includes the following modes for a malware prevention rule:

- Detect Only: Detects malware and does not act
- Detect and Prevent: Detects and blocks malware


Only one malware prevention profile can be attached to a rule. However, a rule can have both a malware prevention profile and an IDS/IPS profile.

8-69 About North-South Malware Prevention

North-south malware prevention detects known malicious files when they enter the perimeter on the NSX gateway firewall.

If the file is unknown, NSX Edge extracts the file and sends it for local and cloud-based analysis.

The verdict is computed, and alerts are generated.



In NSX 4.0.1, only north-south malware detection of the malicious files is supported.


The detection system compares the file hashes (SHA1/MD5/SHA256) to file hashes of known malware.

Local analysis, which is a combination of static analysis and machine learning-based analysis of the files, is performed in the NSX Edge node.

Dynamic analysis is performed in a cloud-based environment and is optional.

8-70 Use Cases for North-South Malware Prevention

North-south malware prevention provides malware detection capabilities at the perimeter of the data center.



North-south malware prevention monitors and generates alerts when users want to download malicious files from an external network or from public clouds.

8-71 Requirements for North-South Malware Prevention

North-south malware prevention has the following requirements:

- The NSX environment must be configured with a valid license for malware prevention.
- At a minimum, NSX Application Platform with the Standard form factor must be deployed in the environment.
- North-South malware prevention is only supported on Tier-1 gateways.
- The form factor of the NSX Edge nodes must be extra large.
- NSX Manager must have Internet access.

Proxy settings can be configured if necessary.

For more information about the types of valid licenses for malware prevention, see "VMware NSX" at

https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/nsx/vmware-nsx-datasheet.pdf.


The sizing requirement for extra large NSX Edge nodes are 16 vCPUs, 64 GB of RAM, and 200 GB of storage.

Beginning with NSX 4.0.1, NSX Malware Prevention can be also configured on a bare-metal edge node.

Cloud-based analysis is optional.

8-72 North-South Malware Prevention Architecture

North-south malware prevention architecture includes the NSX Edge node, NSX Application Platform, NSX Threat Intelligence Cloud, and NSX Advanced Threat Analyzer Cloud.




The north-south malware prevention architecture shares the same components as east-west malware prevention, with the difference that the NSX Edge node is used in place of the ESXI hosts.

8-73 NSX Edge Components

The main components on the edge node for north-south malware prevention are as follows:

- IDS/IPS engine
- Security hub
- RAPID
- ASDS Cache




The main components on the edge node for north-south malware prevention perform the following functions:

- IDS/IPS engine: Extracts files and relays events and data to the security hub
 North-south malware prevention uses the file extraction features of the IDS/IPS engine that runs on NSX Edge for north-south traffic.
- Security hub: Collects file events, obtains verdicts for known files, sends files for local and cloud-based analysis, and sends information to the security analyzer
- RAPID: Provides local analysis of the file
- ASDS Cache: Caches reputation and verdicts of known files

8-74 North-South Malware Prevention Packet Flow for Known Files

The packet flow for north-south malware prevention is similar to east-west malware prevention.




The packet flow process includes the following steps:

- A file transfer of a known file is detected on the NSX Edge node.
- The IDS/IPS engine running on the NSX Edge node extracts the file, computes the hash, and provides information to the security hub.
- The security hub uses the hash to verify whether the file is known to the local ASDS cache.
- The local ASDS cache sends the verdict back.
- If the file is not in the local cache, ASDS queries the ASDS back end internally for the file reputation.
- The ASDS back end sends the verdict back to the security hub, and appropriate action is taken.

8-75 North-South Malware Prevention Packet Flow for Unknown Files

If the security hub cannot retrieve the file reputation from the ASDS back end, the file is sent for local and cloud-based analysis.



If the security hub cannot retrieve the file reputation from the ASDS back end, the file is sent for local and cloud-based analysis:

- The security hub sends the file to the RAPID module to perform local analysis.
- Based on the NSX Policy and local analysis results, RAPID sends the file to NSX Advanced Threat Analyzer Cloud for analysis through NSX Cloud Connector.

This step takes place only if the policy is set up for cloud-based analysis. If cloud-based analysis is not set up, only the local analysis verdict is used and sent back to the security hub.

- NSX Advanced Threat Analyzer Cloud sends the combined verdicts of the local and cloudbased analysis to the security hub, and the appropriate action is taken.
 - If the verdict of a file is malicious and the file's type is Portable Executable, the security hub sends the file's hash to the reputation service to cross-check its reputation. This step is performed to reduce the number of false positives. The file reputation service queries VMware Carbon Black Cloud to retrieve the file reputation.
- The security hub collects verdict information and statistics and sends an event to the security analyzer.
- The security analyzer reports the verdict and statistics to NSX Manager.

The security analyzer polls the security hub for the local and cloud-based analysis verdicts and updates the ASDS back end accordingly. These data are used for future download of the same file.

8-76 Enabling Malware Prevention on Tier-1 Gateways

To enable malware prevention on the selected Tier-1 gateways, you select **Security > Policy**Management > IDS/IPS & Malware Prevention > Settings > Shared and turn on the Malware

Prevention toggle.




8-77 Creating North-South Malware Prevention Profiles

You create a malware prevention profile by selecting Security > Policy Management > IDS/IPS & Malware Prevention > Profiles > Malware Prevention.



8-78 Creating Rules for North-South Malware Prevention

Gateway firewall rules specify the parameters on which north-south malware prevention is applied. The rules are enforced on the selected Tier-1 gateway.




Only one malware prevention profile can be attached to a rule. But a rule can have both a malware prevention profile and an IDS/IPS profile.

You create a rule by selecting Security > Policy Management > IDS/IPS & Malware Prevention > Gateway FW Rules.

8-79 Malware Prevention Dashboard (1)

On the **Potential Malware** tab, you can find the list of all detected files that are potentially harmful to the system for both east-west and north-south traffic.



You can access the UI dashboard by selecting Security > Threat Event Monitoring > Malware Prevention.

On the **Potential Malware** tab, only files with a potentially harmful verdict appear. Files with a benign verdict are not shown on this tab.

The NSX UI shows the number of files detected, the time of the detection, the verdict (malicious, suspicious, or uninspected), and the malware family and class.

Files are categorized based on their score range:

- 70 through 100 (red): Malicious
- 30 through 69 (orange): Suspicious
- 1 through 30 (green): Benign

Benign files appear on the All Files tab, not the Potential Malware Files tab.

-1 (gray): Uninspected

Uninspected files are not analyzed because they appear in the allow list. Files in the allow list are not blocked even if they are classified as suspicious or malicious.

The following filters can be used to search through the inspected files:

- Blocked: Specifies whether the file has been blocked or not by a malware prevention rule
- File Hash: Searches using the hash value, which is a unique value corresponding to the content of the file.
- File Type: Filters based on the type of the file, for example, .doc, .pdf, .exe, and so on
- Malware Class: Defines the type of threat, for example, virus, trojan horse, worm, adware, ransomware, spyware, and so on
- Malware Family: Identifies a specific group of malware files that typically originate from the same source code or are developed by the same authors, for example, valyria, darkside, and so on
- Verdict: Filters based on the decision taken about the files.


8-80 Malware Prevention Dashboard (2)

On the All Files tab, you can find the list of all inspected files.



8-81 About the Allowlist

Files in the allowlist are not blocked even if they are classified as suspicious or malicious.



You can add files to the allowlist after they are detected by the system. You select **Security** > **Threat Event Monitoring** > **Malware Prevention**.

You can list all files that are present in the allowlist by selecting Security > Policy Management > IDS/IPS & Malware Prevention > Settings > Malware Prevention.

Files in the allow list appear as uninspected files (grey color) in the Malware Prevention dashboard.

8-82 Lab 15: (Simulation) Configuring Malware Prevention for Fast-West Traffic

Configure malware prevention for east-west traffic:

- Install Malware Prevention
- Validate the Malware Prevention Deployment from the CLI
- Register the Malware Prevention Service
- Deploy the Service Instances
- Validate the Service Instances Deployments
- Create a Malware Prevention Profile
- Configure the East-West Malware Prevention Rules
- Download Files from the Guest VM
- Review the Malware Prevention Dashboard

8-83 Review of Learner Objectives

- Identify use cases for malware prevention
- Identify the components in the malware prevention architecture
- Describe the malware prevention packet flows for known and unknown files
- Configure malware prevention and validate the configurations

8-84 Lesson 4: NSX Intelligence


8-85 Learner Objectives

- Describe NSX Intelligence and its use cases
- Explain NSX Intelligence system requirements
- Activate NSX Intelligence
- Explain NSX Intelligence visualization, recommendation, and network traffic analysis capabilities

8-86 About NSX Intelligence

NSX Intelligence is a distributed analytics solution that provides visibility and dynamic security policy enforcement for NSX environments. Features of NSX Intelligence include:

- Visualization of the data center objects and traffic flows
- Recommendations for micro-segmentation planning
- Suspicious Traffic Detection to identify abnormal traffic patterns



8-87 Use Cases for NSX Intelligence

NSX Intelligence enables several capabilities for security administrators.







NSX Intelligence enables the following capabilities for security administrators:

- Visibility: Provide insight about the micro-segmentation traffic flows
- Policy recommendations: Generate security recommendations and dynamically enforce security policies
- Suspicious Traffic Detection: Use Al techniques to detect malicious network behaviors

8-88 NSX Intelligence Requirements

The requirements for using NSX Intelligence are as follows:

- NSX Application Platform with an Advanced form factor must be deployed in the environment.
- You must have a valid NSX license to avail the NSX Intelligence features.
- You must have an Enterprise Administrator role to start and stop data collection.


Before deploying NSX Intelligence, ensure that you have the correct NSX license for the NSX Intelligence features that you are planning to use in your environment. For more information about the product offerings for NSX 4.0.x security, see VMware knowledge base article [89137] at https://kb.vmware.com/s/article/89137.

The Enterprise Administrator role starts and stops data collection. Other user roles, such as Security Administrators, can visualize the NSX Intelligence data and create and apply recommendations. However, an Enterprise Administrator role is mandatory to start and stop data collection.

8-89 NSX Intelligence Activation

Before you can start using the visualization, recommendation, and suspicious traffic detection capabilities of NSX Intelligence, you must activate the NSX Intelligence feature from the NSX UI or API.

The activation process runs some prechecks to ensure that the feature can successfully run in the environment.



You can activate NSX Intelligence from the NSX UI by selecting System > Configuration > NSX Application Platform. The NSX Intelligence tile can be found under the features section of NSX Application Platform.

NSX Intelligence is unavailable if NSX Application Platform is not deployed with an Advanced form factor in the environment or if the required licenses are not available.

For more information about upgrading from an earlier version of NSX Intelligence to NSX Intelligence 4.0.1 or later, see Activating and Upgrading VMware NSX Intelligence at https://docs.vmware.com/en/VMware-NSX-Intelligence/4.0/install-upgrade/GUID-9F91CFBC-DE26-451C-90E0-5AC07117BFFD.html

8-90 Validating the NSX Intelligence Deployment

You can use common kubectl commands to better understand the state of the NSX Intelligence deployment:

Validate the successful deployment of the NSX Intelligence pods:

```
kubectl get pods -n nsxi-platform
```

Display the detailed state of all containers within a pod:

```
kubectl describe pod <pod-name> -n nsxi-platform
```

Review the logs for a given pod:

```
kubectl logs <pod-name> -n nsxi-platform
```

The activation of NSX Intelligence on NSX Application Platform deploys additional pods on the Kubernetes cluster. You can use common kubectl commands to better understand the state of the NSX Intelligence deployment and to review the logs for a particular pod, if required.

8-91 Granular Data Collection

NSX Intelligence provides the ability to select the standalone hosts or clusters for which you want to enable data collection.




On new deployments of NSX Intelligence, data collection is enabled by default on all hosts and clusters. As an NSX Administrator, you can choose to collect data for particular standalone hosts or clusters. This method helps reduce the amount of data to capture and process, which improves the amount of resources available in the environment.

You can adjust the NSX Intelligence data collection settings from the NSX UI by selecting System > Settings > NSX Intelligence.

8-92 NSX Intelligence Visualization (1)

You can visualize flows for VMs and security groups by selecting Plan & Troubleshoot > Discover & Plan > Discover & Take Action.



Users can filter data flows based on security groups, VMs, physical servers, and IP addresses.

NSX Intelligence offers a time range for the visualization of traffic flows called Now, which you can use to visualize the most recent flows in your environment.

Traffic flows display public IP addresses used as sources and destinations.

In addition, NSX Intelligence offers enhanced filtering capabilities to more granularly define flows that are displayed in the canvas.

Data flows are divided into the following categories:

- Unprotected: The traffic flow matches the default firewall rule to allow, drop, or reject any traffic from any source and any destination. More granular security policies are needed to secure the environment.
- Blocked: The traffic flow matches a more granular rule than the default rule that drops or rejects traffic.
- Allowed: The traffic flow matches a more granular rule than the default rule that allows traffic.

8-93 NSX Intelligence Visualization (2)

You can examine the details of traffic flows by clicking the corresponding line in the canvas.



The screenshot shows details about an allowed SSH flow between the rhelvm1 and ubuntuvm6 VMs.

Each flow includes the following information:

- Source: Name of the source VM, the source IP address, the user, and the process run
- Destination: Name of the destination VM and the destination IP address.
- Services: Service or services identified in the flow
- App ID: Application ID, if applicable
- FQDN: Fully qualified domain name, if applicable
- Latest Flow: Unprotected, allowed, or blocked
- End Time: Time when the flow ended.

8-94 NSX Intelligence Recommendations (1)

You can start security recommendations by selecting Plan & Troubleshoot > Discover & Plan > Recommendations.

You can initiate multiple recommendations at one time. However, they are processed sequentially.



Many application owners do not have access to the NSX UI. So, you must be able to export the recommendation in a readable format.

In previous versions of NSX, only the JavaScript Object Notation (JSON) format was supported for the export.

You can now export the recommendation in the comma-separated values (CSV) format.


The CSV format stores values in a text file using a comma to separate every value. This format is typically used to store tabular data.

8-95 NSX Intelligence Recommendations (2)

Recommendations analyze traffic data for a given security group or set of VMs for a specified period. Recommendations suggest security groups, services, and distributed firewall rules.

You configure the following parameters to start a recommendation:

- Selected Entitles In Scope
- Time Range
- Connectivity Strategy
- Recommendation Output
- Group Reuse Threshold
- Recommendation Service Type
- Exclusion



You configure the following parameters to start a recommendation:

 Selected Entities in Scope: VMs, physical servers, and security groups can be used as inputs for the recommendation. Security groups can include virtual machines, segment ports, segments, and VIFs.

NSX Intelligence allows the selection of multiple groups as the scope for a new recommendation. These groups must contain no more than 250 effective compute entities in total.

To enhance the fidelity of the recommendations in brownfield deployments, NSX Intelligence also considers the existing distributed firewall policies that are applied to the groups selected as the scope for a new recommendation.

- Time Range: This parameter is the period for which data is analyzed. It ranges from the last 1 month to the last 1 hour.
- Connectivity Strategy:
 - Create Rules For:
 - All Traffic: This default option considers all outbound, inbound, and intra-application traffic flow types.
 - Incoming and Outgoing Traffic: This option considers all traffic flow types that originate both inside and and outside the application boundary.
 - Incoming Traffic: This option only considers traffic flow that originates outside the application boundary.
 - Default Rule:
 - None: This default option does not create any default rule for the security policy.
 - Allowlist: This option creates a default drop rule.
 - Denylist: This option creates a default allow rule.
- Recommendation Output:
 - Compute-based: This default option recommends security groups, including VMs.
 - IP-based: This option recommends security groups, including a static list of IPs.
- Group Reuse Threshold: Users can customize the group reuse threshold to determine
 whether existing groups should preferably be reused or new ones created instead. A low
 threshold of around 10 percent represents an aggressive group reuse, whereas a high
 threshold of 100 percent indicates minimal group reuse. The default group reuse threshold is
 set to 80 percent.

Beginning with NSX 4.0.1, NSX Intelligence can reuse some of the ipset groups already present in the NSX inventory when making recommendations to avoid having too many groups.

Recommendation Service Type:

- L4 Services: This default option generates L4 services and rules as an output for the recommendation.
- L7 Context Profiles: L7 rules, including context-profile information, are recommended for flows with L7 application ID information. If application information is not available, L4 recommendations are generated.

Exclusion:

- Exclude Flows: This setting specifies the traffic flow types to exclude during the
 recommendation analysis. Default values are Broadcast flows and Multicast flows.
 These flow types are not relevant for application category rules. Excluding broadcast
 flows, multicast flows, or both flow types can help optimize the DFW rule
 recommendation analysis.
- Exclude Infrastructure Workload: New to NSX 4.0.1, this feature classifies the VMs in the environment. An algorithm determines which VMs are part of the network infrastructure. The administrator can then decide to exclude those infrastructure VMs from the workload recommendation, which reduces noise from the infrastructure during the recommendation and the visualization of the workloads.

8-96 NSX Intelligence Recommendations (3)

The Monitoring column indicates whether VM changes are detected after the initial recommendation.

If VMs change their group membership after the initial analysis session, a rerun flag is set and users are prompted to rerun the recommendation to analyze changes.

You can turn the Monitoring toggle on or off.



8-97 NSX Intelligence Recommendations (4)

The recommended distributed firewall rules, security groups, and services can be published in NSX Manager.



After the recommendation session is completed, recommendations about the distributed firewall rules, security groups, and services that should be created to secure the environment are provided. You can publish these recommendations in NSX Manager and the distributed firewall rules, security groups, and services are automatically configured for you. You can customize the recommendations before final publication. This customization can include changing the names of the recommended rules and security groups.

Beginning with NSX 4.0.1, you can review and change the group name by directly clicking the group on the **Rules** tab.

8-98 NSX Suspicious Traffic Detection

The NSX Suspicious Traffic Detection feature analyses the data collected by NSX Intelligence and flags suspicious activities using the supported detectors:

- Data Upload/Download
- Destination IP Profiler
- DNS Tunneling
- Domain Generation Algorithm (DGA)
- Horizontal Port Scan
- LLMNR/NBT-NS Poisoning and Relay
- Netflow Beaconing
- Network Traffic Drop
- Port Profiler
- Server Port Profiler
- Remote Services
- Uncommonly Used Port
- Unusual Network Traffic Pattern
- Vertical Port Scan

NSX Intelligence supports the detection of the following threats:

- Data Upload/Download: Detects and alerts if an unusually large amount of data is uploaded or downloaded from a VM
- Destination IP Profiler: Detects and alerts if a VM connects to an IP address that is not part
 of its typical communication pattern
- DNS Tunneling: Detects and alerts about an unusual volume of differing DNS requests towards the same root DNS name

This action might suggest an attempt to exfiltrate data over DNS.


 Domain Generator Algorithm (DGA): Detects and alerts about suspicious DNS traffic from a VM, indicating potential activity from DGA malware

DGAs are used by cybercriminals to prevent their servers from being blocked or taken down. The algorithm produces new domains on demand that a malware sample can use as its Command & Control server.

- Horizontal Port Scan: Detects and alerts if an intruder tries to attack a single port or service across multiple VMs
 - Horizontal Port Scan is also known as sweeping.
- LLMNR/NBT-NS Poisoning Relay: Detects and alerts if a VM sends an unusual number of responses to LLMNR/NBT-NS requests
- Netflow Beaconing: Detects and alerts if a VM indicates beaconing behavior, such as contacting one or more hosts at regular intervals and transferring a similar number of bytes
- Network Traffic Drop: Detects and alerts if an unusually high amount of traffic is dropped by a distributed firewall rule
- Port Profiler: Detects and alerts about suspicious ports accessed from a source VM
- Server Port Profiler: Detects and alerts about suspicious ports accessed on a target machine
- Remote Services: Detects and alerts if suspicious behavior is observed for remote connections such as telnet, SSH, VNC, and remote RDP/RDS sessions.
- Uncommonly Used Port: Detects and alerts if a nonstandard port is used for a given a protocol
 - For example, SSH traffic runs on a port other than the standard port 22.
- Unusual Network Traffic Pattern: Detects and alerts about deviations from predicted network traffic patterns for a given VM
- Vertical Port Scan: Detects and alerts if an intruder tries to attack multiple open ports or services of a target VM

8-99 Configuring Detector Definitions

You can enable the detectors you are interested in on the **Detector Definitions** tab in the NSX UI. All detectors are disabled by default.



Suspicious Traffic Detection is only supported for hosts and clusters that have data collection enabled.


When data collection is enabled, you can enable the detectors you are interested in by selecting Security > Threat Event Monitoring > Suspicious Traffic > Detector Definitions.

Click the pencil icon next to the name of the detector to enter edit mode and turn on the toggle.

All detectors are disabled by default. You must explicitly turn on each detector that you want to use in your NSX environment.

8-100 Visualizing Detected Threats (1)

Detected Threats are displayed on the **Events** tab. Threats are classified according to the MITRE ATT&CK Framework.



You visualize the threat detection events by selecting Security > Threat Event Monitoring > Suspicious Traffic > Events.

The **Events** tab displays all threat detection events identified in the system, classified according to the MITRE ATT&CK framework:

- Persistence events appear in dark orange.
- Credential Access events appear in light blue.
- Discovery events appear in green.
- Command and Control events appear in light purple.
- Lateral Movement events appear in light orange.
- Collection events appear in dark gray.
- Exfiltration events appear in red.

The event classification is also specified in the Type section of the event description.

Events that cannot be clearly mapped to an existing MITRE ATT&CK framework tactic or technique are categorized under the Other category.

Threat detection events are graphically represented by a histogram. Security administrators can specify the period that they are interested in by adjusting the blue vertical lines.

Each threat event type is represented by a dot in the histogram. The size of the dot is proportional to the number of occurrences of an event.

Additional information about each type of event appears in a tabular format.

8-101 Visualizing Detected Threats (2)

Each detected threat can be expanded to retrieve additional details, including its impact score, severity, detector type, affected objects, detected anomalies, and so on.



The impact score for a given event is calculated by combining its severity and the confidence of the detector technique.

The detector type is also displayed in the event details, along with a brief description.

The affected objects, such as target or source VMs, are also displayed here.

Finally, depending on the detector type, the deviation between the normal pattern of behavior and the anomaly is also included.

8-102 Review of Learner Objectives

- Describe NSX Intelligence and its use cases
- Explain NSX Intelligence system requirements
- Activate NSX Intelligence
- Explain NSX Intelligence visualization, recommendation, and network traffic analysis capabilities

8-103 Lesson 5: NSX Network Detection and Response

8-104 Learner Objectives


- Describe NSX Network Detection and Response and its use cases
- Explain the architecture of NSX Network Detection and Response
- Activate NSX Network Detection and Response
- Describe the visualization capabilities of NSX Network Detection and Response

8-105 About NSX Network Detection and Response

NSX Network Detection and Response is an advanced threat prevention platform that provides complete network visibility, detection, and prevention of sophisticated threats:

- Collects network traffic across on-premises networks, cloud, and hybrid cloud infrastructures
- Uses artificial intelligence techniques to analyze network traffic and gain insights into advanced threats
- Helps security teams to visualize the entire attack and trigger the appropriate response







8-106 NSX Network Detection and Response Use Cases

Security teams use NSX Network Detection and Response to perform several functions.





Visualize the Entire Attack



Prevent Intrusions Easter



Reduce False Positives


Security teams can use NSX Network Detection and Response for the following purposes:

- Detect all threat movements: NSX Network Detection and Response can detect threats entering the network perimeter (north-south) as well as attacks that move laterally in the network perimeter (east-west).
- Visualize the entire attack: With NSX Network Detection and Response, you can visualize a
 complete campaign blueprint and a detailed threat timeline across the network so that you
 can quickly understand the scope of an attack and prioritize resources. Additionally, NSX
 Network Detection and Response maps to the MITRE ATT&CK tactics and techniques for
 greater understanding of the key events in a campaign.
- Prevent intrusions faster: NSX Network Detection and Response uses real-time, scalable Al
 and machine learning to detect and stop threats at wire speed.
- Reduce false positives: NSX Network Detection and Response delivers the industry's highest-fidelity insights into advanced threats and reduces false positives by up to 90 percent. NSX Network Detection and Response learns in real time to update detection fidelity.

8-107 NSX Network Detection and Response High-Level Architecture

The high-level architecture of NSX Network Detection and Response contains the following tiers:

- Scalable Analytics tier
- Management tier
- Sensor tler



The high-level architecture of NSX Network Detection and Response contains distinct tiers with the following functions:


- The Scalable Analytics tier is like the brain of NSX Network Detection and Response. It is a
 distributed analytics platform containing multiple nodes that perform deep content
 inspection, network traffic analysis, network and asset profiling, and anomaly detection.
- The Management tier provides the REST API and a web-based UI for all user configurations.
 It also displays alerts and intrusion events.
- The Sensor tier includes one or more sensors, and it is responsible for collecting network traffic across the environment.

8-108 NSX Network Detection and Response in NSX Deployments

The capabilities of NSX Network Detection and Response are closely integrated with onpremises NSX deployments:

- NSX Network Detection and Response collects security events and system configuration data from the NSX Edge nodes, NSX Manager, and ESXI hosts (Sensor tier).
- The collected data is analyzed and correlated using a cloud-based distributed analytics platform called NSX Advanced Threat Analyzer Cloud (Scalable Analytics tier).
- NSX Cloud Connector displays alerts and Intrusion events through a web-based UI (Management tier).


NSX Network Detection and Response is also used for dynamic file analysis or sandboxing when malware prevention is enabled in the NSX environment.



8-109 NSX Network Detection and Response Architecture (1)

NSX Network Detection and Response collects the following data from the NSX platform:

- IDS/IPS events from the ESXI hosts
- Suspicious traffic events from NSX Intelligence
- Files and malware prevention events from the NSX Edge nodes and Security Analyzer
- Site configuration data from NSX Manager




North-south IDS/IPS events and east-west malware prevention events are not yet collected.

8-110 NSX Network Detection and Response Architecture (2)

Data collected from the NSX environment is aggregated and analyzed as follows:

- Security Analyzer receives malware prevention events from the NSX Edge nodes and forwards them to NSX Cloud Connector.
- NSX Cloud Connector gathers IDS/IPS events, malware prevention events and files, and suspicious traffic events from NSX and forwards them to NSX Advanced Threat Analyzer Cloud
- NSX Advanced Threat Analyzer Cloud analyzes and correlates the IDS/IPS, malware prevention, and suspicious traffic events and provides insights about ongoing campaigns.
- Campaign Information appears in the NSX Network Detection and Response UI.



Data collected from the NSX environment is aggregated and analyzed as follows:

- Security Analyzer receives file events from the malware prevention modules in the NSX Edge nodes and persists them locally. These events are then forwarded to NSX Cloud Connector.
- 2. NSX Cloud Connector serves as a gateway between the on-premises deployment of NSX and the NSX Advanced Threat Analyzer Cloud distributed analytics platform. Its purpose is to centralize communication with NSX Advanced Threat Analyzer Cloud services and provide an authenticated channel between clients and the cloud. It gathers IDS/IPS events from the ESXI hosts, malware prevention events and files from the NSX Edge node, and suspicious traffic events from NSX intelligence. The Cloud Connector then forwards all the information to the NSX Advanced Threat Analyzer Cloud for further analysis.
- NSX Advanced Threat Analyzer Cloud analyzes and correlates the IDS/IPS, malware
 prevention, and suspicious traffic events and provides insights about ongoing campaigns. A
 campaign is a set of related events that use specific techniques and can be mapped to the
 MITRE ATT&CK Framework stages to define an attack story.
- Campaign information is displayed in the NSX Network Detection and Response UI, which
 resides in NSX Cloud Connector. The NSX Network Detection and Response UI is separate
 from the NSX UI, but it can be cross-launched from NSX Manager.

8-111 NSX Network Detection and Response Requirements

Before sending events and data to NSX Advanced Threat Analyzer Cloud for campaign correlation, you must activate the NSX Network Detection and Response feature from the NSX UI.

The requirements for activating NSX Network Detection and Response are as follows:

- As a minimum, NSX Application Platform must be deployed with a Standard form factor in the environment. The Advanced and Evaluation form factors are also supported.
- NSX Detection and Response requires an NSX Advanced Threat Prevention license.
- NSX Advanced Threat Analyzer Cloud must be reachable from NSX Application Platform.

By Installing NSX Network Detection and Response, customers accept sending the following data to the public cloud:

- Files for sandboxing or event correlation
- IDS/IPS and malware prevention events for correlation
- Suspicious traffic events for correlation, if NSX Intelligence is available in the environment.


NSX Advanced Threat Analyzer Cloud must be reachable from NSX Application Platform. Deployments without Internet connectivity are not supported.

8-112 NSX Network Detection and Response Activation (1)

You activate NSX Network Detection and Response by selecting System > Configuration > NSX Application Platform > Features.

The NSX Network Detection and Response activation deploys both NSX Cloud Connector and the components required by NSX Network Detection and Response.

As part of the activation, you specify the NSX Advanced Threat Analyzer Cloud instance that you want your environment to connect to.



NSX Cloud Connector is a shared component between NSX Network Detection and Response and NSX Malware Prevention. The NSX Cloud Connector deployment is skipped if the NSX Malware Prevention feature is already configured in the environment.

In environments in which both these features are activated, changing the cloud region requires reinstalling both NSX Network Detection and Response and NSX Malware Prevention. Modifying the cloud region after installation is not supported.

To fully use the capabilities of NSX Network Detection and Response, you enable the following NSX features in the environment:

- NSX Intelligence
- NSX Malware Prevention
- Suspicious Network Activity: Refers to the Suspicious Traffic Detection capabilities provided by NSX Intelligence
- NSX Distributed IDS/IPS

Before proceeding with the activation of NSX Network Detection and Response, the activation wizard verifies that the applied NSX Advanced Threat Prevention license is valid and that NSX Advanced Threat Analyzer Cloud is accessible.

The FQDN for the United States cloud is nsx.west.us.lastline.com and the FQDN for the European cloud is nsx.nl.emea.lastline.com. NDR uses HTTPS port 443 to access NSX Advanced Threat Analyzer Cloud.

Additional cloud regions might become available over time.

8-113 NSX Network Detection and Response Activation (2)

You can validate the successful activation of NSX Network Detection and Response from the NSX UI.



8-114 Validating the NSX Network Detection and Response and NSX Cloud Connector Deployments

To verify that the NSX Cloud Connector and NSX Network Detection and Response components are running, you run the following command in your Kubernetes cluster:

kubectl get pods -n nsxi-platform | egrep "(cloudconnector|nsx-ndr)"

```
kadmin@SA-K8s-CTRL-01:~$ kubect1 get pods -n nswi-platform | egrep "(cloud-connector(nsw-ndr)"
                                                    2/2 Running D
                                                                           14h
cloud-connector-check-license-status-596bff9666-4rgof
                                                                    D
cloud-connector-file-server-Sbb8957bf4-t9smv
                                                     1/1
                                                           Running
                                                                               1.45
cloud-connector-proxy-759cc5b69f-kncvb
                                                     2/2
                                                           Running
                                                                               14h
                                                     0/2 Completed 0
cloud-connector-register-wormx
cloud-connector-update-license-status-6b8fc876c8-ch6k1
                                                    2/2 Running D
nex-ndr-enable-ids-b25vk
                                                     0/1 Completed 0
nsw-ndr-feature-switch-watcher-notifier-ndr-79886f475f-96tqc 1/1 Running D
                                                      0/1 Completed D
nax-ndr-setup-kefke-hq4jq
nsx-ndr-upload-config-db847544b-vm7mw
                                                     2/2
                                                           Running D
                                                                               14h
                                                                      D
nsx-ndr-worker-file-event-processor-8cc6b7f58-nhbgx
                                                    2/2 Ronning
                                                                               14h
                                                    2/2 Running D
nam-ndr-worker-file-event-uploader-77b4f9c95b-narcp
                                                                               Idh
                                                                     D
nsw-ndr-worker-ids-event-processor-df54bb4f9-bnvtk
                                                    2/2
                                                           Running
                                                                               14h
nsx-ndr-worker-monitored-host-uploader-6f6f6c9f75-j62vn
                                                                      D
                                                    2/2 Ronning
nsx-ndr-worker-ndr-event-processor-Sdffb58878-xSdfh
                                                     2/2 Running
                                                                     0
                                                            Running D
nsx-ndr-worker-ndr-event-uploader-685f7976d-2mj8x
                                                     2/2
```

In a healthy environment, all NSX Cloud Connector and NSX Network Detection and Response pods must show a status of Running or Completed. The example output in the screenshot corresponds to an environment with a Evaluation form factor in which NSX intelligence is not installed. In an Advanced form factor deployment with NSX intelligence installed, additional pods are available, such as nsx-ndr-worker-nta-event-processor that is responsible for processing Suspicious Traffic events from NSX intelligence.

8-115 Visualizing and Mitigating Attacks

NSX Network Detection and Response creates attack visualizations that give security teams the required context to quickly understand the scope of an attack and prioritize their response, including:

- Extent and duration of every event
- Active threats and affected hosts
- Communication between local and external systems
- Data sets accessed and harvested




8-116 Accessing the NSX Network Detection and Response UI

The NSX Network Detection and Response UI is installed as a plug-in in the NSX UI during the NSX Cloud Connector deployment.

You can access the NSX Network Detection and Response UI by clicking the GO TO CAMPAIGNS link in the Security Overview page.

Alternatively, you can access the NSX Network Detection and Response UI from the NSX Network Detection and Response tile in NSX Application Platform.



8-117 Campaign Overview: Active Threats and Attack Stages


NSX Network Detection and Response provides a summary of malicious activity in the network by showing the affected hosts and the stages of the active threats.

In alignment with the MITRE ATT&CK framework, NSX Network Detection and Response classifies malicious activity by attack stage.



8-118 Campaign Blueprint

NSX Network Detection and Response includes a dynamic blueprint that shows how an attack enters and moves laterally across the network, including compromised hosts and external communications.




8-119 Campaign Timeline

NSX Network Detection and Response also provides a detailed chronology of each stage of an attack to assist with remediation.



8-120 Reviewing Events

You can review all events discovered in the environment by selecting **Events** in the NSX Network Detection and Response UI.



Depending on the severity of the event and other events occurring in the environment at the time, an event might be correlated into a campaign. Regardless of whether the event is part of the campaigning, it appears on the Events page in the NSX Network Detection and Response UI.

8-121 Lab 16: (Simulation) Using NSX Network Detection and Response to Detect Threats

Install and use NSX Network Detection and Response to detect and visualize advanced threats:

- Install NSX Network Detection and Response
- Validate the NSX Network Detection and Response Deployment from the CLI
- Enable NSX Distributed IDS/IPS for a vSphere Cluster
- Create an NSX Distributed IDS/IPS Profile
- Configure NSX Distributed IDS/IPS Rules
- Generate Malicious Traffic
- Analyze Threat Detection Events and Campaigns

8-122 Review of Learner Objectives

- Describe NSX Network Detection and Response and its use cases
- Explain the architecture of NSX Network Detection and Response
- Activate NSX Network Detection and Response
- Describe the visualization capabilities of NSX Network Detection and Response

8-123 Key Points (1)

- NSX IDS/IPS uses real-time deep packet inspection to identify and prevent attempts at exploiting vulnerabilities in your applications.
- NSX Distributed IDS/IPS protects against malicious activity, including the exploits of known application-level vulnerabilities, application denial of service, lateral movement, and clientside and server-side exploits.
- You enable IDS/IPS at the Tier-1 gateway to protect the perimeter of your NSX environment.
- Behavior-based IDS/IPS helps to detect unusual traffic, malicious attacks, and security breaches in the network when compared to a baseline of normal traffic.
- NSX Application Platform is a container-based solution that is deployed on an existing Kubernetes cluster.
- NSX Application Platform must be deployed before using the NSX security features, such as NSX Malware Prevention, NSX Intelligence, NSX Network Detection and Response, and Metrics
- NSX Malware Prevention detects and prevents malicious file transfers by combining signaturebased detection of known malware with static and dynamic analyses of malware samples.

8-124 Key Points (2)

- East-west malware prevention protects the data center from the spread of internal malware and from malware that passes the network perimeter. It uses the Guest introspection agents, installed on every guest VM, to extract files.
- North-south malware prevention detects known malicious files when they enter the perimeter on the NSX Edge gateway firewall. It uses the IDS/IPS engine of the NSX Edge node to extract files.
- NSX Intelligence is a native distributed analytics solution that provides visibility and dynamic security policy enforcement for NSX environments.
- The NSX Suspicious Traffic Detection feature analyses the data collected by NSX Intelligence and flags suspicious activities using detectors.
- NSX Network Detection and Response is an advanced threat prevention platform that
 provides complete network visibility, detection, and prevention of sophisticated threats.
- NSX Network Detection and Response capabilities are integrated with on-premises NSX deployments.

Module 9

NSX Services

9-2 Importance

NSX provides several services that can help you address operational challenges in the virtual network architecture.

9-3 Module Lessons

- Configuring NAT
- Configuring DHCP and DNS Services
- 3. Configuring NSX Advanced Load Balancer
- IPSec VPN
- L2 VPN

9-4 Lesson 1: Configuring NAT

9-5 Learner Objectives

- Explain the role of network address translation (NAT)
- Distinguish between source and destination NAT
- Describe how reflexive NAT works
- Describe stateful active-active SNAT and DNAT operation
- Explain how NAT64 facilitates communication between IPv6 and IPv4 networks

9-6 About NAT

Network address translation (NAT) is a service that maps one IP address space to another. You can configure NAT on Tier-0 and Tier-1 gateways. NSX supports various NAT rule configurations, depending on the gateway type and the NSX Edge HA mode.

Supported NAT Rules	Gateway Type	NSX Edge HA Mode
SNAT / DNAT / NAT64	Tier-0 and Tier-1	Active-Standby
Reflexive NAT	Tier-0 only	Stateless Active-Active
Stateful Active-Active SNAT / DNAT	Tier-0 and Tier-1	Stateful Active-Active

Network address translation (NAT) was designed originally to conserve the public Internet address space. During the 1990s, Internet providers quickly depleted the available IPv4 address supply. NAT became the primary method for IPv4 address conservation. NAT performs one-to-one mapping (one public IP address is mapped to one private IP address) or one-to-many mapping (one public IP address is mapped to multiple private IP addresses).

You can create different NAT rules:


- Source NAT (SNAT) translates the source IP of the outbound packets to a known public IP address so that the application can communicate with the outside world without using its private IP address, SNAT also tracks the reply.
- Destination NAT (DNAT) enables access to internal private IP addresses from the outside world by translating the destination IP address when inbound communication is initiated. DNAT also manages the reply. For both SNAT and DNAT, users can apply NAT rules based on the 5-tuple match criteria.
- NAT64 is a mechanism for translating IPv6 packets into IPv4 packets.
- Reflexive NAT rules are stateless access control lists (ACLs) that must be defined in both directions. These rules do not track the connection. Reflexive NAT rules are applied when stateful NAT cannot be used.
- Introduced in NSX 4.0.1, stateful active-active SNAT/DNAT is achieved by using an external server IP hash to ensure symmetric northbound and southbound traffic through an edge by punting traffic between edge nodes.

9-7 About SNAT

The SNAT rule changes the source address in the IP header of a packet. The rule can also change the source port in the TCP or UDP headers.

In the diagram, as packets are received from the NAT VM, the T1-GW-01 Tier-1 gateway changes the source IP address of the packets from 172.16.101.11 to 80.80.80.1.

You can selectively bypass an existing SNAT rule for specific traffic by creating a No SNAT rule.



SNAT changes the source address in the IP header of a packet. It can also change the source port in the TCP/UDP headers.


The typical usage is to change a private (rfc1918) address into a public address for packets leaving your network. You can create a rule to either enable or disable the source NAT.

9-8 About DNAT

The DNAT rule changes the destination address in the IP header of a packet. It can also change the destination port in the TCP or UDP headers.

This rule is typically used to redirect incoming packets with a destination public IP address to a private IP address in the network.

You can selectively bypass an existing DNAT rule for specific traffic by creating a No DNAT rule.




9-9 About Reflexive NAT

You can use the reflexive NAT rule when a Tier-0 gateway runs in stateless active-active mode and when stateful NAT might lead to issues because of asymmetric paths.

Reflexive NAT rules are stateless access control lists (ACLs) that must be defined in both directions. These rules do not track connections and are sometimes called stateless NAT.

In the diagram, because the same NAT is applied to the NSX Edge Edge-1 and the NSX Edge Edge-2, asymmetric flows are supported.



When Tier-O is running in a stateless active-active mode, you cannot configure stateful NAT where asymmetrical paths might cause issues. For active-active routers, you can use reflexive NAT, which is also called stateless NAT.

For reflexive NAT, you can configure a single source address to be translated or a range of addresses. If you configure a range of source addresses, you must also configure a range of translated addresses. The size of the two ranges must be the same. The address translation is deterministic. The first address in the source address range is translated to the first address in the translated address range. The second address in the source range is translated to the second address in the translated range, and so on.

In the diagram, the source VM (172.16.101.11) on the internal network sends a packet to an external client (x,x,x,x) on the internet. The packet is routed to the Tier-O gateway hosted on NSX Edge Edge-2, which creates a reflexive NAT entry with source IP 172.16.101.11 and translated IP 80.80.80.1. When the return traffic arrives on NSX Edge Edge-1 (with the destination 80.80.80.1), the same reflexive NAT entry is used to translate 80.80.80.1 to 172.16.101.11. Because the same NAT is applied to NSX Edge Edge-1 and NSX Edge Edge-2, asymmetric flows are supported.

9-10 Configuring SNAT and DNAT

SNAT and DNAT are configured on the Tier-1 or Tier-0 gateway in the following use cases:

- You typically use SNAT to change a private address into a public address for packets leaving your network.
- You typically use DNAT to redirect incoming packets with a destination public address or port to a private IP address or port in your network.



To configure SNAT and DNAT, you provide values for the following options:

- Name: Provide a name for the NAT rule.
- Action: Specify the action of the NAT rule if a match occurs.
- Source IP: Specify a source IP address or an IP address range in CIDR format. If you leave
 this text box blank, the NAT rule applies to all sources outside the local subnet.

- Destination IP: Specify a destination IP address or an IP address range in CIDR format.
- Destination Port: Specify a destination port.
- Translated IP: The new IP address as the result of NAT.
- Service: Select a single service entry on which the NAT rule is applied.
- Applied To: Select objects that this NAT rule applies to. The available objects are gateways, interfaces, labels, service instance endpoints, and virtual endpoints.
- Enabled: To enable or disable the NAT rule.
- (Optional) Logging: Used in analysis and troubleshooting. The default is no logging.
- Priority: A lower value means a higher priority. The default is 0.
- Firewall includes the following options:
 - Match External Address: To match the firewall rule with the external address of the NAT rule
 - Match Internal Address: To match the firewall rule with the Internal address of the NAT rule
 - Bypass: To skip the firewall stage
- (Optional) Tags: Specify to group NSX objects.

9-11 Configuring Reflexive NAT

When a Tier-O gateway runs in stateless active-active mode, you can use reflexive NAT:

- The address translation is sequential, for example, the first address in the source range is translated to the first address in the translated range, and so on.
- The two ranges (source range and translated range) must be of equal size.




To configure reflexive NAT, you provide values for the following options:

- Name: Provide a name for the NAT rule.
- Action: Specify the action, Reflexive, if a match occurs.
- Source IP: Specify a source IP address or an IP address range in CIDR format. If you leave
 this text box blank, the NAT rule applies to all sources outside the local subnet.
- Destination IP: Specify a destination IP address or an IP address range in CIDR format.
- Destination Port: Specify a destination port.
- Translated IP: The new IP address as the result of NAT.
- Service: Select a single service entry on which the NAT rule is applied.
- Applied To: Select objects that this NAT rule applies to. The available objects are gateways, interfaces, labels, service instance endpoints, and virtual endpoints.
- Enabled: To enable or disable the NAT rule.
- (Optional) Logging: Used in analysis and troubleshooting. The default is no logging.
- Priority: A lower value means a higher priority. The default is 0.

- Firewall includes the following options:
 - Match External Address: To match the firewall rule with the external address of the NAT rule
 - Match Internal Address: To match the firewall rule with the Internal address of the NAT rule
 - Bypass: To skip the firewall stage
- (Optional) Tags: Specify to group NSX objects.

9-12 Stateful Active-Active Services

NSX version 4.0.1 introduces a new architecture that supports stateful active-active services. This architecture includes SNAT and DNAT. This architecture supports stateful active-active services by pinning specific flows to an individual edge.



Northbound and southbound traffic path selections are enforced by a common ECMP mechanism, which is a hash of the external server IP.

In the diagram, by enforcing a common ECMP hash based on the external server IP, with the ability to punt traffic between edges, the stateful session is maintained with bidirectional traffic through Edge-4.

Northbound path selection:


- The compute hypervisor has 4-way ECMP to the edge SRs.
- The compute hypervisor performs a 5-tuple hash, based on protocol number, source address, destination address, source port, and destination port.
- In this example, Edge-4 is selected as the northbound path.

Southbound path selection:

- Both Tor-Left and Tor-Right have 4-way ECMP to the Edge SRs.
- In this example, Tor-Right has performed a hardware vendor hash, selecting Edge-3.
- An IP hash is performed, based on the external server source IP, and traffic is punted from Edge-3 to Edge-4.

9-13 Stateful Active-Active Interface Group

Stateful active-active introduces the requirement for an Interface Group to which stateful services, such as NAT, are applied.




The stateful active-active interface group has the following requirements:

- The Interface Group has exactly one Interface from every edge node in the edge cluster.
- A separate interface Group is required for each uplink interface on an edge.
- For Edges with two uplinks, two Interface Groups are required
- In the scenario, each edge has one uplink interface, and one interface group is required.

9-14 Stateful Active-Active SNAT

Stateful active-active SNAT is similar to the active-standby SNAT, where an SNAT rule carries the translated IP address and all nodes share the same SNAT IP in the A/A cluster.




Symmetric traffic path selection is achieved by enforcing a common ECMP hash based on the external server target IP.

The SNAT rule changes the source address in the IP header of a packet. In the example, the SNAT session is pinned to Edge-4 to consistently translate SNAT 172.16.101.11 traffic to 80.80.80.1.

9-15 Stateful Active-Active Configuration

On a new NSX gateway, enable Stateful Active-Active using the Stateful toggle, when the HA mode is set to Active-Active.



The feature does not support converting from a stateless to a stateful Active-Active gateway configuration. A new gateway must be deployed.

After you enable the gateway stateful services and save the configuration, you cannot modify the HA mode for the gateway.

9-16 Configuring Interface Groups

From the gateway configuration, interfaces and interface groups, create an external interface group to include all gateway uplink interfaces.



In the example, the Tier-O Gateway has two uplink interfaces. You must add both the interfaces to the interface group.

9-17 Configuring Stateful SNAT and DNAT

Stateful active-active SNAT and DNAT configuration are similar to the active-standby SNAT and DNAT configuration, except that the NAT rule is applied to the interface group.




Stateful active-active SNAT and DNAT configuration are similar to the active-standby SNAT and DNAT configuration, except for the following parameter:

Applied To: Select the interface group to which this NAT rule applies.

9-18 About NAT64

NAT64 is a mechanism for translating IPv6 packets into IPv4 packets:

- NAT64 allows IPv6-only clients to communicate with IPv4 servers.
- Changes are not needed in the IPv6 or IPv4 nodes.
- NAT64 is supported on Tier-0 and Tier-1 gateways.
- NAT64 is stateful and requires the Tier-O gateway to be deployed in active-standby mode.
- NAT64 requires the Tier-1 gateway to be configured with an active-standby edge cluster.



The NAT64 mechanism enables IPV6 to IPV4 connectivity. NAT64 is based on the RFC 6146 and RFC 6145 standards.

NAT64 allows an IPv6-only client to initiate communications to an IPv4-only server. IPv6 must initiate the traffic

NAT64 translates IPv6 packets to IPv4 packets and forwards them to the IPv4 network. This functionality is designed so that changes are not required to either IPv6 or IPv4 nodes.

9-19 Configuring NAT64 Rules

To configure NAT64 rules:

- Navigate to Networking > Network Services > NAT.
- Select NAT64 from the View drop-down menu.
- Specify the Tier-0 or Tier-1 gateway where you want to create the rule and click ADD NAT64 RULE.



In the example, NAT64 rules are created on T0-GW-01.

You configure the following parameters in a NAT64 rule:

- Name: Name for the NAT64 rule.
- Action: The only supported action is NAT64, which translates between IPv6 and IPv4 addresses.
- (Optional) Source: IPv6 or IPv6 address range in CIDR format.
- Destination: IPv6 or IPv6 address range in CIDR format. The prefix must be /96 because the destination IPv4 IP is embedded as the last 4 bytes in the IPv6 address.
- Translated: This address is the IPv4 address or address pool for the source IPv6 address.
 The IPv4 address(es) must have a route enabled at the Tier-1 gateway if this is configured on the Tier-0 gateway or the return path is unknown.
- Applied To: NAT64 rules can only be applied to uplink interfaces or tags.
- Enabled: To enable or disable the NAT64 rule.
- (Optional) Logging: Used in analysis and troubleshooting. The default is no logging.
- Priority: A lower value means a higher priority. The default is 0.

- Firewall includes the following options:
 - Match External Address: To match the firewall rule with the external address of the NAT rule
 - Match Internal Address: To match the firewall rule with the internal address of the NAT rule
 - Bypass: To skip the firewall stage
- (Optional) Tags: Specify to group NSX objects.

9-20 Lab 17: Configuring Network Address Translation

Configure source and destination network address translation rules on the Tier-1 gateway:

- Prepare for the Lab
- Create a Tier-1 Gateway for Network Address Translation
- Create a Segment
- Attach a VM to NAT-Segment
- Configure NAT
- Configure NAT Route Redistribution
- Verify the IP Connectivity

9-21 Review of Learner Objectives

- Explain the role of network address translation (NAT)
- Distinguish between source and destination NAT
- Describe how reflexive NAT works
- Describe stateful active-active SNAT and DNAT operation
- Explain how NAT64 facilitates communication between IPv6 and IPv4 networks

9-22 Lesson 2: Configuring DHCP and DNS Services

9-23 Learner Objectives

- · Explain how DHCP and DHCP Relay are used for IP address allocation
- Configure DHCP services in NSX
- Describe how to use a DNS forwarder service
- Configure DNS forwarder services in NSX

9-24 About DHCP

DHCP allows clients to automatically obtain network configuration settings, such as IP addresses, subnet masks, default gateways, and DNS configuration, from a DHCP server.

DHCP Discover: "Locate all available DHCP servers"



DHCP Offer: "DHCP Server IP is 172.16.0.1, Proposed IP address 172.16.0.100/24."

The DHCP protocol eliminates the need for individually configuring network devices manually.

The operation can be summarized as follows:

- The DHCP client broadcasts a DHCP Discover message to locate all available DHCP servers on the subnet.
- A DHCP server broadcasts a DHCP Offer message, informing the client that it is available, proposing a client IP address, subnet mask, default gateway IP address, DNS IP address, IP lease time, and a DHCP server IP address.
- The DHCP client broadcasts a DHCP Request message to the server, requesting the proposed IP network configuration data.
- The DHCP server broadcasts an acknowledgment.
- The DHCP client configures its network interface with the proposed IP network configuration.


The IP allocation process is based on broadcast traffic.

A DHCP server must be available to host a range of addresses for each client subnet.

9-25 About DHCP Relay

A DHCP relay agent offers a more centralized approach to providing DHCP services across multiple subnets.

A DHCP relay agent forwards requests and replies between a DHCP server and a DHCP client, offering the flexibility of placing the DHCP server on a remote network.



From the DHCP client perspective, the IP allocation process remains broadcast based.

The DHCP relay agent converts the local DHCP broadcast message into a unicast message and sends the unicast message to the DHCP server.

A DHCP relay agent is required for each DHCP client subnet.

By using DHCP relay agents, a DHCP server can service multiple DHCP client subnets.

9-26 DHCP in NSX

NSX supports three types of DHCP services.


DHCP Service Type	escription		
DHCP Local Server	A DHCP server that is centrally managed by NSX. This server is local to a single segment.		
DHCP Relay	A DHCP relay agent is local to a single segment that relays client requesto an external DHCP server.		
Gateway DHCP	CP A DHCP server that is centrally managed by NSX. This server is available to all segments that are connected to the gateway.		

9-27 DHCP Local Server

DHCP Local Server is a DHCP service managed by NSX that is local to the segment and not available to the other segments in the network.

- It provides a dynamic IP assignment service only to the VMs that are attached to the segment.
- It runs as a service (service router) in the edge nodes of an NSX Edge cluster.
- The IP address of a local DHCP server must be in the subnet that is configured on the seament.
- It supports Tier-0 and Tier-1 gateways and segments not connected to a gateway.
- It can service VLAN-backed or overlay-backed segments.

DHCP Local Server Topology



In this configuration, the DHCP requests are managed in the NSX environment without relying on an external DHCP server.

DHCP Local Server runs as a service (service router) in the edge nodes of an NSX Edge cluster.


For standalone segments, DHCP Local Server is selected by default.

9-28 DHCP Relay

A DHCP relay agent is local to a single segment and relays client requests to an external DHCP server.

The DHCP servers can be located inside or outside of the SDDC.

The DHCP client segment must be connected to either a Tier-0 or Tier-1 gateway.




9-29 Gateway DHCP

Gateway DHCP is a DHCP server that is centrally managed by NSX. This server is available to all segments that are connected to the gateway.

By default, segments that are connected to a Tier-0 or Tier-1 gateway use Gateway DHCP.

Gateway DHCP Server Topology



The IP address of a Gateway DHCP server can be different from the subnets that are configured in the segments.

Individual segments connected to a gateway configured for Gateway DHCP can be selectively configured to use different DHCP types, Local DHCP Server, or DHCP Relay.

9-30 DHCP Workflow



9-31 Creating a DHCP Profile

To create a DHCP Profile, click **Networking > Networking Profiles > DHCP > ADD DHCP PROFILE** and select either **DHCP Relay** or **DHCP Server** as the required profile type.



To create a DHCP Profile, you provide values for the following options:


- Name: You use this name to identify the DHCP profile.
- Profile Type: Select either DHCP Relay or DHCP Server.
- Server IP Address is based on DHCP type
- (Optional) Tags: Add tags to label static bindings so that you can quickly search or filter bindings, troubleshoot and trace binding-related issues, or do other tasks.

For DHCP Server profiles, you provide values for these additional options:

- Edge Cluster: Select an NSX Edge cluster from the drop-down menu.
- Auto Allocate Edges: Yes (default) The first edge node becomes the active edge for DHCP services, and the second edge node becomes the standby edge for DHCP services.
- Standby Relocation: No (default) If the edge node where the active or standby DHCP server is running fails, a new standby DHCP server is created on another edge node to maintain high availability.

9-32 Assigning the DHCP Profile to a Segment

To assign the DHCP profile to a segment, click **Networking > Segments > Edit an existing segment > SET DHCP CONFIG** and select a configured profile.




To assign the DHCP profile to a segment, you provide values for the following options:

- DHCP Type: Select either DHCP Relay or DHCP Server.
- DHCP Profile: Select the configured DHCP profile.

9-33 Assigning the DHCP Profile to a Gateway

To assign the DHCP Profile to a gateway, click **Networking** > **Tier-0** or **Tier-1 Gateway**, edit the gateway, set the DHCP Config, and select a configured profile.




To assign the DHCP profile to a gateway, you provide values for the following options:

- Type: Select either DHCP Relay or DHCP Server.
- DHCP Profile: Select the configured DHCP profile.

9-34 Configuring DHCP Services

To configure DHCP services, click **Networking > Segments > Edit an existing segment > EDIT DHCP CONFIG.**



To configure DHCP services, you provide values for the following options:


- DHCP Type: Select the existing DHCP Relay or DHCP Server by name.
- DHCP Profile: Select the configured DHCP profile.
- DHCP server Address: If you are configuring a DHCP local server, server IP address is required.
- DHCP Ranges: Optionally enter the DHCP IP pool range.
- Lease Time: Optionally enter the amount of time in seconds for which the IP address is bound to the DHCP client.
- DNS Servers: Optionally enter the IP address of the domain name server (DNS) to use for name resolution. A maximum of two DNS servers are permitted.

9-35 About DNS Services

A DNS server translates domain names to IP addresses.

A DNS zone is a distinct portion of the domain name space in DNS.

Depending on the DNS zone, a DNS forwarder can forward DNS queries to specific DNS servers.




The resource to resolve, vmbeans.lab.com, is in the lab.com DNS namespace. The DNS forwarder can be configured to forward DNS queries to the lab.com DNS zone servers.

9-36 About DNS Forwarder

In NSX, the DNS client requests can be forwarded to the external DNS server by configuring a DNS forwarder in Tier-0 or Tier-1 gateways. The DNS forwarder provides the following functions:

- Forwards DNS requests from clients to upstream DNS servers
- Can optionally forward DNS requests to FQDN zones
- Caches the responses received from the upstream servers, reducing system load and improving performance




9-37 DNS Workflow



9-38 Creating DNS Zones

A default DNS zone is required. Optionally, you can configure one or more FQDN DNS zones. To create DNS zones, click **Networking > DNS > DNS Zones > ADD DNS ZONE**.



To create DNS zones, you provide values for the following options:

- Name: You use this name to identify the DNS zone.
- Domain: Select Any for the default zone. Select an FQDN for the domain for an FQDN zone.
- DNS servers: Enter the IP address of up to three remote DNS servers for this DNS zone.
- (Optional) Source IP: You must specify a source IP if the DNS forwarder service listener IP is an internal address that is not reachable from the external upstream DNS server.
- (Optional) Tags: Specify to group NSX objects.

When you configure a DNS zone, you can specify a source IP for a DNS forwarder to use when forwarding DNS queries to an upstream DNS server. If you do not specify a source IP, the DNS query packet source IP will be the DNS forwarder's listener IP.

9-39 Creating DNS Forwarder Services

To create a DNS forwarding services, click **Networking > DNS > DNS Services > ADD DNS SERVICE**.



To create DNS forwarder services, you provide values for the following options:

- Name: You use this name to identify the DNS forwarder service.
- Tier-O/Tier-1 Gateway: Specify the gateway for the service.
- DNS service IP: Clients send DNS queries to this IP address, which is also known as the DNS forwarder's listener IP.
- Default DNS Zone: Select the default DNS zone.
- (Optional) FQDN Zones: Select up to five FQDN zones.
- Admin Status: Enable or disable the DNS forwarder service.
- (Optional) Log Level: Used in analysis and troubleshooting. The default log level is info.
- Cache Size: Specify the DNS forwarder cache size.
- (Optional) Tags: Specify to group NSX objects.

When you configure a DNS zone, you can specify a source IP for a DNS forwarder to use when forwarding DNS queries to an upstream DNS server. If you do not specify a source IP, the DNS query packet source IP will be the DNS forwarder's listener IP.

9-40 Establishing Forwarder Connectivity

The external upstream DNS servers require IP connectivity to the DNS forwarder, which can be achieved using route advertisement or SNAT.




If the DNS listener IP is not routed, forwarder connectivity can be established by:

- Using gateway route advertisement to announce the listener IP.
- Using gateway SNAT to translate the listener IP to a public IP.

9-41 Configuring DHCP to Allocate DNS

To optionally configure DHCP services to allocate DNS Servers, click **Networking > Segments > Edit an existing segment > EDIT DHCP CONFIG,** and configure DNS Servers.



9-42 Review of Learner Objectives

- Explain how DHCP and DHCP Relay are used for IP address allocation
- Configure DHCP services in NSX
- Describe how to use a DNS forwarder service
- Configure DNS forwarder services in NSX

9-43 Lesson 3: Configuring NSX Advanced Load Balancer


9-44 Learner Objectives

- Describe NSX Advanced Load Balancer and its use cases
- Explain the NSX Advanced Load Balancer architecture
- Deploy NSX Advanced Load Balancer
- Explain the NSX Advanced Load Balancer components and how they manage traffic
- Configure virtual IP addresses, server pools, and virtual services

9-45 About NSX Advanced Load Balancer

VMware will use NSX Advanced Load Balancer as part of its load-balancing strategy.

NSX Advanced Load Balancer provides multicloud load balancing, web application firewall, application analytics, and container ingress services across data centers and clouds.



9-46 Benefits of NSX Advanced Load Balancer

NSX Advanced Load Balancer offers several advantages.



NSX Advanced Load Balancer offers the following advantages:

- End-to-end automation: NSX Advanced Load Balancer fully automates the life cycle management and placement of load-balancing components.
- Higher performance: NSX Advanced Load Balancer provides optimal traffic flows with no traffic hairpinning. Additionally, it supports ECMP-based active-active scale-out mode.
- Easy to troubleshoot: NSX Advanced Load Balancer offers native built-in log analysis and rich analytics tools that provide end-to-end visibility of the environment. This improved visibility reduces the troubleshooting time from days to minutes.
- On-demand scalability: The NSX Advanced Load Balancer platform automatically scales horizontally based on the traffic needs and rebalances the load across all components to ensure high performance.

9-47 NSX Advanced Load Balancer Feature Edition Comparison (1)

The table provides a feature comparison between the Basic and Enterprise licenses of NSX Advanced Load Balancer.



NSX Data Center Advanced and NSX Data Center Enterprise Plus editions include the NSX Advanced Load Balancer Basic entitlement. The Basic entitlement provides load-balancing features that are equivalent to the native NSX load balancer.

The NSX Advanced Load Balancer Enterprise edition requires an additional license and provides all features available in NSX Advanced Load Balancer, including GSLB, WAF, multicloud support, and so on.

For more information about the licenses and features available for NSX Advanced Load Balancer, see "NSX Advanced Load Balancer Editions" at https://avinetworks.com/docs/22.1/nsx-license-editions/


9-48 NSX Advanced Load Balancer Feature Edition Comparison (2)

Pfet	Feature	(Body)	Februaries
Cartainer Hightsi	Server Tues (III	Yes	Yes
	Committee	Almind Ingress	Tut Ingress DNS are USLE integration High Ells cluster is (mail) AZ ausport
Section (Section Pathon)	Assemble	Base LB agreestration	Pully must remain it as enture Milds! Must starte and overse, was out and part those plans; Made data and part those plans;
	Additional Management of the Control	Admir Tambie one	Active/Active depayments BOT and stood notice SCHIT support framework hard temperature
	Facility and selection	×	Plantins DN across terunts/strum/as grozes
	Commence :	No Assess Class?	monte imagnitions with all major states & private illustration and contact an increase our preference.
	<u> </u>	×	Come Management of a 5,5mm (Mountly Thrond Statistics)
All Control of the Co	An annual of Annual of	×	Americal application between plant in the control of the control o

9-49 NSX Advanced Load Balancer Architecture

The NSX Advanced Load Balancer architecture consists of two main components:

- NSX Advanced Load Balancer Controller:
 - Central repository for the configuration and policies related to load-balancing services.
 - Provides a user interface to perform configuration and management tasks.
 - Responsible for deploying Service Engines.
 - Typically deployed as a three-node cluster for high availability.
- Service engine (SE):
 - Performs application load-balancing operations.
 - Collects real-time analytics from application traffic flows.




The NSX Advanced Load Balancer architecture is built on software-defined principles. It separates the data and control plane to deliver scalable application load balancing. The platform provides a dynamic, centrally managed pool of load-balancing resources for virtual machines and containers. Since NSX 4.0.0.1, the configuration of NSX Advanced Load Balancer using the NSX UI and NSX API is deprecated. This feature should not be used because it will be removed completely in future releases. The installation of the NSX Advanced Load Balancer Controller cluster and the cross-launch of the NSX Advanced Load Balancer UI from NSX will continue to be supported. However, all configuration tasks for load balancers integrated with NSX environments should be performed directly through the NSX Advanced Load Balancer UI or API.

NSX Advanced Load Balancer Controller continues to interact with vCenter Server and NSX Manager through APIs for autodiscovery of SDDC objects such as ESXI hosts, segments, Tier-1 Gateways, and so on.

9-50 NSX Advanced Load Balancer Deployment Workflow

To deploy NSX Advanced Load Balancer:

- The user uploads the NSX Advanced Load Balancer Controller .cva file to NSX Manager.
- NSX Manager Invokes the NSX Advanced Load Balancer Controller node deployment in vCenter Server.
- vCenter Server deploys NSX Advanced Load Balancer Controller.
- On startup, NSX Advanced Load Balancer Controller registers with NSX Manager by using the Avi Lifecycle Manager (Avi LCM).
- After registration, Avi LCM performs the basic configuration of NSX Advanced Load Balancer Controller.
- In a cluster deployment, the NSX Advanced Load Balancer Controller Instances 2 and 3 are deployed, and the cluster is created.
- Service engine VMs are created or deleted based on the load-balancer configuration.



The NSX Advanced Load Balancer Controller OVA must be version 20.1.6 or later. Previous versions of the image are not accepted. Only one OVA file can be uploaded at a time. The disk space required for the OVA bundle is about 4 GB.

Before registration of NSX Advanced Load Balancer Controller with the Avi Lifecycle Manager, trust is established between the two entities using certificates.


During the provisioning phase, the Avi Lifecycle Manager configures the following parameters on the controller:

- Admin user password
- DNS and NTP servers
- Controller cluster information, including the cluster VIP if specified during the deployment

9-51 NSX Advanced Load Balancer Consumption Workflow

Since NSX 4.0.0.1, all services related to NSX Advanced Load Balancer are configured directly using the NSX Advanced Load Balancer UI or API:


- The user configures NSX Advanced Load Balancer through the NSX Advanced Load Balancer UI or API.
- NSX Advanced Load Balancer Controller receives the load balancing configuration and disseminates it to the service engines.
- The Service Engines perform load balancing operations and manage all client and serverfacing network interactions.



9-52 Requirements for NSX Advanced Load Balancer

Before using NSX Advanced Load Balancer in your NSX environment, you must complete the following tasks:


- Deploy the NSX Advanced Load Balancer Controller cluster for the NSX UI.
- Create a cloud connector from the NSX Advanced Load Balancer UI.
- (Optional) Create a service engine group from the NSX Advanced Load Balancer UI.



A cloud connector is used to integrate the NSX Advanced Load Balancer platform with the NSX environment. The cloud connector defines the connectivity information for the service engines, including the NSX segments and the Tier-1 gateway they are connected to.

9-53 Deploying the NSX Advanced Load Balancer Controller Cluster

You can deploy the NSX Advanced Load Balancer Controller cluster from the NSX UI by navigating to System > Configuration > Appliances > NSX Advanced Load Balancer.



Before deploying the NSX Advanced Load Balancer Controller cluster:

- vCenter Server must be registered with NSX as a Compute Manager.
- A virtual IP address must be configured for the NSX Advanced Load Balancer Controller cluster.

Each NSX Advanced Load Balancer Controller cluster requires only one management IP address. This IP address is used to configure the controller. The management IP address is also used by the controller to communicate with the service engines.

In a cluster deployment, the management IP addresses for all controllers must belong to the same subnet.


An NSX Advanced Load Balancer Controller cluster includes one or three nodes. All nodes must be deployed individually. To ensure that the cluster quorum is maintained, the deployment of the second node is queued until the third node is deployed.

9-54 Service Engines Deployment and Connectivity

The NSX Advanced Load Balancer Controller cluster creates or deletes service engine VMs. based on the load-balancer configuration.

Each service engine requires two IP addresses:

- Management IP address: Used for communication with the NSX Advanced Load Balancer Controller cluster
- Data plane IP address: Used for communication with the server pool network.



9-55 Creating a Cloud Connector (1)


Before NSX Advanced Load Balancer Controller can create service engine VMs, you must create a cloud connector for the NSX environment.

The cloud connector defines the connectivity information for both the management and the data plane networks of the service engines, including:

- NSX transport zone
- NSX seament
- Tier-1 gateway used as the default gateway

In a typical deployment, separate NSX segments are used for the management and data plane interface of the service engines.

All required NSX constructs must be manually preconfigured in NSX Manager before creating the cloud connector.




Placing the service engines on NSX segments directly connected to a Tier-0 gateway is not supported.

Both NSX overlay and NSX VLAN-backed segments can be used to configure the management and data plane network for the service engines.

All required NSX constructs, including transport zones, Tier-1 gateway, and segments, must be manually preconfigured in NSX Manager before creating the cloud connector.

9-56 Creating a Cloud Connector (2)

You create a Cloud Connector from the NSX Advanced Load Balancer UI by navigating to Infrastructure > Clouds.



The cloud connector wizard first connects to NSX Manager to fetch the network constructs available and presents them as options to configure the SE management and data plane networks.


Only one management network is supported for all SE groups created for the NSX environment.

9-57 Creating a Service Engine Group

A service engine group is a method of grouping service engines to provide data plane isolation and redundancy:

- Multiple SE groups might exist in the NSX environment.
- SE groups are used to manage load-balancing traffic for a given load balancer service.
- If a service engine fails, another service engine within the same SE group takes over.

Creating a service engine group is optional. A default service engine group is automatically created for the NSX cloud. However, you can modify the existing service engine group or create another service engine group.



To create a service engine group from the NSX Advanced Load Balancer UI, navigate to Infrastructure > Cloud Resources > Service Engine Group.

The high availability mode of the SE group controls the behavior of the SE group if an SE failure occurs. It also controls how the load is distributed across SEs.

The following high availability modes are available:

- Legacy high availability active-standby mode: This mode is primarily intended to mimic a legacy appliance for easy migration to NSX Advanced Load Balancer. It deploys two services engines; one in active and the other in standby mode.
- Elastic high availability N+M mode: This default mode deploys N active SEs for loadbalancing purposes, but also deploys M additional SEs within the service group as a buffer to absorb any SE failures.
- Elastic high availability active-active mode: This high availability mode load balances services
 across a minimum of two SEs that are both in active mode.

Active-standby is the only availability mode supported with the Basic Edition licensing. To avail of N+M and active-active high availability modes, you must purchase an NSX Advanced Load Balancer Enterprise license.

The VS Placement across SEs option determines whether the creation of load-balancing services also creates new SEs or uses those already available.

The following options exist:

- Compact: Attempts to place load balancer services on already existing service engines. This
 option is default for elastic high availability N+M and legacy high availability active-standby
 mode.
- Distributed: Maximizes load-balancing performance by deploying new service engines and avoiding placements on existing SEs. This option is the default for elastic high availability active-active mode.

For more information about the SE group configuration options, see "Service Engine Group" at https://avinetworks.com/docs/22.1/service-engine-group/.


9-58 NSX Advanced Load Balancer Components

NSX Advanced Load Balancer Includes several components:

 A virtual service is a software construct containing a virtual IP address, a port, and a protocol.

The virtual service is associated with a single SE group and attached to a Tier-1 gateway.


- A server pool is associated with a virtual service and includes the group of servers responsible for load balancing the client requests.
- A health monitor is attached to a server pool and verifies the status of the servers within it.
- A persistence profile is attached to a server pool and reconnects clients to the same pool member.



A virtual service is a software construct containing a virtual IP address, a port, and a protocol. External clients use this combination to access the servers behind the load balancer.

9-59 NSX Advanced Load Balancer Topologies

NSX Advanced Load Balancer supports different topologies for the SE and server pool deployment.



Service Engines on a dedicated segment

Service Engines on a shared segment

NSX Advanced Load Balancer supports two different topologies for the SE and server pool deployment:

- Service engines on a dedicated segment: This option allows you to manage the IP address
 assignments for the SE data plane interfaces and the server pool separately. The NSX
 segment used for the SE data plane must be created in NSX before creating an NSX Cloud
 Connector in NSX Advanced Load Balancer Controller.
- Service engines on a shared segment: With this option, the SE data plane interfaces share
 the same address space as the server pool servers and reside in the same NSX segment.


9-60 VIP Placement and Route Redistribution

Before deploying a virtual service, the Tier-1 and Tier-0 gateways must be configured to distribute the LB VIP.

During the virtual service configuration, the following actions are automatically performed:

- The VIP is placed on one or more service engines depending on the high availability mode configured.
- VIP static routes are created on the Tier-1 gateway where the SE data plane network is connected.

The VIP static routes include as many next hops as available service engines.




The NSX administrator is expected to configure the Tier-1 gateway to advertise the virtual service VIP with the Tier-0 gateway. For north-south reachability of the VIP, the administrator should also configure the Tier-0 gateway to redistribute the VIP to the external router through BGP.

After a virtual service is placed on an SE group, NSX Advanced Load Balancer Controller creates VIP static routes on the Tier-1 gateway where the SE data plane network is connected. The VIP static routes include as many next hops as available service engines. These static routes do not need to be advertised or redistributed, because they are only locally relevant to the Tier-1 gateway and the back-end service engines for ECMP purposes.

9-61 North-South Traffic

An external client request to the VIP is managed as follows:


- The external client traffic enters through the uplink of the Tier-O gateway.
- The Tier-0 gateway forwards the traffic to the appropriate Tier-1 gateway.
- Using the static VIP routes, the Tier-1 gateway routes the request to the VIP on the service engines.
- The service engines forward the traffic to the back-end server pool through the Tier-1 gateway.



9-62 East-West Traffic (1)

An Internal client request to a VIP connected to the same Tier-1 gateway is managed as follows:


- The request is sent to the Tier-1A gateway where the client network is connected.
- Using the static VIP routes, the Tier-1A gateway routes the request to the VIP on the service engines.
- The service engines forward the traffic to the back-end server pool through the Tier-1A gateway.



9-63 East-West Traffic (2)

An internal client request to a VIP connected to a different Tier-1 gateway is managed as follows:

- The request is sent to the Tier-1B gateway where the client network is connected.
- The traffic is routed to the Tier-O gateway, which forwards it to the Tier-1A gateway.
- Using the static VIP routes, the Tier-1A gateway routes the request to the VIP on the service engines.
- The service engines forward the traffic to the back-end server pool through the Tier-1A gateway.



9-64 Accessing the NSX Advanced Load Balancer UI

Since NSX 4.0.0.1, the configuration of all NSX Advanced Load Balancer components must be performed directly through the NSX Advanced Load Balancer UI or API, even if the solution is integrated with an NSX environment.


You can cross-launch to the NSX Advanced Load Balancer UI from NSX or log in directly to the deployed NSX ALB Controller.



9-65 Creating a Virtual IP Address

You create a virtual IP address that you can associate with a virtual service.

During the creation of the virtual IP, you specify the Tier-1 gateway to which the service engine data plane network is connected.




You create a Virtual IP address by navigating to **Applications > VS VIPs** in the NSX Advanced Load Balancer UI.


During the creation of the virtual IP, you specify the cloud connector and the Tier-1 gateway to which the service engine data plane network is connected.

The Virtual IP address can be static or autoallocated from a preconfigured pool. It is possible to configure IPv4, IPv6 addresses, or both.

9-66 Creating a Server Pool

When configuring a server pool, you specify the pool members, load-balancing algorithm, health monitors, persistence profiles, and other parameters.





You create a Server Pool by navigating to **Applications > Pools** in the NSX Advanced Load Balancer UI.

You specify the following parameters during the creation of a server pool:

- Name: A unique name for the server pool.
- Cloud: The cloud connector details for the NSX environment.
- VRF Context: Virtual Routing Framework (VRF) is a method to isolate traffic in a system.
 VRF is also called a route domain in the load balancer community. A global VRF context is created by default. Network administrators might create custom VRF contexts to isolate traffic between different tenants or subsets.
- Default Server Port: New connections to servers will use this destination service port. The default port is 80.
- Load-balancing algorithm: The selected load-balancing algorithm controls how the incoming connections are distributed among the servers in the pool.
- Tier-1 gateway (logical router): Specify the Tier-1 gateway that you want to attach the server pool to. This value matches the Tier-1 gateway specified for the virtual service and VIP.

- Servers or pool members can be configured according to the following criteria:
 - IP Address, range, or DNS name: Manually enter the IP address, DNS name, and port for each member.
 - IP Group: Select an existing IP pool profile or create a profile.
 - Security Group: Select a native NSX grouping.
- You can verify server health by applying one or more health monitors in the Health Monitor section. Active monitors generate traffic from each service engine and mark a server up or down based on the response. The passive monitor listens only to client-to-server communication.
- In the Profiles and Policies section, you can configure a persistence profile to ensure that
 subsequent connections from the same client connect to the same server. Persistent
 connections are critical for most servers that maintain the client session information locally,
 such as HTTP applications that need to keep users' information for some time. This section
 also allows you to configure autoscaling policies for the load balancer.
- In the SSL section, you can enable SSL encryption between the service engines and the back-end servers.

For more information about the server pool configuration options, see the Avi Networks website at https://avinetworks.com/docs/22.1/architectural-overview/applications/pools/.

9-67 Configuring Load-Balancing Algorithms


Load-balancing algorithms control how the incoming connections are distributed among the servers in the pool. The table includes a summary of the commonly used algorithms.

Methods	
New connections are sent to the server that currently has the least number of outstanding concurrent connections.	
New connections are sent to the next eligible server in the pool in sequential order.	
New connections are sent to the server that is currently providing the fastest response to new connections or requests.	
New connections are distributed across the servers using a hash key.	
New connections are sent to the server with the lightest load, regardless of the number of connections that server has.	
Instead of trying to distribute all connections or requests across all servers, NSX Advanced Load Balancer determines the fewest number of servers required to satisfy the current client load.	

For a full list of supported load balancing algorithms, see "Load Balancing Algorithms" at https://avinetworks.com/docs/22.1/load-balancing-algorithms/

9-68 Configuring Health Monitor Profiles

Health monitors validate the status of the servers in a pool to make forwarding decisions. You can configure health monitor profiles for HTTP, HTTPS, TCP, and UDP applications amongst others.



In the server pool creation wizard, you use the **Health Monitor** tab to configure active and passive health monitors for the pool members.

Some active health monitors that are preconfigured in the system are as follows:

- System-HTTP: Sends a request to a web server and validates either the HTTP response code or the HTML response data.
- System-HTTPS: Used to validate the health of HTTPS encrypted web servers. It relies on the same mechanism as the HTTP monitor.
- System-Ping: Validates the status of the server by sending a ping command and receiving a reply.
- System-TCP: Validates that the server can successfully establish a TCP connection.
- System-UDP: Sends a UDP datagram to the server and matches the server's response against the expected response data.
- System-DNS: Queries name servers for an A record and matches the resolved response against an expected IP address.

You can also create custom health monitor profiles for your applications.

9-69 Configuring Persistence Profiles

Persistence profiles ensure the stability of stateful applications by directing all related connections to the same back-end server.

NSX Advanced Load Balancer supports the following types of persistence profiles:

- Client IP
- HTTP Cookle
- Custom HTTP Header
- App Cookle
- TLS



In the server pool creation wizard, you use the **Profiles/Policies** tab to configure persistence profiles.

NSX Advanced Load Balancer supports the following types of persistence profiles:

- Client IP: NSX Advanced Load Balancer maintains a table with the client's source IP address.
 The client's source IP address is kept in the table for a given period of time. While the IP address remains in the table, any new connection by the user will be sent to the same server.
- HTTP Cookie: The NSX Advanced Load Balancer service engines insert an HTTP cookie into a server's first response to a client. In this mode, no configuration changes are required on the back-end servers.
- Custom HTTP Header: This method allows to specify an HTTP header for persistence.
- App Cookie: Rather than NSX Advanced Load Balancer inserting its own cookie into HTTP responses for persistence, this mode relies on an existing cookie that is inserted by the back-end server.
- TLS: Applicable to virtual services that are terminating SSL or TLS. This method embeds
 user persistence information within TLS session information.

9-70 Configuring Server Pool Security Settings

You can encrypt the traffic between the service engines and the back-end servers by enabling SSL in the server pool.

The chosen SSL profile defines which ciphers and SSL versions are used to encrypt the traffic.



In the server pool creation wizard, you use the SSL tab to configure security settings for the server pool.


You can encrypt the traffic between the service engines and the back-end servers by specifying the following parameters:

- SSL Profile: The SSL profile defines which ciphers and SSL versions are supported to encrypt the traffic.
- Server SSL Certificate Validation PKI Profile: This option validates the certificate presented by the server against the selected PKI profile. When not enabled, the service engine automatically accepts the certificate presented by the server when sending health checks.
- Service Engine Client Certificate: When establishing an SSL connection with a server, either
 for normal client-to-server communications or when executing a health monitor, the service
 engine presents this certificate to the server.

Additionally, you can also enable Common Name Check and TLS SNI from the SSL tab.

9-71 Creating a Virtual Service

When creating a virtual service, you specify its virtual IP, port and protocol combination, and the server pool responsible for managing the back-end traffic.



You create a server pool by navigating to **Applications > Virtual Services** in the NSX Advanced Load Balancer UI. The NSX Advanced Load Balancer UI offers two methods for creating a virtual service: basic setup and advanced setup. The advanced setup allows users more granular control over the settings of the virtual service, including the reuse of preconfigured or existing server pools. The same information included in the advanced setup is available whenever a Virtual Service is edited, regardless of how it was created (Basic or Advanced). In both the basic and advanced setups, you must specify the cloud connector before configuring the virtual service.

In an advanced setup, you specify the following main parameters during the creation of a virtual service instance:

- Name: Provide a unique name for the new virtual service.
- Vs VIP: Select an existing virtual service IP address from the drop-down menu or create an address.
- TCP/UDP Profile: Use the TCP/UDP profile to set the virtual service to listen for UDP Fast Path, TCP Fast Path, or TCP Proxy. If the application profile is set to an HTTP profile, the TCP/UDP profile must be set to proxy the TCP connections.


- Application Profile: Determines the behavior of the virtual services, based on application type. Some commonly used profiles are as follows:
 - System-DNS: Default for processing DNS traffic and uses UDP port 53.
 - System-FTP: Used for load balancing FTP workloads for both active and passive FTP modes.
 - System-HTTP: Default for processing nonsecure layer 7 HTTP traffic and uses port 80.
 - System-L4-Application: The virtual service listens for layer 4 requests on the port that you specify in the Service Port field. Select this option to use the virtual service for non-HTTP applications, such as mail or a database.
 - System-Secure-HTTP: Default for processing secure layer 7 HTTPS traffic. Uses port 443.
 - System-Secure-HTTP-VDI: Default for processing virtual desktop HTTP traffic.
 - System-SSL-Application: The virtual service listens for secure layer 4 requests.
 Selecting this option autopopulates port 443 in the Service Port field.
 - System-Syslog: Default for processing Syslog traffic.
- Services: Specify a port number or range for the virtual service.
- Pool/Pool Group: Select the pool to use for this virtual service using the Pool drop-down
 menu. A pool may only be associated with one virtual service. However, the same servers
 can belong to multiple pools using unique port combinations.

For more information about the Virtual Service configuration options, see Create a Virtual Service at https://avinetworks.com/docs/22.1/architectural-overview/applications/virtual-services/create-virtual-service/.

9-72 Validating Virtual Services and Server Pools

On the **Applications** tab in the NSX Advanced Load Balancer UI, you can review the following details for each virtual service and server pool instance:

- Analytics
- Logs
- Health
- Clients/servers
- Security (virtual service)
- Events
- Alerts



On the **Applications** tab in the NSX Advanced Load Balancer UI, you can review the following details for each virtual service and server pool instance:

- Analytics: Provides insights into performance through the real-time analysis of key performance indicators.
- Logs: Logs are Indexed on NSX Advanced Load Balancer Controller. Logs can be viewed
 and filtered locally in the NSX Advanced Load Balancer UI.
- Health: The health score denotes both the responsiveness of the virtual service or pool, and any vulnerabilities.

- Clients/servers: The Clients tab is available for virtual services and displays information about clients accessing that service. The Servers tab is available for server pools and displays information about the status, health, and throughput of each member.
- Security: This tab is only available for virtual services. It provides detailed security information, including SSL and DDoS information.
- Events: Events are used to provide a history of relevant changes that have occurred in a virtual service or service pool.
- Alerts: Alerts are intended to inform administrators of significant events within a virtual service or service pool.

9-73 Lab 18: Configuring NSX Advanced Load Balancer

Configure NSX Advanced load-balancing services:

- Prepare for the Lab
- Create Segments for NSX Advanced Load Balancer
- Deploy NSX Advanced Load Balancer Controller
- Access the NSX Advanced Load Balancer UI
- Create a Cloud Connector for NSX
- Configure Service Engine Networks and Routing
- Test the Connectivity to Web Servers
- Create a Virtual Service
- Configure Route Advertisement and Route Redistribution for the Virtual IP

9-74 Review of Learner Objectives

- Describe NSX Advanced Load Balancer and its use cases
- Explain the NSX Advanced Load Balancer architecture
- Deploy NSX Advanced Load Balancer
- Explain the NSX Advanced Load Balancer components and how they manage traffic
- Configure virtual IP addresses, server pools, and virtual services

9-75 Lesson 4: IPSec VPN

9-76 Learner Objectives

- Explain how IPSec-based technologies are used to establish VPNs
- Compare policy-based and route-based IPSec VPN
- Describe IPSec VPN requirements in NSX
- Create and configure IPsec VPN tunnels
- Identify Traceflow and Live Traffic Analysis enhancements for VPN tunnels

9-77 Use Cases for IPSec VPN

IPSec VPN has several use cases:

- Enables businesses to interconnect remote IP networks securely
- Extends IP networks to remote offices
- Uses a routing-based connection between different, nonoverlapping IP subnets
- Provides a secure communication channel for other nonsecure protocols, such as Generic Routing Encapsulation (GRE)



IPSec VPN secures the traffic that flows between two networks. These networks are connected over a public network through IPSec gateways called endpoints.

NSX Edge supports site-to-site IPSec VPN between an NSX Edge Instance and remote IPSeccapable gateways.

NSX Edge can be one or both endpoints, supporting site-to-site IPsec VPN with another NSX Edge or another vendor's IPsec gateway.

9-78 IPSec VPN Protocols and Algorithms

IPSec VPN includes several protocols and algorithms:

IPSec Service	IPSec Algorithms and Protocols	
Key management	Internet Key Exchange (IKE)	
Authentication	Preshared key	
	Certificates	
Encryption	Advanced Encryption Standard (AES)	
Data integrity	Secure Hash Algorithm (SHA)	

IPSec is not a single protocol. It is a suite of protocols designed to provide confidentiality, authentication, and integrity for a VPN.

To accomplish these goals, IPSec uses Internet Key Exchange (IKE) to:

- Manage the connection to a peer.
- Define security associations used to secure and validate data exchanges.
- Define security protocols used to carry IP traffic over the VPN.
- IKE runs over UDP port 500. If NAT is detected in the gateway, the port is set to UDP 4500.
- IKEv1 (RFC 2409) and IKEv2 (RFC 5996) are supported.

Security Associations (SA): An SA is a basic component of IPSec and contains information about the security parameters negotiated between peers.

The following types of SAs are available:

- IKE (or ISAKMP) SA
- IPSec SA


The IKE SA is used for the control plane of the VPN and contains a combination of mandatory and optional values:


- Encryption Algorithm: Mandatory
- Hash Algorithm: Mandatory
- Authentication Method: Mandatory
- Diffie-Heilman Group: Mandatory
- Lifetime: Optional

9-79 IPSec VPN Methods

IPSec VPN tunnel packets can use different types of headers:

- The authentication header (AH) provides data integrity and authentication without encryption.
- The encapsulating security payload header (ESP) provides encryption, data integrity, and authentication.





The authentication header does not provide encryption, whereas the encapsulating security payload header enables the encryption of the protected payload.

9-80 IPSec VPN Modes

IPSec VPN supports the following modes:

- Transport mode:
 - Preserves the original IP header
 - Typically used for remote-access VPN
- Tunnel mode:

New IP Header


ESP

Header

IP

TCP

- Encapsulates the entire IP datagram with a new header
- Used for site-to-site VPN between IPSec-enabled gateways
- Mode used by NSX





Mode

Auth

ESP

Trailer

Data

9-81 IPSec VPN Types

The following IPSec VPN types are available.

- Policy-based VPN:
 - A VPN policy is used to determine which traffic is protected by IPSec and passes through the VPN tunnel.
 - This static configuration requires modification of the VPN policy when network topology changes occur.
- Route-based VPN:
 - The remote IPSec VPN gateway is a BGP peer and protected local and remote networks are learned based on routes exchanged by using BGP.
 - Routes are learned over a specific interface called a Virtual Tunnel Interface (VTI).



Additional route-based VPN properties:


- All packets routed through the VTI are protected by using IPSec.
- OSPF dynamic routing is not supported for routing through IPSec VPN tunnels.
- Tunnel redundancy is supported:
 - By using BGP only.
 - Do not use static routing or OSPF to achieve VPN tunnel redundancy.
 - The primary tunnel is active, whereas the secondary tunnel is standby.
 - If the peer is unreachable on the primary tunnel, the secondary tunnel becomes active.

9-82 NSX IPSec VPN Deployment

When you deploy IPSec VPN, you should consider several factors:

- IPSec VPN services are available on both Tier-1 and Tier-0 gateways.
- Protected networks must be segments created through the NSX UI or policy APIs.
- Segments can be connected to either Tier-0 or Tier-1 gateways to use VPN services.
- Tenants with overlapping networks require NAT on Tier-0 gateways.
- VPN-based dynamic routing for VTI is supported on Tier-0 gateways only.
- NSX supports site-to-site IPSec VPNs in tunnel mode.
- IPSec tunnels use the DPDK-accelerated performance.

9-83 IPSec VPN: High Availability



IPSec VPN high availability has the following characteristics:

- VPN supports active-standby high availability on the Tier-1 and Tier-0 gateways.
- VPN services use gateway-level fallover.
- VPN configuration and VPN state are synchronized.
- Tunnels are re-established without renegotiation on failover.

Additional IPSec VPN high availability properties:

- This feature is supported for both policy-based and route-based IPSec VPN services.
- You can use high availability virtual IP addresses for external connections.

9-84 IPSec VPN Workflow



All Step 2 profile configurations are optional.

In Step 4, you add an IPSec VPN session that is either policy-based or route-based.

9-85 Configuring an IPSec VPN Service

To configure an IPSec VPN service, you select **Networking** > **VPN** > **VPN Services** > **ADD SERVICE** > **IPSec**. Then you specify the values for the options.



You configure the following settings for an IPSec service:

- Name: You enter a name for the IPSec service.
- Tier-0/Tier1 Gateway: From the Tier-0/Tier-1 Gateway drop-down menu, select a Tier-0/Tier-1 gateway to associate with this IPSec VPN service.
- Admin Status: This option enables or disables the IPSec VPN service. By default, the value is set to Enabled to enable the service on the Tier-O gateway.
- IKE Log Level: This option enables VPN service logging. The IKE logging level determines
 the amount of information that you want collected for the IPSec VPN traffic. The default is
 set to the Info level.
- Session Sync: This option enables or disables the stateful synchronization of VPN sessions.
 By default, this parameter is set to Enabled.
- Tags: You enter a value for tags if you want to include this service in a tag group.

9-86 Configuring DPD Profiles

Dead peer detection (DPD) is a method for detecting whether an IPSec connection is alive. To configure a DPD profile, you select **DPD PROFILES** on the **PROFILES** tab.



This step is optional after you configure an IPSec service.

A DPD profile specifies the number of seconds to wait between probes to detect whether an IPSec peer is alive.

NSX provides a system-generated DPD profile, nsx-default-l3vpn-dpd-profile, which is assigned by default when you configure an IPSec VPN service.

To configure a DPD profile, you select values for the following options:

- Name: You use the name to identify the service.
- DPD Probe Mode:
 - Periodic: For a periodic DPD probe mode, a DPD probe is sent every time the specified DPD probe interval time is reached.
 - On Demand: For an on-demand DPD probe mode, a DPD probe is sent if no IPSec packet is received from the peer site after an idle period. The value in the DPD Probe Interval text box determines the idle period used.
- DPD Probe Interval (sec): You provide a value in seconds to define at which times a DPD detection packet should be sent.
- Retry Count: The number of allowed retries.
- Admin Status: This setting enables or disables the profile.
- Tags: For cloud-based installations, almost every entity can hold a tag.

9-87 Configuring IKE Profiles

The Internet Key Exchange (IKE) profile defines the algorithms that are used to authenticate, encrypt, and establish a shared secret. To configure an IKE profile, you select **IKE PROFILES**.



This step is optional after you configure an IPSec service.

You specify the following settings to configure an IKE profile:

- Name: You use the name to identify the service.
- IKE Version: The options are IKE V1, IKE V2, or IKE FLEX. The selection depends on your business requirements.
- Encryption Algorithm: The encryption algorithm used during the internet Key Exchange (IKE) negotiation.
- Digest Algorithm: The secure hashing algorithm used during the IKE negotiation.
- Diffie-Hellman: The cryptography schemes that the peer site and the NSX Edge instance
 use to establish a shared secret over an insecure communications channel.
- SA Lifetime (sec): The lifetime (in seconds) of the security associations (individual communicating peer identifiers) after which a renewal is required.
- Tags: For cloud-based installations, almost every entity can hold a tag.

9-88 Configuring IPSec Profiles

The IPSec profile defines the security parameters used for negotiations to establish and maintain a secure tunnel between two peers. To configure the IPSec profile, you select **IPSEC PROFILES**



This step is optional after you configure an IPSec service.

You provide values for the following settings to configure the IPSec profile:

- Name: You use the name to identify the service.
- Encryption Algorithm: The encryption algorithm used during the Internet Protocol Security (IPSec) negotiation.
- Digest Algorithm: The secure hashing algorithm used during the IPSec negotiation.
- PFS Group: This setting specifies the Perfect Forward Secrecy (PFS) group, which adds
 protection to the keys used for building secure channels. You can enable or disable this
 option.
- SA Lifetime (sec): The setting specifies the lifetime (in seconds) of the security associations (individual communicating peer identifiers) after which a renewal is required.
- DF Bit: This setting defines whether the encrypted traffic should copy the Don't Fragment (DF) bit from the inner payload to the encrypted traffic.
- Tags: For cloud-based installations, almost every entity can hold a tag.

9-89 Adding a Local Endpoint

To add the local endpoint, you select **Local Endpoints**. This step is required in preparation for configuring an IPSec VPN session.



You provide values for the following settings to add the local endpoint:

- Name: You use the name to Identify the local endpoint.
- VPN Service: This setting is the predefined service to use with this session.
- IP Address: The IP address of the local endpoint.
 - For an IPSec VPN service running on a Tier-0 gateway, the local endpoint IP address must be different from the Tier-0 gateway's uplink interface IP address.
 - For an IPSec VPN service running on a Tier-1 gateway, the route advertisement for IPSec local endpoints must be enabled in the Tier-1 gateway configuration.
- Site Certificate: Local site certificate, used for certificate-based authentication mode for the IPSec VPN session.
- Trusted CA / Self Signed Certificates: Remote site certificate, used for certificate-based authentication mode for the IPSec VPN session.
- Certificate Revocation List: A list of digital certificates that were revoked before their scheduled expiration date and should no longer be trusted.
- Local ID: Used for identifying the local NSX Edge instance. This local ID is the peer ID
 configured on the remote site. The local ID can be any string but is typically the public IP
 address of the VPN or a fully qualified domain name (FQDN) for the local VPN service.
- Tags: For cloud-based installations, almost every entity can hold a tag.

9-90 Configuring IPSec VPN Sessions

You add an IPSec session to define the VPN type: either policy-based or route-based. You can find this configuration option by selecting **Networking** > **VPN** > **IPSec Sessions** > **ADD IPSEC SESSION**.



9-91 Configuring Policy-Based IPSec VPN Sessions (1)

When you select the policy-based option, IPSec tunnels are used to connect multiple local subnets that are behind the NSX Edge Instance, with peer subnets on the remote VPN site.




To configure the policy-based IPSec session, you specify the following settings:

- Name: You use the name to identify the service when you need to use it.
- Type: This setting is already defined by the previous selection.
- VPN Service: This setting is the predefined service to use with this session.
- Local Endpoint: This setting is the earlier configured local endpoint for use with this
 configuration session.
- Remote IP: The setting specifies the IP address of the remote IPSec-capable gateway for building the secure connection.
- Authentication Mode: This setting defines whether to use the preshared key (PSK) or certificate-based connection authentication.
- Local Networks and Remote Networks: These settings define the interesting traffic that should be encrypted through this VPN session.
- Pre-shared Key: This setting specifies the string to define the key if the authentication mode is PSK.
- Remote ID: This setting defines the identifier of the remote peer for verifying the authenticity of the peering.
- Tags: For cloud-based installations, almost every entity can hold a tag.

9-92 Configuring Policy-Based IPSec VPN Sessions (2)

In the **Advanced Properties** section, you select the predefined IKE, IPsec, and DPD profiles from the drop-down menus. You also define which side initiates the connection.



9-93 Configuring Route-Based IPSec VPN Sessions

When you select the route-based option, tunneling is provided on traffic that is based on routes that were learned dynamically over a virtual tunnel interface (VTI).




To configure a route-based IPSec session, you define the following settings:

- Name: You use the name to identify the service for later use.
- Type: This setting is already defined by the previous selection.
- VPN Service: This setting is the predefined service to use with this session.
- Local Endpoint: This setting defines an earlier configured local endpoint to use with this session configuration.
- Remote IP: This setting defines the IP address of the remote IPSec-capable gateway for building the secure connection.
- Compliance suite: You can specify a security compliance suite such as CNSA, FIPS, Foundation, PRIME, or Suite-B to configure the security profiles used for an IPSec VPN session.
- Authentication Mode: This setting defines whether to use a preshared key (PSK) or certificate-based connection authentication.
- Admin Status: This setting enables and disables the service.
- Tunnel Interface: This setting defines the IP address of the local virtual tunnel interface (VTI) that is created to use with this session.
- Pre-shared Key: This setting provides the string for defining the key if the authentication mode is PSK.
- Remote ID: This setting specifies the identifier of the remote peer for verifying the authenticity of the peering.
- Tags: For cloud-based installations, almost every entity can hold a tag.

9-94 Traceflow and Live Traffic Analysis Support for VPN

NSX 4.0.0.1 introduces the following enhancements to observe live packets in a VPN tunnet:

- IPSec components are displayed in packet observations in Traceflow and Live Traffic Analysis.
- You can observe the ESP packets on a Tier-0 uplink interface in Live Traffic Analysis.
- You can observe the cleartext packets on an IPsec Tunnel Interface.



9-95 Review of Learner Objectives

- Explain how IPSec-based technologies are used to establish VPNs
- Compare policy-based and route-based IPSec VPN
- Describe IPSec VPN requirements in NSX
- Create and configure IPsec VPN tunnels
- · Identify Traceflow and Live Traffic Analysis enhancements for VPN tunnels

9-96 Lesson 5: L2 VPN


9-97 Learner Objectives

- Describe L2 VPN technologies in NSX
- Identify various supported L2 VPN endpoints
- Create and configure L2 VPN secure connections

9-98 About Layer 2 VPN

Layer 2 VPN connectivity enables you to extend L2 networks across data centers securely.

- Extends L2 networks (overlay-backed segments and VLAN-backed segments) across sites
 on the same broadcast domain
- Enables VM mobility, such as vSphere vMotion migration and disaster recovery without IP address changes
- Enables hybrid cloud solutions



L2 tunnels are established between L2 VPN endpoints that are interconnected over L3 networks.


The NSX L2 VPN architecture includes:

- L2 VPN server: The client is connected to this destination.
- L2 VPN client: This source initiates communication with the destination L2 VPN server.

9-99 L2 VPN Architecture

NSX L2 VPNs are established using the following layers of encapsulation:

- An outer Internet Protocol Security (IPSec) tunnel to provide secure encrypted communication
- An Inner Generic Routing Encapsulation (GRE) tunnel to support multicast traffic




Combined, GRE tunnels over IPSec provide the secure extension of multicast traffic between sites, required to stretch an L2 broadcast domain. On their own:

- IPsec tunnels are unable to support multicast traffic
- GRE tunnels are not secure.

9-100 L2 VPN Edge Packet Flow

The inbound and outbound L2 VPN Edge packet flows illustrate the sequence of operations that includes GRE and IPSec tunneling.



For outbound L2 VPN traffic (traffic from the Internal network behind the edge node) that is destined for any remote L2 network, the first step is to decapsulate the Geneve frames. The destination address of the internal frame designates whether traffic goes through the local bridge port toward remote sites or is locally managed. Further steps are inserting the ID for the proper VLAN and sending the traffic to the local VTI interface to encapsulate into GRE, which gets protected by IPSec and is forwarded to any given destination.

In the inbound direction, when receiving L2 VPN traffic that is identified by the IPSec engine, the traffic requires IPSec decryption first and GRE decapsulation. After being sent to the bridge interface, traffic is sent to local networks. The required Geneve encapsulation parameters are based on the actual tunnel IDs for the traffic

9-101 L2 VPN Considerations

When deploying L2 VPN, you must consider several important points:

- L2 VPN services are supported on both Tier-0 and Tier-1 gateways.
- Segments can be connected to either Tier-1 or Tier-0 gateways to use L2 VPN services.
- Only one L2 VPN service (either client or server) can be configured for either Tier-0 or Tier-1 gateway.
- An L2 VPN session can extend up to 512 L2 segments.




In NSX, layer 2 VPN has the following additional characteristics:

- The hub-and-spoke topology is supported.
- Tunnel redundancy is not supported. L2 VPN relies on edge high availability.
- L2 VPN interoperability is available only in NSX and does not support third-party interoperability.
- The supported number of L2 VPN server and client sessions depends on the size and type
 of the edge node. To review recommended L2 VPN configuration limits, see the VMware
 Configuration Maximums tool at https://configmax.vmware.com/home.

9-102 Recommended L2 VPN Clients

The following L2 VPN clients are supported:



The following L2 VPN clients are recommended:

- NSX Managed NSX Edge in a separate NSX Managed environment.
 - Overlay and VLAN segments can be extended.
- Autonomous Edge:
 - Enables L2 VPN access from a non-a NSX environment to NSX environments.
 - Deployed by using an OVF file on a host that is not managed by NSX.
 - Only VLAN segments can be extended.


Additional L2 VPN Client details including the supported software versions are available in the NSX Administration Guide at https://docs.vmware.com/en/VMware-

NSX/4.0/administration/GUID-86C8D6BB-F185-46DC-828C-1E1876B854E8.html.

9-103 About Autonomous Edge

Autonomous edge has the following characteristics:

- Is deployed by using an OVF file on a host that is not managed by NSX.
- Extends VLAN-backed segments
- Acts as an NSX Edge gateway, which can be deployed on on-premises data centers and public clouds (for example, Amazon AWS and Microsoft Azure)
- Runs independently without the management plane/central control plane installed in the NSX domain
- Is powered with the high-performing Data Plane Development Kit (DPDK)



9-104 Sample L2 VPN Network Topology

This sample L2 VPN network topology and IP addressing assignments serve as a guide to illustrate parameters used in the L2 VPN configuration steps.



At each and of the tunnel the VPN Tunnel ID is mapped to the Segment VNI or VLAN ID.

9-105 L2 VPN Server Workflow

The L2 VPN Server uses the IPSec features of the gateway. To configure an L2 VPN server, you must first configure a supporting IPSec VPN and then an L2 VPN.



9-106 Configuring an IPSec VPN Service

To configure the IPSec VPN service to be used for L2 VPN, you select **Networking > VPN > VPN Services > ADD SERVICE > IPsec**, to associate the service with a Tier-0 or Tier-1 gateway.



To configure the IPSec for L2 VPN service, you provide values for the following options:

- Name: You use this name to identify the IPSec for L2 VPN service.
- Service Type: IPSec has been selected.
- Tier-0/Tier-1 Gateway: Specify the gateway for the service.
- Admin Status: To enable or disable the IPSec for L2 VPN session.
- Session Sync: To enable or disable the stateful synchronization of the IPSec for L2 VPN session.
- IKE Log Level: The Internet Key Exchange log level. The default is Info level.
- Tags: Enter a value for Tags if you want to include this service in a tag group.
- Global Bypass Rules: Enter the list of local and remote subnets between which IPSec protection is bypassed.

9-107 Adding a Local Endpoint

To add the IPSec endpoint to be used for L2 VPN, you select **Networking > VPN > Local Endpoints.** You click **ADD LOCAL ENDPOINT** to define the local side of the IPSec connection.



To add the local endpoints, you provide values for the following options:

- Name: You use this name to identify the local endpoints.
- VPN service: This setting specifies which IPSec VPN service to use with the endpoint.
- IP Address: This setting is the IPSec tunnel local endpoint IP address.
- Site Certificate: You use this setting with certificate-based authentication to specify which
 certificate to use with this endpoint.
- Trusted CA / Self Signed Certificate: A public key certificate.
- Local ID: This setting specifies the IPsec ID of the local side. The local ID is usually the same
 as the local IP address.
- Tags: For cloud-based installations, almost every entity can hold a tag.

9-108 Adding an L2 VPN Server Service

To add the L2 VPN Server service, you select **Networking > VPN > VPN Services**. To begin, you click **ADD SERVICE** and select **L2 VPN Server** to define an L2 VPN Server service.



To add the L2 VPN server service, you provide values for the following options:

- Name: You use this name to identify the L2 VPN server service.
- Service Type: L2 VPN Server has been selected.
- Tier-0/Tier-1 Gateway: Specify the gateway for the service.
- Sessions: To apply an L2 VPN server session to this L2 VPN server service.
- Hub & Spoke: By default, the value is set to Disabled, which means the traffic received from the L2 VPN clients is only replicated to the segments connected to the L2 VPN server. If this property is set to Enabled, the traffic from any L2 VPN client is replicated to all other L2 VPN clients.
- Tags: For cloud-based installations, almost every entity can hold a tag.

9-109 Adding an L2 VPN Server Session

To add the L2 VPN Server session, you select **Networking > VPN > L2 VPN Sessions**. You click **ADD L2 VPN SESSION** to complete the L2 VPN Server configuration.

A system-generated IPSec VPN session supports the L2 VPN session.




To add the L2 VPN session, you provide values for the following options:

- Name: You use this name to identify the L2 VPN session.
- Mode: Server has been selected.
- VPN service: This setting specifies which L2 VPN server service to use with this L2 VPN session.
- Local Endpoint/IP: This setting specifies which local endpoint to use with this L2 VPN session.
- Remote IP: The IP address of the client-side IPSec tunnel endpoint.
- Segments: Used to connect logical segments to this L2 VPN server service.
- Pre-shared Key: A shared secret common to both L2 VPN client and the L2 VPN server configurations.
- Admin Status: To enable or disable the L2 VPN server session.
- Remote ID: This setting specifies the IPsec ID of the remote side. The remote ID is usually
 the same as the remote IP address.
- Tunnel Interface: The IP address of the server-side GRE tunnel endpoint.
- Tags: For cloud-based installations, almost every entity can hold a tag.

9-110 Attaching Segments to the L2 VPN (1)

You identify the segments that should be extended through the L2 VPN tunnels. You can click the vertical ellipsis icon to edit an existing segment or create a segment.

You can optionally configure a local egress gateway IP so that all VMs on the segment use it as their default gateway.



When you attach the L2 VPN segments, the following key settings are available:

- L2 VPN: This setting defines the previously configured L2 VPN session. The segment that is
 defined is used through that session.
- VPN Tunnel ID: This number is used to identify the communicating local and remote L2 networks. The same ID on both sides means that they are on the same L2 broadcast domain.
- Local Egress Gateway IP: The IP address of the local gateway that the VMs on the segment use as their default gateway. The same IP address can be configured in the remote site on the extended segment.

9-111 Attaching Segments to the L2 VPN (2)

You can also attach segments on the L2 VPN Session page. After selecting the edit mode and clicking **SEGMENTS**, you can add segments and define the tunnel ID for each segment.



9-112 L2 VPN Client Workflow

Perform the following steps to deploy an L2 VPN Client.



For steps 1 and 2, follow the L2 VPN server configuration steps, switching local and remote endpoint IPs.

9-113 Configuring an L2 VPN Client Service

To configure an L2 VPN Client service, you select **Networking > VPN > VPN Services**. To begin, you click **ADD SERVICE** and select **L2 VPN Client** to define an L2 VPN Client service.




To configure the L2 VPN Client service, you provide values for the following options:

- Name: You use this name to identify the L2 VPN client service.
- Service Type: L2 VPN client has been selected.
- Tier-0/Tier-1 Gateway: Specify the gateway for the service.
- Tags: For grouping NSX objects.

9-114 Adding an L2 VPN Client Session (1)

Before creating an L2 VPN Client session, you must obtain or download the peer code from the L2 VPN server.



The peer code is a Base64-encoded configuration string that is available from the L2 VPN server through the **DOWNLOAD CONFIG** option or through a REST API call.

9-115 Adding an L2 VPN Client Session (2)

You configure the local and remote IP addresses and the peer code, which was retrieved in a previous step. You can also click the **Admin Status** toggle to enable or disable the session.




To configure the L2 VPN client session, you provide values for the following options:

- Name: You use this name to identify the L2 VPN client session.
- Mode: Client has been selected.
- VPN service: This setting specifies which L2 VPN client service to use with this L2 VPN session.
- Local Endpoint/IP: This setting specifies which L2 VPN client local endpoint to use with this L2 VPN session.
- Remote IP: The IP address of the server-side IPSec tunnel endpoint.
- Peer Configuration: The peer code downloaded from the L2 VPN server.
- Admin Status: To enable or disable the L2 VPN client session.

9-116 Attaching Segments to the L2 VPN Client

On the Set Segments page, you define the network and the tunnel ID. The same configuration is also available by selecting the **SEGMENTS** option.



9-117 Lab 19: Deploying Virtual Private Networks

Configure the VPN tunnel and verify the operation:

- Prepare for the Lab
- Deploy a New NSX Edge Node to Support the VPN Deployment
- Configure a New Edge Cluster
- Deploy and Configure a New Tier-O Gateway and Segments for VPN Support
- Create an IPSec VPN Service
- Create an L2 VPN Server and Session
- Configure a Predeployed Autonomous Edge as an L2 VPN Client
- Verify the Operation of the VPN Setup

9-118 Review of Learner Objectives

- Describe L2 VPN technologies in NSX
- Identify various supported L2 VPN endpoints
- Create and configure L2 VPN secure connections

9-119 Key Points (1)

- NAT can be configured on Tier-0 and Tier-1 gateways.
- Typically, source translation is used to change a private address to a public address for packets leaving your network.
- Typically, destination translation is used to redirect incoming packets with a destination of a
 public address to a private IP address in your network.
- Reflexive NAT can be used when a Tier-0 gateway runs in stateless active-active mode with asymmetric traffic paths.
- Tier-0 and Tier-1 gateways in stateful active-active mode support stateful services including SNAT and DNAT.
- The NAT64 mechanism translates IPv6 packets to IPv4 packets.
- DHCP is a standard networking protocol for dynamically distributing network configuration parameters, such as IP addresses for interfaces.
- A DNS is a computer application that implements a service for resolving a computer name to an IP address.

9-120 Key Points (2)

- Since NSX 4.0.0.1 the configuration of NSX Advanced Load Balancer using the NSX UI and NSX API is deprecated. Network administrators should configure load-balancers integrated with NSX environments directly through the NSX Advanced Load Balancer UI or API.
- NSX Advanced Load Balancer includes multiple components such as virtual IP address, virtual service, and a server pool with associated health and monitor profiles.
- IPSec VPN services are available on Tier-0 gateways to interconnect different IP networks.
- Using the GRE over IPSec, L2 VPN tunnels can be used to extend layer 2 networks.

Questions?

Module 10

NSX User and Role Management

10-2 Importance

You must manage users and roles to enforce the least user privilege and provide clear separation of duties. By integrating NSX with VMware Identity Manager or LDAP, you can configure role-based access control (RBAC) for external users.

10-3 Module Lessons

- Integrating NSX with VMware Identity Manager
- Integrating NSX with LDAP
- Managing Users and Configuring RBAC

10-4 Lesson 1: Integrating NSX with VMware Identity Manager

10-5 Learner Objectives

- Describe the purpose of VMware Identity Manager
- Identify the benefits of integrating NSX with VMware Identity Manager
- Configure the Integration between NSX and VMware Identity Manager
- Verify the integration between NSX and VMware Identity Manager

10-6 About VMware Identity Manager


VMware Identity Manager is an Identity as a service (IDaaS) solution.

VMware Identity Manager provides the following services for software as a service (SaaS), web, cloud, and native mobile applications:

- Application provisioning
- Conditional access controls
- Single sign-on (SSO)

VMware products can use VMware Identity Manager as an enterprise SSO solution.

VMware Identity Manager is based on the OAuth 2.0 authorization framework.



To verify the version compatibility between NSX and VMware Identity Manager, use the VMware Product Interoperability Matrix at https://interopmatrix.vmware.com/interoperability.

VMware Identity Manager is now offered as Workspace ONE Access.

10-7 Benefits of Integrating VMware Identity Manager with NSX

The Integration of VMware Identity Manager with NSX provides the following benefits related to user authentication:


- Support for extensive authentication, authorization, and accounting (AAA) systems, including:
 - RADIUS
 - Smart cards and common access cards
 - RSA SecureID
 - LDAP and OpenLDAP based on Active Directory (AD)
- Enterprise SSO:
 - Common authentication platform across multiple VMware solutions
 - Seamless SSO experience

NSX has its own native LDAP and Active Directory Integration, but VMware Identity Manager also offers this capability.

10-8 Prerequisites for VMware Identity Manager Integration

The following prerequisites must be met before integrating VMware identity Manager with NSX:

- Deploy the VMware Identity Manager appliance from an OVF template.
- Configure time synchronization for the VMware identity Manager virtual machine.
- Perform an initial configuration of the VMware Identity Manager appliance.



The following steps must be completed before integrating VMware Identity Manager with NSX:

- From the vSphere Client, deploy the VMware Identity Manager appliance from an OVF template.
- Synchronize the VMware Identity Manager virtual machine time with the ESXI host where it is running.
 - Right-click the VM and select Edit Settings > VM Options.
 - Scroll down to the Time section and select the Synchronize guest time with host check box.
- After deploying the VMware Identity Manager appliance, use the Setup wizard available at https://<VMware Identity Manager FQDN>.
 - a.. Set passwords for the admin, root, and remote SSH user.
 - b... Select a database.

10-9 Configuring VMware Identity Manager

After the initial setup, connect to the VMware Identity Manager administration console.

In the console, you configure the following settings:

- Identity sources
- Authentication methods
- Access policies



You can access the VMware identity Manager Administration console at https://<VMware_identity_Manager_FQDN>:443/SAAS/admin.

The following identity sources are supported in VMware identity Manager:

- Active Directory (AD) over LDAP or AD with integrated Windows authentication
- LDAP
- Local directory

To configure authentications methods, select **Identity & Access Management > Authentication Methods**.

To define access policies, select Identity & Access Management > Policies.

Administrators can configure rules that specify the network ranges and types of devices that users can use to sign in.

10-10 Overview of the VMware Identity Manager and NSX Integration


After both NSX and VMware Identity Manager appliances are deployed and configured, you can integrate these components:

- Create an OAuth client for NSX in VMware Identity Manager.
- Obtain the SHA-256 certificate thumbprint for the VMware identity Manager appliance.
- Configure the VMware Identity Manager details in NSX.

10-11 Creating an OAuth Client

Before enabling the integration of VMware Identity Manager and NSX, you must register NSX as a trusted OAuth client in VMware Identity Manager:

- From the VMware Identity Manager administration console, click the Catalog tab.
- Select Settings > Remote App Access.
- On the Clients page, click Create Client.



VMware identity Manager uses the OAuth 2.0 authorization framework to enable NSX and its users to access specific data and services.

Before enabling the integration between VMware Identity Manager and NSX, you must register NSX as a trusted OAuth client in VMware Identity Manager. When configuring NSX details, you select **Service Client Token** from the **Access Type** dropdown menu. This selection indicates that the application, NSX in this example, accesses the APIs for itself. The application does not access the APIs for a particular user.

You must specify a client ID to uniquely identify NSX. You need this value to enable the VMware Identity Manager integration.

You must also click **Generate Shared Secret**. You need this value to enable the VMware Identity Manager integration.

Leave the default settings for all other options.

On the Create Client page, you can optionally set the token time-to-live values by specifying the access, refresh, and idle timers.

10-12 Obtaining the SHA-256 Certificate Thumbprint

Before you configure the integration between NSX and VMware Identity Manager, you must obtain the SHA-256 certificate thumbprint of the VMware Identity Manager appliance:

- Use SSH to log in to the VMware Identity Manager appliance.
- Run the sudo -s command to gain root access.
- Navigate to the VMware Identity Manager configuration directory.

```
cd /usr/local/horizon/conf
```

Retrieve the SHA-256 certificate thumborint of VMware Identity Manager.

```
openssl x509 -in sa-nsxvidm-01.vclass.local_cert.pem -noout -sha256 -fingerprint
```


You need the SHA-256 certificate thumbprint value when you enable the VMware Identity Manager integration.

Use SSH to log in to the VMware Identity Manager appliance with user sshuser.

10-13 Configuring the VMware Identity Manager Details in NSX

To enable the VMware Identity Manager Integration from the NSX UI:

- Select System > Settings > User Management > Authentication Providers.
- Click the VMware Identity Manager tab.
- Click EDIT.
- Provide the configuration information and click SAVE.



In the **OAuth Client ID** and **OAuth Client Secret** text boxes, you enter the client ID and shared secret that you generated when you created the OAuth client for NSX in VMware Identity Manager.

In the SSL Thumbprint text box, you enter the SHA-256 certificate thumbprint value that you generated from the VMware identity Manager appliance command line.

The value entered in the **NSX Appliance** text box must be used to access NSX Manager after the integration. If you enter the fully qualified domain name (FQDN) of NSX Manager and try to access the appliance through its IP address, the authentication fails.

If a virtual IP (VIP) is set up in the NSX Management cluster, you cannot use the external load balancer integration even if you enable it. You can either have VIP or the external load balancer while configuring VMware Identity Manager, but not both. Disable VIP if you want to use the external load balancer.

10-14 Verifying the VMware Identity Manager Integration


You can validate the successful communication between NSX and VMware Identity Manager from the NSX UL



Navigate to System > Settings > User Management > Authentication Providers > VMware Identity Manager to validate the VMware Identity Manager Integration. If the integration is successful, the VMware Identity Manager Integration appears as Enabled. The VMware Identity Manager appliance, the OAuth Client ID, and the NSX Appliance fields are populated.

10-15 Default UI Login


The default login page appears when integration with VMware Identity Manager is not enabled.



The default login page also appears if integration with VMware identity Manager is configured, but VMware identity Manager is down or not reachable at the time of the login.


10-16 UI Login with VMware Identity Manager

After the integration with VMware identity Manager is enabled, you are redirected to the VMware identity Manager login page for authentication.



10-17 Local Login with VMware Identity Manager

For troubleshooting or administration, you might need to bypass VMware Identity Manager when the integration is enabled. Go to https://<NSX_Manager_FQDN>/logIn.jsp?local=true.



10-18 Review of Learner Objectives

- Describe the purpose of VMware Identity Manager
- Identify the benefits of integrating NSX with VMware Identity Manager
- Configure the Integration between NSX and VMware Identity Manager
- Verify the integration between NSX and VMware Identity Manager

10-19 Lesson 2: Integrating NSX with LDAP

10-20 Learner Objectives

- Identify the benefits of integrating NSX with LDAP
- Describe the LDAP authentication architecture
- Configure the Integration between NSX and LDAP

10-21 About LDAP

The Lightweight Directory Access Protocol (LDAP) is an Internet protocol for accessing and managing distributed directory services.

Distributed directory services store information about users and groups, the network infrastructure, and network services.

NSX supports the following directory services or identity sources:

- Active Directory over LDAP
- OpenLDAP



10-22 Benefits of Integrating LDAP with NSX

Integrating LDAP with NSX offers the following benefits:

- Reuses the existing directory service infrastructure
- Does not require the deployment of the VMware Identity Manager appliance
- Is simple to configure and manage

10-23 Authentication with LDAP

LDAP authentication operates as follows.



After Integrating LDAP with NSX, the authentication process is as follows:

- The user initiates a login request from the UI or the API.
- NSX Manager receives the login request and creates an LDAP bind request to the appropriate identity source, for example, Active Directory.
- The Identity source returns an LDAP bind response to NSX Manager.
- 4. The bind response might succeed or fail during authentication. If the status is success, NSX Manager provides the appropriate access privilege based on the assigned RBAC role for the user or group. If the status is failure, NSX Manager displays an authentication error.

10-24 Adding an Identity Source

You can add up to three Identity sources.



You can add a new identity source by navigating to System > Settings > User Management > Authentication Providers > LDAP

You specify the following settings as part of the identity source configuration:


- Name: The name of the identity source.
- Domain Name: The domain that you want to add as an identity source.
- Type

The following types of identity sources are available:

- Active Directory over LDAP
- Open LDAP
- LDAP servers: The connection settings to your LDAP servers.
- Base DN: The point from where a server searches for users.

10-25 Configuring the LDAP Server

As part of adding an identity source, you specify the connection settings to the LDAP server.



Three LDAP servers are allowed per LDAP domain. The servers are tried in order with each request.

You specify the following settings when configuring the connection to the LDAP server:

- Hostname/IP: The fully qualified domain name or the IP address of the LDAP server. In LDAPS configurations, the FQDN must match the host name in the LDAP server certificate.
- LDAP Protocol: LDAP (unsecured) or LDAPS (secured).
- Port: The default LDAP port 389 and LDAPS port 636 are used for the directory sync. Do not change the default values.
- Use StartTLS: This toggle is available only for the LDAP protocol. If enabled, SSL/TLS is
 used to establish a secure connection.
- Certificate: The LDAP server provides a certificate as part of the ADD/Check status workflow. Accept the certificate.
- Bind Identity: Domain account with read permission for all objects in the domain tree.
- Password: Password for the bind identity account.

To verify that you can connect to the LDAP server, click **Check Status** (under Connection Status).

10-26 UI Login with LDAP

You log in as an Active Directory or OpenLDAP user by specifying the domain name on the NSX login page.



If a domain is not provided, local authentication is performed.

10-27 Review of Learner Objectives

- Identify the benefits of integrating NSX with LDAP
- Describe the LDAP authentication architecture
- Configure the integration between NSX and LDAP

10-28 Lesson 3: Managing Users and Configuring RBAC


10-29 Learner Objectives

- Identify the different types of users in NSX.
- Recognize permissions and roles available in NSX
- Create and configure custom roles
- Explain object-based RBAC in a multitenancy environment
- Assign roles to users

10-30 NSX Users

The following types of users can access the NSX environment:

- Local users
- Principal identity users
- External users:
 - Active Directory and OpenLDAP users
 - Users managed by VMware Identity Manager



Local users have been preconfigured by the system:

- admin has the Enterprise Administrator role and cannot be modified.
- audit has the Auditor role. It can be renamed but it cannot be configured with other roles.
- guestuser1 and guestuser2 can be renamed and configured with RBAC roles.

Principal identity users are unique users. These users own the object that they create and ensure that the object can only be modified or deleted by the owning principal identity. Principal identity users are authenticated by client certificate. The authentication is local to NSX Manager. Principal identities are usually used by third-party management platforms, such as VMware Integrated OpenStack, Tanzu Kubernetes Grid Integrated Edition, VMware Aria Automation, and so on, but they are also used by NSX Application Platform.

The napp_platform_egress, napp_platform_ingress, and napp_platform_kafka principal identity users are added when NSX Application Platform is installed. Each of these users is configured with a specific role, which cannot be modified.

OpenID Connect is a simple security layer built on top of the OAuth2 authentication protocol. NSX supports the OAuth2 workflow that involves redirecting a user to an external identity provider (iDP). OpenID Connect users are authenticated by the IDP and the information about the authenticated sessions are sent to NSX. IDP partners are not supported.

10-31 Activating Guest Users

The two local users, guestuser1 and guestuser2, are inactive by default and can be activated by using the NSX UI or the API.

You can activate the two guest users by navigating to **System > Settings > User Management > Local Users.**



The two local users, guestuser1 and guestuser2, can only be activated by the admin user.


Unlike the admin and audit users, guestuser1 and guestuser2 can be configured with RBAC roles.

10-32 Using Role-Based Access Control

RBAC enables you to restrict system access to users based on their role in the company.

Users are assigned roles, and each role has specific permissions:

- Local users, admin, and audit are preconfigured with specific roles that cannot be modified.
- Guest users, principal identity users, and external users can be configured with any of the built-in roles or custom roles.



Role-based access control (RBAC) is a method to enforce the least privilege and separation of duties principles.

10-33 Built-In Roles (1)

NSX provides built-in roles.

Role	Description
Auditor	Read permissions on all features
Enterprise Admin	Full access permissions on all features
GI Partner Admin	Role used for third-party endpoint protection service insertion
LB Admin	Read permissions on all networking services and full access permissions on load-balancing features
LB Operator	Read permissions on all networking services and load-balancing features
Netx Partner Admin	Role used for third-party Network Introspection service insertion
Network Admin	Full access permissions on all networking services

10-34 Built-In Roles (2)

Role	Description
Network Operator	Read permissions on all networking services and execute permissions on monitoring and troubleshooting tools
Org Admin	Performs CRUD operations in the allocated Org object in a multitenancy environment
Project Admin	Performs CRUD operations in the allocated Project object in a multitenancy environment
Security Admin	Full access permissions on all security configurations
Security Operator	Read permission on all security configurations and execute permissions on monitoring and troubleshooting tools
Support Bundle Collector	Permission only to generate a support bundle
VPN Admin	Read permissions on all networking services and full access permissions on VPN features
630	

10-35 Custom Role-Based Access Control

The custom role-based access control (Custom RBAC) feature enables you to create custom roles in addition to the existing built-in roles.

Custom RBAC has the following common use cases:


- Provide flexibility to create custom roles and grant custom permissions to NSX users.
- Extend the current RBAC capabilities beyond the built-in roles with preconfigured permissions.
- Help companies to meet regulatory guidelines and compliance requirements for RBAC.



When default roles cannot enforce least user privileges and clear separation of duties, Custom RBAC helps enforce these options by providing more granularity. Custom RBAC provides flexibility and customization opportunities for specific deployment considerations and use cases.

10-36 Creating Custom Roles (1)


You can either clone an existing role or create a role in the NSX UI.



Only an Enterprise administrator can create a custom role. However, Enterprise administrators can delegate the custom role creation to another custom role.

10-37 Creating Custom Roles (2)

You can set permissions for the new role.




The following NSX features are not supported with the Custom RBAC role:

- Upgrade
- Migrate
- Fabric
- TraceFlow
- NSX Intelligence
- Inventory of physical servers and containers

10-38 Multitenancy Hierarchy Model in NSX 4.0.1

NSX 4.0.1 has expanded its data model to support multitenancy by introducing the following new constructs:

- ORG: Represents the provider space
- Project: Represents the tenant space



ORG and Project are referred to as objects that contain the NSX resources.

Default ORG is a multitenancy organization container created by default by the NSX system. The ORG can represent:

- Provider spaces in a service provider scenario
- Companies that might have different departments or business units

Project is a container that provides isolation for specific networking and security constructs in an organization.

The Project can represent:

- Tenant spaces in a service provider scenario
- Different departments or business units in a company

An ORG can have more than one Project. Each Project has its own separate and isolated space.


Projects are created by an Enterprise Admin or an ORG Admin.

10-39 Object-Based RBAC in a Multitenancy Environment

In NSX 4.0.1, the following new user roles are available for RBAC:

- ORG Admin:
 - Creates Projects
 - Provides role binding between Project Admin and Project object
 - Performs CRUD operations in the allocated ORG object
- Project Admin:
 - Performs CRUD operations in the allocated Project object

Object-based RBAC is only configurable through NSX API in NSX 4.0.1



Role binding is used to bind a user role to an object. The user can perform CRUD operations in an object based on its role binding.

Role binding is hierarchical in nature. Users at any level can see all the role bindings for themselves and the users below them.

Users at any level can provide role bindings for the users below them. For example, an Enterprise Admin can provide role bindings between an ORG Admin user and an ORG and between a Project Admin user and a Project.


An Enterprise Admin can perform CRUD operations across the entire NSX infrastructure.

Object-based RBAC is only configurable through NSX API in NSX 4.0.1. NSX UI support is expected in future releases.

10-40 Role Assignment

You can add, change, and delete role assignments for users or user groups:

- Select System > Settings > User Management > User Role Assignment.
- Click ADD ROLE FOR PROVIDERS and select a user type.
 - LDAP
 - VMware Identity Manager
- Select the search domain.
- Search for the users or user groups that you want to assign the roles to.
- Select a role or roles and click SAVE.



10-41 Lab 20: Managing Users and Roles

Integrate NSX Manager with Active Directory over LDAP:

- Prepare for the Lab
- Add an Active Directory Domain as an Identity Source
- Assign NSX Roles to Domain Users and Test Permissions
- Modify an Existing Role and Test the Role Permissions

10-42 Review of Learner Objectives

- Identify the different types of users in NSX.
- Recognize permissions and roles available in NSX
- Create and configure custom roles
- Explain object-based RBAC in a multitenancy environment
- Assign roles to users

10-43 Key Points

- VMware Identity Manager provides application provisioning, conditional access controls, and SSO for SaaS, web, cloud, and native mobile applications.
- You can integrate NSX with VMware Identity Manager and configure RBAC for users that VMware Identity Manager manages.
- You can add Active Directory over LDAP or OpenLDAP identity sources to NSX and configure RBAC for these users.
- Role-based access control (RBAC) enables you to restrict system access to users based on their role in the company.
- The Custom RBAC feature in NSX enables administrators to create custom roles in addition to the existing built-in roles.

Questions?

Module 11

NSX Federation

11-2 Importance

In NSX, Federation enables network administrators to define global configuration settings and policies that span multiple sites. You must understand the Federation architecture and configuration, logical switching, logical routing, and security to successfully configure NSX Federation in your environment.

11-3 Module Lessons

- Federation Architecture
- Installing and Onboarding Federation
- Federation Networking
- 4. Federation Security


11-4 Lesson 1: Federation Architecture

11-5 Learner Objectives

- Describe Federation and its use cases
- Describe the requirements and limitations of Federation
- Describe the Federation configuration workflow

About NSX Federation 11-6

NSX Federation offers globally managed network and security services across multiple locations, improving scale and design flexibility. NSX supports up to eight federated locations.



(Same Global Manager)

NSX Federation supports multilocation deployments and offers:

- A single interface for managing networking and security
- Improved network agility and business continuity
- Consistent policy configuration and enforcement
- Simplified disaster recovery and avoidance

Having multiple federated sites brings the following benefits:

- High availability of applications
- Better application response time
- Cost-effective hosting solution based on the application criticality
- Configuration during mergers or acquisitions

11-7 NSX Federation Components: Global Manager

The NSX Federation architecture has two main components: Global Manager and Local Manager.

A Global Manager appliance provides the GUI and the REST APIs for configuring objects across geographical sites. Global Manager has the following characteristics:

- Global configurations are replicated to the standby Global Manager appliance.
- Global configurations are sent to the relevant Local Manager appliances.
- Configurations that require protection against site failures must be configured in Global Manager.
- During site onboarding, you can move the existing Local Manager configuration to Global Manager.





11-8 NSX Federation Components: Local Manager

Local Manager provides management for a single site. UI and API access are available through the Local Manager virtual IP address.

Local Manager performs the following functions:

- Realizes global configurations that are sent by Global Manager
- Continues to accept user configurations pertaining to the local site
- Owns the Infrastructure preparation and configuration, for example, transport node preparation, edge configurations, transport zones, IP address pools, and so on






11-9 NSX Federation Components: Global Manager and Local Manager Clusters

Global Manager and Local Manager have the following characteristics:

- Global Manager and Local Manager appliances are deployed in a VM form factor.
- The Global Manager node does not include a control plane.
- Global Manager appliances are deployed as three-node clusters that can be in a single location or spread across three locations.
- Local Manager appliances are deployed as three-node clusters in each location for high availability and scalability.
- Local Manager appliances can selectively override some global object configuration parameters.



Global Manager nodes are deployed with an Open Virtualization Format (OVF) template and form a cluster with three nodes.


The Global Manager node includes a management plane. It does not include a control plane.

Local Manager can selectively override some global object configuration. For example, Local Manager can configure site-specific Border Gateway Protocol (BGP) timer configurations on a global Tier-0 gateway.

11-10 Federation Configuration Types

NSX Federation supports the following types of configurations:

- Global configuration:
 - The GM node receives the configuration.
- Local configuration;
 - The LM cluster receives the configuration on a given location.



11-11 Ownership of Logical Configuration (1)

Network objects that are created by GM are owned by GM:

- Tier-0 and Tier-1 gateways, segments, segment profiles, and so on, can be configured from GM.
- GM is the single source of truth for these objects. These objects can only be modified or deleted by GM.

Global Manager


The GM-created objects are visible to LM but are read-only.



11-12 Ownership of Logical Configuration (2)

Network objects that are created by LM are owned by LM:

- Tier-0 and Tier-1 gateways, segments, segment profiles, and so on, can be configured from LM.
- LM is the single source of truth for these objects. These objects can only be modified or deleted by LM.
- The LM-created objects are not visible to GM.
- During the onboarding process, you can move the object to GM. Then, LM loses ownership
 of these objects.



The following configurations can be imported from the Local Manager into Global Manager:


- T0 gateway
- T1 gateway
- Firewall security policies
- Services
- Security profiles
- Context profiles
- Services
- Groups
- NAT
- DHCP
- DNS
- Gateway profiles
- Time-based firewall (supports import/onboard)

For more information about the local manager configurations supported for importing into global manager, see NSX-T Data Center Installation Guide at: https://docs.vmware.com/en/VMware-NSX/4.0/installation/GUID-388CE659-3FE3-4EF4-ABA3-AE3FCAA191E9.html

11-13 Infrastructure Ownership

LM always manages infrastructure objects:


- These objects include transport nodes, edge nodes, transport zones, and so on.
- LM completely owns infrastructure preparation.
- Objects can only be created, modified, or deleted by LM.



11-14 Global Configuration

The global configuration workflow is as follows:


- The user sends the configuration to the active GM.
- The active GM stores locally and replicates the configuration to the standby GM.
- The active GM sends it to the relevant LMs.



11-15 Local Configuration

The local configuration workflow is as follows:


- A user can also send the configuration to LM at each location, as required.
- Each LM stores the configuration locally.
- No configuration is sent to GM.



11-16 Federation Configuration Example

In the configuration example, networks are created and realized only in two locations (Location A and Location B).

TO, T1, and their associated segment are stretched to Location A and Location B.




Blue Segment associated with T1 is also stretched to Location A and Location B.

Two VMs (VM1 and VM2) are connected to Blue_Segment.

This configuration is realized in GM and LMs across locations.

11-17 Federation Configuration Example Workflow (1)

As the T0, T1, and their associated segment are only stretched to Locations A and B, the configuration is not sent to LM at Location C.




In the example, the user wants to apply a configuration that is stretched to Location A and Location B only:

- The user sends the configuration by using REST over HTTPS to the active GM node in Location A, which stores the configuration locally.
- The active GM replicates the configuration to the standby GM.
- GM pushes the configuration to LM clusters of Location A and Location B, because the configuration is intended only for them. The configuration is not sent to LM at Location C.

11-18 Federation Configuration Example Workflow (2)

LMs of each location initiate a sync operation to exchange the Control Plane configuration of a given topology.



In the example, the control plane of LM in each location learns the IP and MAC information of the VMs associated with Blue_Segment:

- LMs at each location initiate a sync operation to exchange the IP and MAC information of VMs associated with Blue_Segment.
- Each LM stores the realization information of a given topology discovered through the sync operation. As a result, the Control Plane of LM in each location has the IP and MAC information of all the VMs associated with Blue_Segment across Location A and Location B.

LM does not update control plane information to GM, because GM has no control plane. However, the GM UI enables you to see the inventory of all locations by querying the relevant LMs.

11-19 Review of Learner Objectives

- Describe Federation and its use cases
- Describe the requirements and limitations of Federation
- Describe the Federation configuration workflow

11-20 Lesson 2: Installing and Onboarding Federation

11-21 Learner Objectives

- Describe the prerequisites for Federation
- Configure GM as active and standby in the primary and secondary location
- Describe location onboarding

11-22 NSX Federation: Prerequisites

NSX Federation requires several prerequisites.

Location-to-Location Latency, Stretched Networking Use Case	Up to 150 ms
Location-to-Location Latency, Stretched Security Use Case	Up to 500 ms
WAN Bandwidth	No congestion for the management plane or the data plane
WAN MTU	1,700 bytes
Data Plane	For different Internet Service Providers (ISPs), a public address must be advertised from both locations
Compatibility	All appliances in an NSX Federation environment must have the same version installed.

Location-to-location traffic must satisfy the following characteristics:

- A maximum 150 ms Latency (RTT) between locations in the stretched networking use case, where networking objects, such as Tier-0 gateways, Tier-1 gateways, or segments span multiple sites.
- A maximum 500 ms Latency (RTT) between locations in the stretched security use case, supported since NSX-T Data Center 3.2.1. If present, networking objects do not span multiple sites.

- Allow IP and firewall connectivity across sites:
 - Allow Management traffic between GM and LM.
 - Allow Data plane traffic between edge nodes (RTEP).
 - Do not configure NAT for Management and RTEP networks.
- WAN bandwidth requirement:
 - No congestion for the Management plane (GM to LM traffic).
 - No congestion for the data plane.
- WAN MTU requirement:
 - MTU of 1,700 bytes or more to avoid the edge node RTEP traffic fragmentation.
- Version requirement:
 - All appliances in an NSX Federation environment must have the same version installed.
- Global Manager supports only Policy Mode. NSX Federation does not support Manager Mode.

11-23 Onboarding Process

The onboarding process involves infrastructure onboarding and location onboarding.

The onboarding process includes several steps.



The infrastructure onboarding process is repeated for each location until all locations are on board.


RTEP configuration is required for stretched networking functionality.

Location onboarding promotes existing objects on a local manager into the global manager configuration.

The onboarding process does not affect running workloads.

11-24 Active Global Manager Configuration






The **Location Manager** tab is the central management point for configuring GMs (active and standby).

You can add different locations from this tab to GM. The managers in these locations are Local Manager (LM).

You can configure the standby GM from the active GM site to provide flexibility for configuring the Federation component from a single console.

11-25 Adding Standby Global Manager (1)


You must add a standby GM, If the active GM falls, the standby GM can be activated from its UI.



You can configure the standby GM from the active GM site to provide flexibility for configuring the Federation component from a single console.

11-26 Adding Standby Global Manager (2)

You provide the Location B GM information. The certificate thumbprint of the GM is required.




Site-B-GM> get certificate api thumbprint f7c7ld32ab32d4e8371416d305ccla5a5d4df7792b7b0cb236ab50da2f8f1bad

The compatibility check is mandatory because it checks version compatibility between the two GMs.

From the CLI, run get certificate api thumbprint to fill the details of the SHA-256 thumbprint in the UI.

11-27 Adding a Location

You add the location from ADD ON-PREM LOCATION, and provide the details about NSX Manager at Location A.




The SHA-256 certificate thumbprint is required.

You can obtain the certificate thumbprint by using the following command on NSX Manager: get certificate api thumbprint

11-28 Validating the Local Manager


The Location Manager tab on LM provides information about the Sync status of LM to GM.



The **Location Manager** tab also provides details about the active and standby GM. The details are read-only.

11-29 Location Onboarding

You can optionally onboard a location by promoting LM objects to the GM.



LM objects can be imported either during onboarding or after onboarding. The screenshot shows an option to import after your location is onboarded.

If the default transport zone is not found for the Local Manager at site Site-A-LM, then wait for default transport zone discovery or reload the enforcement point and retry.

A mandatory check verifies that the LM is backed up and is a prerequisite to onboarding a location.

The import takes time. During this time, no downtime occurs.

You can use this time to prepare the Networking and Security workloads and configuration.


import of configurations into the Global Manager is blocked if you have any of the following configurations in your Local Manager. You must remove unsupported configurations to proceed with importing. After your supported Local Manager configurations are successfully imported into Global Manager, you can add the configurations for any of the unsupported features back into your Local Manager.

The following Local Manager configurations are not supported for importing into Global Manager:

- DHCP dynamic binding
- Distributed IDS
- Distributed security for vCenter VDS Port Groups
- Endpoint protection
- Forwarding policies
- Guest Introspection
- Identity firewall
- IDS/IPS
- L2 Bridge
- Load balancer
- Malware prevention
- Metadata proxy
- Multicast
- Network detection and response
- Network introspection
- Routing protocols (OSPF)
- Routing VPN and EVPN
- Service Insertion
- TO VRF
- TLS inspection
- URL filtering

11-30 Onboarding Preparation







When LM objects are imported to GM, you must mark them to differentiate these objects.

To distinguish better, you add either a prefix or a suffix. You can define the prefix or suffix.

In this example, Site B onboarded objects are prepended with the text Site-B-LM- for location tracking.

11-31 Verifying Onboarding





In this example, onboarded objects, Site-B-LM-siteb-t0gw-01 and Global-IP-address-Group, are now global.

11-32 Review of Learner Objectives

- Describe the prerequisites for Federation
- Configure GM as active and standby in the primary and secondary location
- Describe location onboarding

11-33 Lesson 3: Federation Networking

11-34 Learner Objectives

- Describe the stretched networking concepts in Federation
- Explain the supported Tier-0 and Tier-1 stretched topologies
- Explain Layer 2 concepts related to NSX Federation
- Identify network types that can be interconnected using L2 Bridging in federated environments

11-35 Stretched Networking (1)

NSX Federation supports the stretching of the following networking constructs across locations:

- Tier-0 gateways (T0)
- Tier-1 gateways (T1)
- Segments

The following services are supported on stretched T1 gateways:

- NAT
- Gateway firewall
- IPv6
- DHCP
- DNS





The Global Manager UI enables the user to create stretched networks, for example, segments, gateways and security policies, and rules.


11-36 Stretched Networking (2)

GM can stretch the Tier-0 and Tier-1 gateways across locations:

 All segments that are connected to the downlink port of the stretched T1 gateway are automatically stretched across locations.

The edge nodes are used to forward cross-location traffic:

- This method eliminates the need for tunnels between hypervisors across locations.
- The edge nodes use RTEPs to provide cross-location communication.



Remote Tunnel Endpoints (RTEPs) are created from the LM UI.

On the edge nodes, RTEPs are configured to forward the traffic across sites.


NSX does not encrypt the traffic, but the user can encrypt this cross-location traffic with a thirdparty tool.

11-37 Tier-0 and Tier-1 Gateways: Logical Topologies (1)

GM supports stretch networks for cross-location communication and nonstretched networks for a local location.

A GM-supported configuration has the following features:


- Tier-0 and Tier-1 gateways can still be local to Location Manager.
- Tier-0 and Tier-1 gateways can stretch across locations.
- Segments connected to the stretched Tier-O/Tier-1 gateways are automatically stretched across locations.
- Segments connected to nonstretched Tier-0/Tier-1 gateways are local to the location.



11-38 Tier-0 and Tier-1 Gateways: Logical Topologies (2)

GM-supported configuration:

- Tier-O and Tier-1 gateways can be stretched to all or some of the locations.
- Segments associated with stretched for the Tier-O/Tier-1 gateway are also stretched to the same span:
 - The span of a segment is equal to the span of the Tier-0/Tier-1 gateway.
- If the scope of the Tier-1 gateway is equal to or is a subset of the Tier-0 gateway, a stretched Tier-0 gateway can connect to a nonstretched Tier-1 gateway.
- For a Tier-1 gateway without services, the span of Tier-1 is equal to the span of the Tier-0 gateway.



The example includes the stretched network objects, including the stretched segments and stretched Tier-0 and Tier-1 gateways.


Spans have the following properties:

- The span of the segment is equal to the span of the Tier-0 and Tier-1 gateways.
- The span of the Tier-0 gateway is equal to the span of the Tier-1-no services gateway, or conversely.
- The span of the Tier-0 gateway is equal to or a subset of the span of the Tier-1 gateway.

11-39 Tier-0 and Tier-1 Gateways: Logical Topologies (3)

In the example, the span of the TO-Stretched and both T1-Stretched gateways are not the same. These connections are not possible.

However, the span of T1-Not Stretched is a subset of the T0-Stretched span. The connection is possible.



The example shows how a stretched and nonstretched network can co-exist.


A nonstretched Tier-1 gateway is local to Location A only and Tier-0 gateway is stretched across Location A and Location B. The connection is possible because the nonstretched Tier-1 gateway is a subset of the Tier-0 Stretched gateway.


A connection is not possible between T1-Stretched-1 and T1-Stretched-2 to T0-Stretched because T0-Stretched is not stretched to Location C.

11-40 Single-Location Tier-0 Gateway Deployments

GM supports the following configuration for Tier-0 gateways for each location:

- Tier-0 is deployed in A/A mode in a location:
 - All Tier-0 gateways are active on all the NSX Edge nodes in the edge cluster.
 - The traffic ingresses or egresses from all the edge nodes with the active Tier-0 gateway.
- Tier-0 is deployed in A/S mode in a location:
 - One Tier-0 gateway is active on one NSX Edge node and one standby gateway exists in the edge cluster.
 - The traffic Ingresses or egresses from the edge nodes with the active Tier-0 gateway.






11-41 Single-Location Tier-1 Gateway Deployments

GM supports the configuration for the Tier-1 gateway for each location.

Tier-1 is deployed in A/S mode in a location:


- The Tier-1 gateway is active on one NSX Edge node and is standby on another node.
- The traffic ingresses or egresses from the edge node deployed with the active Tier-0 gateway.



11-42 Multilocation Tier-0 and Tier-1 Gateway Deployments (1)

A location is configured as either primary or secondary:

- Only one location can be configured as primary, and all other locations are secondary.
- TO/T1 can be deployed in primary and secondary (P/S) locations:
 - Traffic to T0 from T1 and segments is contained in the location.
 - One location is active to send northbound egress traffic for a destination.




Location A is primary, and all northbound egress traffic is routed by using the primary Tier-0 gateways edge node.

The Tier-1 gateways at Location B and Location C send the traffic through RTEPs to the Location A edge uplink.

11-43 Multilocation Tier-0 and Tier-1 Gateway Deployments (2)

TO can be deployed in all primary locations (All_P):

- Traffic to Tier-0 gateway from Tier-1 gateway and segments is contained in the location.
- All locations have an active Tier-O gateway to send northbound egress traffic.




For all primary locations, the northbound egress traffic is routed locally.

11-44 Multilocation TO-Stretched Gateway Modes (1)

Tier-0 is deployed in A/A mode in P/S locations:

- For the Tier-0 gateway configured in the active-active HA mode, the Tier-0 gateway will be active on all the NSX Edge nodes in the edge cluster.
- For the Tier-0 gateways configured in the primary-secondary mode, the north-south traffic from all locations is forwarded to the Tier-0 gateways configured in a primary location.
- The Tier-O gateway in the active-active high availability mode does not support stateful NAT. However, stateless NAT can be used.




In the diagram:

- Location A is primary, and Location B is secondary.
- Tier-0 gateways are deployed in active-active mode for both locations.
- RTEP IPs are configured on EN1, EN2, EN3, and EN4.
- EN3 and EN4 at Location B egress northbound traffic to either EN1 or EN2, or both, on Location A.

11-45 Multilocation TO-Stretched Gateway Modes (2)

Tier-0 is deployed in A/A mode in all primary locations:

- For the Tier-0 gateway configured in active-active HA mode, the Tier-0 gateway will be active on all the NSX Edge nodes in the edge cluster.
- For the Tier-0 gateway configured in a primary setup, northbound and southbound are sent and received through their respective Tier-0 gateways in their location.
- This configuration is also called A/A Local Egress.
- Tier-0 does not support services in this deployment mode:
 - Stateless NAT can be used.




In the diagram:

- Both locations are primary.
- Tier-0 gateways are deployed in active-active mode for both locations.
- In this use case for local egress, all Tier-O sites route the northbound traffic locally.

11-46 Multilocation T1-Stretched Gateway Modes (1)

Tier-1 is deployed without services:


- Tier-1 is deployed without services in all locations.
- Only Distributed Router (DR) is realized because services are not configured and Service Router (SR) does not exist.
- Tier-1 routes the northbound traffic directly to TO-Stretched on the hypervisor.
- For workloads connected to T1-Stretched across locations, the communication happens through edge nodes (RTEP) configured for L2 Stretch.
- The T1-no-Services gateway does not require edge nodes for routing, but edge is required for L2 stretching.



11-47 Multilocation T1-Stretched Gateway Modes (2)

Tier-1 is deployed in A/S mode in P/S locations:

- The edge nodes in each location must support Tier-1.
- Services are enabled on Tier-1 in this deployment mode.
- Active edge nodes in both primary and secondary locations can receive southbound traffic for their respective locations.
- Only the primary location's active edge node can send the northbound traffic of both locations to stretched Tier-O.




The concept is similar to active-standby Tier-0, where the active node on the primary site routes the traffic.

RTEP configured on Tier-0 routes the traffic across locations.

11-48 About RTEP

RTEP is an edge node IP that is used for edge node communication across sites:

- RTEPs use Geneve encapsulation to communicate.
- Only one RTEP IP can be configured per edge node.
- The RTEP configuration is done from the LM.



To configure RTEP:

- Click Site-A-GM from UI Selector.
- 2. Click the System tab to display the configuration for Location Manager.
- Click Location Manager to display details about the location where you want to create the RTEP.
- On Site-A-LM, click Networking.

You can select the LM edge node cluster where you need to configure RTEP.

Click CONFIGURE

You can navigate to Local Manager for the selected site.


11-49 Stretched Layer-2 Network

The cross-location layer-2 communication is provided by the edge nodes of an edge cluster in each location.

This method avoids the management of many tunnels and BFD sessions between all hosts across locations.

The L2 stretched communication workflow includes:

- The source ESXI host sends frames to the edge node (TEP-TEP communication).
- The source edge node forwards a frame to the destination edge node (RTEP-RTEP communication).
- The destination edge node forwards a frame to the destination ESXI host (TEP-TEP communication).




The MAC address of a remote VM is learned through RTEP on the edge nodes. The edge node passes this information to its local transport node.

11-50 Stretched L2: VNI Mapping

When a stretched segment is created from GM:

- GM requests each LM to create a local segment:
 - The VNI selection is local to LM and can be different VNIs in different locations.
- GM selects the VNI for RTEP.
- LM does the TEP-VNI:RTEP-VNI mapping:
 - This mapping is only available in edge nodes.



Example: If you create a stretched segment from GM and it is stretched to two locations, then the VNI id can be 5001 on one location. On another location, the stretched segment VNI id can be 6003.

VNI can be different, but UUID of the stretched segment will be same across all locations.


11-51 About Federation L2 Bridging

NSX-T Data Center 3.2.2 supports L2 Bridging across locations.

L2 Bridging can be used to interconnect:

- A VLAN-backed network to an overlay segment at the same location
- 2. A VLAN-backed network to a stretched overlay segment across locations
- A VLAN-backed network to a VLAN-backed network, over a stretched overlay segment, across locations

NSX Federated Environment




A stretched Tier-0 or stretched Tier-1 gateway is used to interconnect stretched overlay segments across locations.

11-52 Components of Federation L2 Bridging

Federation L2 Bridge components include:


- RTEP (R): An edge node interface used for edge node communication across sites
- Bridge: A connecting point between an overlay segment and a VLAN-backed network
- L2forwarder (L2): An edge node that is responsible for forwarding bridged traffic
- Edge Bridge Profile: Contains the L2 Bridge configuration details, including location, Edge cluster, primary node, and secondary node.



11-53 Federation L2 Bridge Communication

Federation L2 Bridge Communication is as follows:

- When VM1's MAC address is learned by the Bridge, Edge 2 will forward that MAC address to Edge 3 through the RTEP (R).
- The communication across locations, occurs between RTEP (R) Interfaces on the Active Edge Node 1 and Active Edge Node 3.



11-54 Federation Traceflow

Traceflow analysis is supported across multiple locations, is initiated from the GM, and can assist with network troubleshooting in federated deployments:



Traceflow identifies the path that a packet takes to reach its destination or, conversely, where a packet is dropped along the way. Each entity reports the packet handling on input and output, to determine whether issues occur when receiving a packet or when forwarding the packet.

In this Inter-location Traceflow example:

- Traceflow is from source location Site A, VM sa-web-01 to destination Site B, VM sb-web-01
- Traffic between sites is Geneve encapsulated between RTEP interfaces

11-55 Review of Learner Objectives

- Describe the stretched networking concepts in Federation
- Explain the supported Tier-0 and Tier-1 stretched topologies
- Explain Layer 2 concepts related to NSX Federation
- Identify network types that can be interconnected using L2 Bridging in federated environments

11-56 Lesson 4: Federation Security

11-57 Learner Objectives

- · Explain the Federation security use cases
- Explain the security configuration workflows
- Describe the Federation security components


11-58 Stretched Security in NSX Federation

The guiding security principles for NSX also apply to NSX Federation, except that the boundaries are extended to support multiple locations.

NSX Federation supports stretching of the following security constructs across multiple locations to centralize and simplify security policies:

- Distributed firewall
- Gateway firewall
- FQDN filtering
- L7 App ID context support





The following security features are not supported:

- Identity Firewall
- NSX Distributed IDS/IPS
- Network Introspection
- Endpoint protection
- Malware prevention
- Network detection and response

11-59 Use Cases for Security

Federation security has the following use cases:

- Provide consistent policy across deployments managed using NSX Federation
- Effective disaster recovery ensuring continuity of established security framework
- Extension of network and security framework to another location for increased capacity




Federation provides central security for data centers managed on premises, such as data center extension with a centralized pane.

11-60 Security Configuration Workflow

The security configuration workflow includes the following steps:

- The user sends the configuration to Global Manager (GM).
- GM sends the configuration to Local Manager (LM) for each location.
- LM forwards the configuration to MP/CCP.
- The central control planes (CCPs) at each location synchronize the configuration across locations.
- 5. Each CCP sends the consolidated configuration to the data plane.



11-61 About Regions

A Region is a collection of locations, which is used to create focused groups for security and networking policies.

Regions can be global, regional, or local.

Some regions are created automatically after the onboarding process in GM.

You can create more regions.



Global, Site-A-LM, and Site-B-LM are created by the system while registering and onboarding to GM of LM. Up to eight locations can appear on this page based on the number of supported locations in NSX Federation.

Region A+B is created manually. In the use case, regional groups were added based on the regions. Locations A and B were manually added.

11-62 About Groups

The NSX objects can be grouped to use in firewall rules that apply to Global, Regional, or Local regions.



Groups are defined based on the region. They can be Global, Regional, or Local.

The Global groups span across all sites in Federation and have the following characteristics:


- Created from the Global Manager NSX UI.
- Global groups are owned by NSX and are created with discovered objects. They can also be imported from a file (IP or MAC list).
- Users can only create, update, and delete Global Groups from the GM.
- Both GM and LM can read Global groups.
- Global groups cannot be used as nested members in LM-based policies.

The Regional groups span across two or more locations, for example, Location A+B.

The Local groups are specific to the site, for example, Region-A-Web-Servers, Region-B-Web-Servers, and Region-A-App-Servers.

11-63 GM Groups and Span Example (1)

In this example, GM groups were created with various location spans.




In the example, groups are as follows:

- Group1: Global, stretched to all sites, Locations A, B, and C
- Group2: Regional, stretched to Locations A and B
- Group3: Local to Location A
- Group4: Local to Location B

The GM sends groups to each LM within the group's span.

11-64 GM Groups and Span Example (2)

The span of a group must align with the span of the objects to be included in that group.




In the example, Group1 was created using Segment S1 as its membership:

- The span of Group1 is Location A and Location B.
- Segment S1 spans Location A, B, and C.
- The span of the group does not align with the span of the required objects, resulting in a security exposure.
- The required security policy is not applied to VM3.

11-65 Group Membership Criteria

You dynamically assign members to a group based on one or more criteria




A criterion can have one or more conditions:

- NSX uses the logical AND operator after each condition within a criterion.
- Multiple criteria can be combined where OR is the default logical operator between conditions, and where an option exists to use the logical AND operator.

In this example, VMs with a computer name that starts with web are dynamically added to the Global-Web-Servers global group.

11-66 Group Membership Based on the VM Tag (1)


Group membership criteria can include VM-based tags.



In the example, the membership criteria for the group, Global-Web-Servers, is based on the VM tag web.

11-67 Group Membership Based on the VM Tag (2)

Tags are applied to VMs even after migration with vSphere vMotion to Location B, which allows them to maintain group membership.




In the example, the Global-Web-Servers group members are maintained when VM1 is migrated using vSphere vMotion to Location B.

11-68 About GM-Based Policy

A GM-based policy is a collection of one or more firewall rules.

You can configure the following settings for a policy:

- Applied To in the policy is set to DFW:
 - All logical switch ports (VMs and containers) of the span receive the rules in that section.
- Applied To in the policy is set to Group:
 - All group members (VMs and containers) receive the rules in that section.



GM can create FW in all categories except Ethernet and Emergency.

You can create a section and apply it to a location, groups, and so on.

When **Applied To** in the policy is set to **Group**, this setting overrides the **Applied To** for the DFW rule.

11-69 About GM-Based Rules

GM-based rules enforce traffic based on objects and tags.


You can configure the following settings for a rule:

- Applied To in the rule is set to DFW:
 - All logical switch ports (VMs and containers) of the span receive the rules within that section.
- Applied To in the rule is set to Group:
 - All group members (VMs and containers) receive the rules within that section.



11-70 Overlap of GM and LM Sections

GM and LM can create gateway firewall and distributed firewall sections in the same category. For GM and LM sections created in the same category, GM sections are always at the top.




Except the Ethernet and Emergency categories where only LM can create sections:

- GM and LM sections are individually created under the Tier-0 or Tier-1 gateway firewall and under the distributed firewall.
- GM rules are always above LM.

11-71 GM DFW Management Enhancement


From NSX 4.0.1, DFW can be centrally enabled and disabled on all registered LMs from the GM.



You cannot override the global DFW configuration setting from a LM.

11-72 Global DFW Exclusion List Enhancement

From NSX 4.0.1, you can globally exclude a set of virtual machines from distributed firewall protection centrally from the GM.




A DFW exclusion list excludes a set of virtual machines from distributed firewall protection.

Two exclusion lists exist on the LM. One list is from the GM and the other list is its own locally defined exclusion list. All excluded members of these two lists are sent to the Management Plane as a unified exclusion list, to be applied to the LM.

11-73 Time-Based DFW Rules Enhancement

From NSX 4.0.1, you can globally define time-based DFW rules, centrally from the GM, which are pushed to all registered LMs.



With time windows, security administrators can restrict traffic from a source or to a destination, for a specific time period.

11-74 Review of Learner Objectives

- Explain the Federation security use cases
- Explain the security configuration workflows
- Describe the Federation security components

11-75 Key Points

- Federation provides consistent policy and operational simplicity with multisite functionality for data centers.
- GM performs onboarding for multiple sites.
- GM creates stretched networks by using Tier-0 and Tier-1 routers.
- The edge node delivers cross-location communication to avoid connecting hypervisors across sites.
- The guiding security principles for NSX also apply to NSX Federation, except that the boundaries are extended to support multiple locations.

Questions?