

410.4

Workstations and Servers

SANS

Copyright © 2019, Justin Searle. All rights reserved to Justin Searle and/or SANS Institute.

PLEASE READ THE TERMS AND CONDITIONS OF THIS COURSEWARE LICENSE AGREEMENT ("CLA") CAREFULLY BEFORE USING ANY OF THE COURSEWARE ASSOCIATED WITH THE SANS COURSE. THIS IS A LEGAL AND ENFORCEABLE CONTRACT BETWEEN YOU (THE "USER") AND THE SANS INSTITUTE FOR THE COURSEWARE. YOU AGREE THAT THIS AGREEMENT IS ENFORCEABLE LIKE ANY WRITTEN NEGOTIATED AGREEMENT SIGNED BY YOU.

With the CLA, the SANS Institute hereby grants User a personal, non-exclusive license to use the Courseware subject to the terms of this agreement. Courseware includes all printed materials, including course books and lab workbooks, as well as any digital or other media, virtual machines, and/or data sets distributed by the SANS Institute to the User for use in the SANS class associated with the Courseware. User agrees that the CLA is the complete and exclusive statement of agreement between The SANS Institute and you and that this CLA supersedes any oral or written proposal, agreement or other communication relating to the subject matter of this CLA.

BY ACCEPTING THIS COURSEWARE, YOU AGREE TO BE BOUND BY THE TERMS OF THIS CLA. BY ACCEPTING THIS SOFTWARE, YOU AGREE THAT ANY BREACH OF THE TERMS OF THIS CLA MAY CAUSE IRREPARABLE HARM AND SIGNIFICANT INJURY TO THE SANS INSTITUTE, AND THAT THE SANS INSTITUTE MAY ENFORCE THESE PROVISIONS BY INJUNCTION (WITHOUT THE NECESSITY OF POSTING BOND), SPECIFIC PERFORMANCE, OR OTHER EQUITABLE RELIEF.

If you do not agree, you may return the Courseware to the SANS Institute for a full refund, if applicable.

User may not copy, reproduce, re-publish, distribute, display, modify or create derivative works based upon all or any portion of the Courseware, in any medium whether printed, electronic or otherwise, for any purpose, without the express prior written consent of the SANS Institute. Additionally, User may not sell, rent, lease, trade, or otherwise transfer the Courseware in any way, shape, or form without the express written consent of the SANS Institute.

If any provision of this CLA is declared unenforceable in any jurisdiction, then such provision shall be deemed to be severable from this CLA and shall not affect the remainder thereof. An amendment or addendum to this CLA may accompany this courseware.

SANS acknowledges that any and all software and/or tools, graphics, images, tables, charts or graphs presented in this courseware are the sole property of their respective trademark/registered/copyright owners, including:

AirDrop, AirPort, AirPort Time Capsule, Apple, Apple Remote Desktop, Apple TV, App Nap, Back to My Mac, Boot Camp, Cocoa, FaceTime, FileVault, Finder, FireWire, FireWire logo, iCal, iChat, iLife, iMac, iMessage, iPad, iPad Air, iPad Mini, iPhone, iPhoto, iPod, iPod classic, iPod shuffle, iPod nano, iPod touch, iTunes, iTunes logo, iWork, Keychain, Keynote, Mac, Mac Logo, MacBook, MacBook Air, MacBook Pro, Macintosh, Mac OS, Mac Pro, Numbers, OS X, Pages, Passbook, Retina, Safari, Siri, Spaces, Spotlight, There's an app for that, Time Capsule, Time Machine, Touch ID, Xcode, Xserve, App Store, and iCloud are registered trademarks of Apple Inc.

PMP and PMBOK are registered marks of PMI.

SOF-ELK® is a registered trademark of Lewes Technology Consulting, LLC. Used with permission.

SIFT® is a registered trademark of Harbingers, LLC. Used with permission.

Governing Law: This Agreement shall be governed by the laws of the State of Maryland, USA.



Workstations and Servers

© 2019 Justin Searle | All Rights Reserved | Version E01_01

The **SANS ICS410**, ICS/SCADA Security Essentials course, was developed by a collection of experts whose diverse work experiences, knowledge, and skills truly blend together to cover the very specific content areas for this course.

Justin Searle is the Director of ICS Security at InGuardians, specializing in ICS security architecture design and penetration testing. Justin led the Smart Grid Security Architecture group in the creation of NIST Interagency Report 7628 and has played key roles in the Advanced Security Acceleration Project for the Smart Grid (ASAP-SG), National Electric Sector Cybersecurity Organization Resources (NESCOR), and Smart Grid Interoperability Panel (SGIP). Justin has taught courses in hacking techniques, forensics, networking, and intrusion detection for multiple universities, corporations, and security conferences. He is currently a Senior Instructor for the SANS Institute. In addition to electric power industry conferences, Justin frequently presents at top international security conferences such as Black Hat, DEFCON, OWASP, Nullcon, and AusCERT. Justin co-leads prominent open source projects including The Control Thing Platform, Samurai Web Testing Framework (SamuraiWTF), Samurai Security Testing Framework for Utilities (SamuraiSTFU), Yokoso!, and Laudanum. Justin has an MBA in International Technology and is a CISSP and SANS GIAC certified Incident Handler (GCIH), Intrusion Analyst (GCIA), Web Application Penetration Tester (GWAPT), and GIAC Industrial Control Security Professional (GICSP).

Dr. Eric Cole is an industry-recognized security expert with over 20 years of hands-on experience. Dr. Cole has experience in information technology with a focus on helping customers focus on the right areas of security by building out a dynamic defense. Dr. Cole has a master's degree in computer science from NYIT and a doctorate from Pace University with a concentration in information security. He served as CTO of McAfee and Chief Scientist for Lockheed Martin. Dr. Cole is the author of several books, including *Advanced Persistent Threat*, *Hackers Beware*, *Hiding in Plain Sight*, *Network Security Bible*, 2nd Edition, and *Insider Threat*. Eric is the inventor of over 20 patents and is a researcher, writer, and speaker. He is also a member of the Commission on Cyber Security for the 44th President and several executive advisory boards. Dr. Cole is the founder and an executive leader at Secure Anchor Consulting, where he provides leading-edge cybersecurity consulting services, expert witness work, and leads research and development initiatives to advance the state-of-the-art in information systems security. Dr. Cole is actively involved with the SANS Technology Institute (STI). He is a SANS faculty Fellow who works with students, teachers, and develops and maintains courseware.

Eric Cornelius is the Director of Critical Infrastructure and Industrial Control Systems (ICS) at Cylance, Inc. He is responsible for the thought leadership, architecture, and consulting implementations for the company. His leadership keeps organizations safe, secure, and resilient against advanced attackers. Previously, Eric served as the Deputy Director and Chief Technical Analyst for the Control Systems Security Program at the US Department of Homeland Security. As an active researcher in the field of cybersecurity since 2002, Eric supported many "boots-on-the-ground" engagements involving penetration testing, forensics, and malware analysis. Through these engagements, he aided multiple government, military, and private sector organizations in protecting their networks and industrial control systems. In addition to his years of technical leadership, Eric literally wrote the book on incident response in the ICS arena. Eric's extensive knowledge of critical infrastructure and those who attack it will be brought to bear at Cylance as he leads a team of experts in securing America's critical systems.

Contributing Authors

Michael Assante is currently the SANS project lead for Industrial Control System (ICS) and Supervisory Control and Data Acquisition (SCADA) security. He served as Vice President and Chief Security Officer of the North American Electric Reliability Corporation (NERC), where he oversaw industry-wide implementation of cybersecurity standards across the continent. Prior to joining NERC, Michael held a number of high-level positions at Idaho National Labs and he served as Vice President and Chief Security Officer for American Electric Power. His work in ICS security has been widely recognized and he was selected by his peers as the winner of *Information Security Magazine's* security leadership award for his efforts as a strategic thinker. The RSA 2005 Conference awarded him its outstanding achievement award in the practice of security within an organization. He has testified before the US Senate and House and was an initial member of the Commission on Cyber Security for the 44th Presidency. Prior to his career in security, Michael served in various naval intelligence and information warfare roles and he developed and gave presentations on the latest technology and security threats to the Chairman of the Joint Chiefs of Staff, Director of the National Security Agency, and other leading government officials. In 1997, he was honored as a Naval Intelligence Officer of the Year.

Tim Conway is currently the Technical Director of ICS and SCADA programs at SANS. He is responsible for developing, reviewing, and implementing technical components of the SANS ICS and SCADA product offerings. He was formerly the Director of CIP Compliance and Operations Technology at Northern Indiana Public Service Company (NIPSCO) where he was responsible for Operations Technology, NERC CIP Compliance, and the NERC training environments for the operations departments within NIPSCO Electric. Tim was previously an EMS Computer Systems Engineer at NIPSCO for eight years, with responsibility over the control system servers and the supporting network infrastructure. He previously served as the Chair of the RFC CIPC, is the current Chair of the NERC CIP Interpretation Drafting Team, a current member of the NESCO advisory board, the current Chair of the NERC CIPC GridEx 2013 Working Group, and the current Chair of the NBISE Smart Grid Cyber Security panel.

TABLE OF CONTENTS**PAGE**

Patching ICS Systems	4
Defending Microsoft Windows	21
EXERCISE 4.1: Baselining with PowerShell	43
Defending Unix and Linux	51
Endpoint Security Software	69
EXERCISE 4.2: Configuring Host-Based Firewalls	87
Event Logging and Analysis	95
EXERCISE 4.3: Windows Event Logs	110
Honeypots	117
Attacks on the Perimeter	131
EXERCISE 4.4: Finding Remote Access	144

This page intentionally left blank.

Course Roadmap

Day 1: ICS Overview

Day 2: Field Devices and Controllers

Day 3: Supervisory Systems

Day 4: Workstations and Servers

Day 5: ICS Security Governance

1. Introduction
2. Patching ICS Systems
 - Patch Decision Tree
 - Vendors, CERTS, and Security Bulletins
3. Defending Microsoft Windows
 - Windows Services
 - Windows Security Policies and GPOs
 - **Exercise 4.1: Baselining with PowerShell**
4. Defending Unix and Linux
 - Differences with Windows
 - Daemons, SystemV, and SystemD
 - Lynix and Bastille
5. Endpoint Security Software
 - Antivirus and Whitelisting
 - Application Sandboxing and Containers
 - **Exercise 4.2: Configuring Host-Based Firewalls**
6. Event Logging and Analysis
 - Windows Event Logs and Audit Policies
 - Syslog and Logrotate
 - **Exercise 4.3: Windows Event Logs**
7. Connectivity to the Internet
 - Honeypots
 - Attacks on the Perimeter
 - **Exercise 4.4: Finding Remote Access**

This page intentionally left blank.

Patching ICS Systems

Applicable Standards:

- **NIST CSF v1.1:** PR.IP-12
- **ISA/IEC 62443-2-3**
- **ISO/IEC 27001:2013:** A.12.6.1, A.14.2.3, A.16.1.3, A.18.2.2, A.18.2.3
- **NIST SP 800-53 Rev. 4:** RA-3, RA-5, SI-2
- **CIS CSC:** 4
- **COBIT 5:** BAI03.10, DSS05.01, DSS05.02

This page intentionally left blank.

ICS UPDATES AND PATCHES

Workstations and servers run control processes

Many companies do not update their ICS systems

- Can be fragile when it comes to system updates
- Many uptime limitations exist
- Patching windows often do not exist
- Concerns with process stability after patching

If IT must be careful with patching, ICS support teams must be doubly so

ICS systems are often systems of systems and

- Patching one system may cause issues with interdependent systems
- Not all ICS systems can be replicated in testing environments
- Adequately testing patches can be difficult

Above all else, work with your vendors when it comes to patching



Workstations and servers running control processes can be fragile when it comes to system updates. Many companies do not update their ICS systems due to uptime limitations and concerns with process stability after patching. If traditional IT departments must be careful with applying patching, ICS support teams must be doubly so. Not all ICS systems can be replicated in testing environments, so adequately testing patches can be difficult. Also, ICS systems are often systems of systems and patching one system may cause issues with other interdependent systems. The challenges are great, but the risk of doing nothing is greater, as unpatched systems are trivial to compromise.

Patching in the ICS realm is different than your average IT system. When availability is king, taking a system offline to patch means a loss of redundancy for the patch period in a best case. If the patch takes a long time, has unintended consequences, or experiences a number of other difficulties, the system may lose survivability and stability.

The urgency of a patch needs to be weighed against the potential conflicts and ramifications that introducing a patch to a system can have. It is not uncommon for patches to introduce more bugs. (Anytime you make a system more complex, you increase potential problems.) Furthermore, these bugs may not show up immediately on the test systems but can later cause problems when introduced into the production environment. A rushed timeline exacerbates this issue.

Equipment vendors have historically been slow to release patches for vulnerabilities in their products or don't at all. If a serious vulnerability is found in Windows, you can expect a rapid patch resolution. With ICS systems, this is unfortunately not the case.

These factors, combined with additional difficulties discussed elsewhere like network isolation, geographic isolation, and physical barriers, make patching ICS equipment non-trivial.

Caution: Ensure you follow your organization's operational patch management practices. Work with your control system vendor and operations management when planning and implementing any of the items discussed in this module.

ICS VENDOR SUPPORT

Cybersecurity of ICS requires partnership with your ICS vendors

ICS vendors are the experts of their products

- Expertise on how the system is designed and supposed to function
- Road maps to the ongoing functionality and security
- Implement and possibly design the protocols their products use

ICS vendors support your patching efforts in many ways

- Provide security updates and patches to their own products
- Provide guidance on which patches are safe to apply on hosts and connected systems
- Tools to help maintain security or functionality for their products
- Product-specific training and services around patching
- Security advisories and reporting services

→ They need to certify updates and offer confirmation waiting 2-3 months and patching may be a good policy.

A supplier, a manufacturer, or simply a vendor of ICS components or systems should be a partner and important resource to your security efforts and program. Many vendors have security programs and points of contact responsible for interfacing with customers and addressing security issues involving their products. Since the 2010 identification of Stuxnet, there has been an increase in the resources dedicated to security reporting and remediation across ICS suppliers.

Many vendors provide programs for customers to engage in and will provide services such as alerting customers of identified security vulnerabilities through "Knowledge Bulletins" or working jointly with CERTS (like the ICS-CERT) to develop alerts and advisories.

Others will provide base support and will require maintenance contracts and specific services to provide additional security help. The procurement process can be used to establish expectations. Some of your requirements might be satisfied by their standing security programs and others will require additional support for a price. There are many resources available to establish expectations and align requirements.

Vendors also possess expertise in the protocols they support and implement. For example, if you install SIEMENS, you will be using PROFIBUS, ProfiNET or Foundation Field Bus.

VENDOR VULNERABILITY HANDLING

ICS vendors differ in their vulnerability-handling process

- Many vendors struggle with what to release publicly vs. privately
- Community pushing for responsible vulnerability disclosure and handling
- Organizations such as CERT/CSIRT try to help manage these processes
- ICS-CERT in the US has become one of the primary handlers for ICS

ICS-CERT attempts to coordinate between a security researcher and the vendor

- Receive reports from researcher, customer, or other CERT
- Analysis, handling, disclosure, patch development and testing, patch release
- Establishes secure communications (PGP keys, S/MIME/email) between all parties

Share info
only with customer
should be in good
way to follow.

- 1) it is private,
2) can't put a
security perimeter to
protect my system
eg it is public, attach

Vulnerability-handling processes can differ by vendor, but there are some basic process steps to expect. This process sets the means and philosophy as to how a vendor will work with security researchers to learn about security vulnerabilities. The handling process establishes how an organization can be contacted (for example, directly by a researcher or by inclusion of a researcher in a program—some even pay for this type of knowledge or go through a third party to coordinate, like a CERT) and what to expect when reporting security vulnerabilities. There has been an ongoing debate about the merits of open disclosure and responsible disclosure. Many in the ICS community believe in a managed responsible disclosure process. Historic shortfalls in vendors communicating their security weaknesses to their customers and addressing tough issues have fed some arguments that open/public disclosure ensures understanding and is a more powerful force to drive change.

Reference:

Examples include <https://new.abb.com/about/technology/cyber-security>

ICS SUPPLIER SECURITY ADVISORIES

ICS advisories typically provide

- Discovery and publication dates
- Links to related advisory updates (in case you miss one)
- Suggestions of risk, often with CVSS scores and CVE/CWE links
- Summary of risk and impact
- List of affected products (often not entirely complete)
- Technical description of the weakness
- Vulnerability characterization (remote code execution, Denial of Service, etc...)
- Suggested solutions and mitigation options
- Additional resources and links
- Often include recognition for those that discovered the flaw

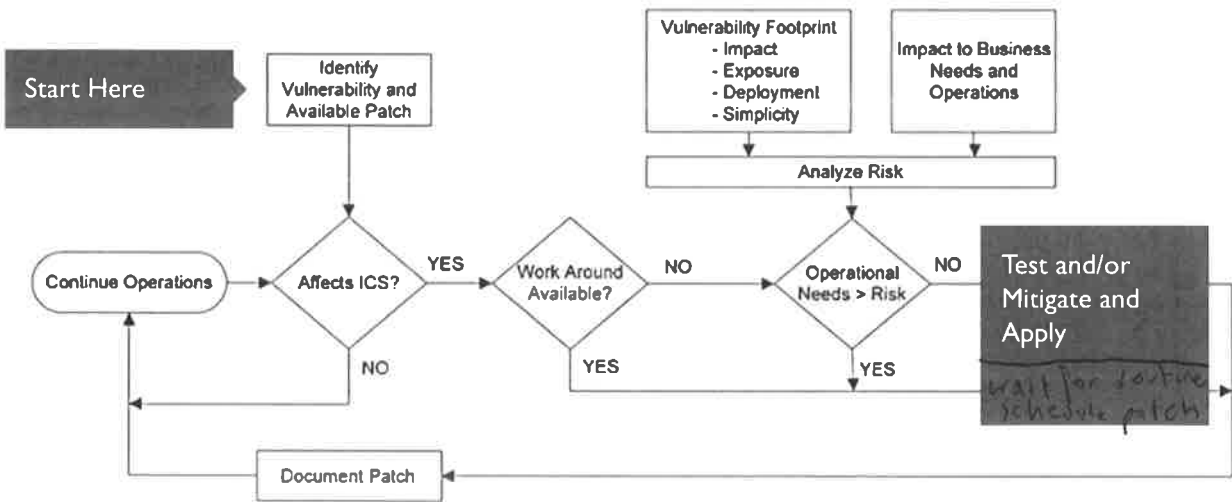
Advisories are only as good as the authors who write them

... and your ability to understand them

Advisories typically cover some basic elements to inform customers of identified vulnerabilities and security threats to a particular product or family of products. The advisory attempts to provide the necessary information so the end user can make risk decisions and implement mitigations or solve the problem by removing a vulnerability with the resources and support provided by the vendor.

The advisory covers the basic elements included on the slide. Some elements are more challenging than others, and the advisory is only as good as the level of understanding of the authors and necessary support teams within an organization to address the security concern. Many vulnerability advisories have been quick to identify a minimal list of products affected and have missed other products that possessed the same vulnerability.

DHS-SUGGESTED PATCH DECISION TREE



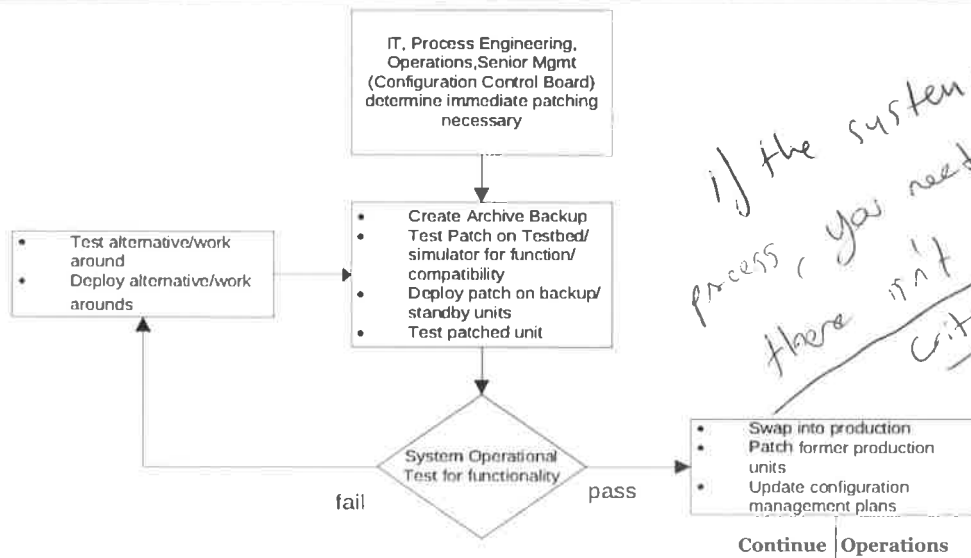
This patch decision tree comes from the US Department of Homeland Security's *Recommended Practice for Patch Management of Control Systems*. This document is the recommended process for all organizations in the US that have ICS-related infrastructures.

Questions about patch necessity need to be asked: Can the patch application be delayed to a later date during a maintenance window? Does the exploit allow an intruder access into other sensitive systems? Does this ICS support life or health? Is there a workaround that we could use instead of patching?

Reference:

<https://goo.gl/8CoLBD>

DHS-RECOMMENDED PATCH MANAGEMENT



if the system is critical for process, you need redundancy. if there isn't is it really critical for you??

SANS

ICS410 | ICS/SCADA Security Essentials 11

DHS Control System Security Program – Recommended Practice for Patch Management of Control Systems

If the urgency determination requires immediate action and a workaround solution is either not available or not the best option, then the following actions should be taken:

1. Where possible, create a backup/archive and verify its integrity by deploying it on a standby system.
2. Create a checklist/procedure for patch activities and deploy the patch on the standby system.
3. Test the patched standby system for operational functionality and compatibility with other resident applications.
4. Swap the patched standby system into production and keep the previous unpatched production system as a standby for emergency patch regression.
5. Closely monitor the patched production system for any issues not identified during testing.
6. Patch the standby system (old production) after confidence is established with the production unit.
7. Update software configuration management plan and related records.

Reference:

<https://goo.gl/8CoLBD>

STUXNET PATCH EVOLUTION

Vulnerability	0-Day Release	Discovery Date	Patch Date
Task Scheduler EoP CVE-2010-3888 MS10-092	March 1, 2010	June 2010	Dec 12, 2010
Keyboard EoP CVE-2010-2743 MS10-073	March 1, 2010	June 2010	Oct 12, 2010
Print Spooler RCE CVE-2010-2729 MS10-061	June 22, 2009	June 2010	Sep 14, 2010
Shortcut .lnk RCE CVE-2010-2568 MS10-046	March 1, 2010	June 2010	Sep 8, 2010

For an example of the importance of a good patch management program, we can look at the latest 0-day exploit in Stuxnet. As we discussed earlier in the class, Stuxnet contained four 0-day exploits in Windows, two for the initial infection and two for privilege escalation. Even after the vulnerabilities were discovered in June 2010, it still took Microsoft 3–6 months to develop patches and release them. By this point, the Print spooler had already been in the wild for well over a year, and the others for at least 6 months. That gives us some insight to the vendor's patch management program. Now what about your patch management program? What percentage of your IDS systems have applied these patches released so many years ago?

EoP – Elevation of Privilege or sometimes referred to as Escalation of Privilege

RCE – Remote Code Execution

labeled as EoP → important by MS. That's why there is a long period between discovery and patching.

If you patch whenever a patch is published in that second you are still late roughly 200 days. so were then 200 days passed until a known vuln. is patched.

EMAIL/NEWSFEED BULLETINS

Easiest way to stay on top of new patches, SPs, exploits, and news

Almost impossible to do your job well if you're not subscribing to some form of security bulletin and/or patch announcement service!

Subscribe to all of your vendor announcement services!

Don't forget the OS vendors

- <https://www.microsoft.com/security/default.aspx>
- <https://access.redhat.com/security/updates/advisory/>

Many excellent bulletins are free!

- <https://ics-cert.us-cert.gov/alerts>
- <http://www.sans.org/newsletters/>
- <https://secuniaresearch.flexerasoftware.com/advisories> *→ add*
- <https://www.packetstormsecurity.com>

There are also many industry and international sources of information

Vendor websites and subscription services are great places to start. These typically offer the most updated and detailed information available for each product; however, beyond your major vendors, this often doesn't scale well for large organizations, and many times not even for small organizations. The sheer number of vendors we deal with, and the wide variety of formats their information is presented in makes using this information extremely difficult. Sometimes, other bulleting services provide a more unified approach. For the ICS industry, ICS-CERT has become one of the main cybersecurity bulleting services used throughout the world.

If you'd also like to browse security sites to stay informed beyond your normal vendor resources, then a good place to start is <https://packetstormsecurity.com>. The Packetstorm website is easy to search, contains a variety of articles from different "perspectives" on security, and usually has a link to any hacking/security tool you are likely to try to find. Another great community resource is <https://www.exploit-db.com>, which contains details on known exploit code for each vulnerability.

NATIONAL COMPUTER SECURITY INCIDENT RESPONSE TEAMS (CSIRT)

Albania: ALCIRT	Czech Republic: CSIRT.CZ	Kazakhstan: KZ-CERT	New Zealand: NCSC	Spain: CERT-SeguridadIndustria
Algeria: DZ-CERT	Czech Republic: GovCERT.CZ	Kenya: KE-CIRT/CC	Nigeria: NgCERT	Spain: ESP DEF CERT
Argentina: ICIC CERT	Denmark: GovCERT.DK	Korea South: KN-CERT	Norway: EkomCERT	Sri Lanka: Sri Lanka CERT/CC
Armenia: CERT AM	Ecuador: EcuCERT	Korea South: KrCERT/CC	Norway: MiljoCERT	Sweden: CERT-SE
Australia: CERT Australia	Egypt: EG-CERT	Kosovo: KOS-CERT	Norway: NorCERT	Switzerland: GovCERT.ch
Austria: CERT.at	Estonia: CERT-EE	Laos: LaoCERT	Oman: OCERT	Switzerland: SWITCH-CERT
Austria: GovCERT.AT	Ethiopia: Ethio-CER2T	Latvia: CERT.LV	Panama: CSIRT Panama	Taiwan: TWNCERT
Azerbaijan: CERT.GOV.AZ	Finland: NCSC-FI	Lithuania: CERT-LT	Paraguay: CERT-Py	Tanzania: TZ-CERT
Bangladesh: bdCERT	France: CERT-FR	Lithuania: LT-CERT	Peru: PeCERT	Thailand: ThaiCERT
Bangladesh: BGD e-GOV CIRT	Georgia: CERT.GOV.GE	Luxembourg: CIRCL	Poland: CERT.Gov.PL	Tonga: CERT.to
Belgium: CERT.be	Germany: CERT-Bund	Luxembourg: GOVCERT.LU	Poland: CERT.PL	Tunisia: tunCERT
Brazil: CERT.br	Ghana: CERT-GH	Luxembourg: NCERT.LU	Portugal: CERT.PT	Turkey: TR-CERT
Brazil: CTIR Gov	Guatemala: CERT-GT	Macau: MOCERT	Portugal: CNCS	Uganda: Ug-CERT
Brunei Darussalam: BruCERT	Hong Kong: GovCERT.HK	Malaysia: MyCERT	Qatar: Q-CERT	Ukraine: CERT-UA
Bulgaria: CERT Bulgaria	Hong Kong: HKCERT	Malta: CSIRTMalta	Romania: CERT-RO	United Arab Emirates: aeCERT
Burkina Faso: CIRT.BF	Iceland: CERT-IS	Mauritius: CERT-MU	Russia: GOV-CERT.RU	United Kingdom: NCSC
Cambodia: CamCERT	India: CERT-IN	Mexico: CERT-MX	Saudi Arabia: CERT-SA	United States: US-CERT
Canada: CCIRC	Indonesia: ID-SIRTII/CC	Mexico: TIC CERT	Singapore: SingCERT	Uruguay: CERTuy
Chile: CLCERT	Iran: CERTCC MAHER	Moldova: CERT-GOV-MD	Slovakia: CSIRT.SK	Uzbekistan: UZ-CERT
China: CNCERT/CC	Israel: CERT-IL	Montenegro: CIRT.ME	Slovakia: GOV CERT SK	Venezuela: VenCERT
Colombia: colCERT	Italy: CERT-PA	Morocco: maCERT	Slovakia: SK-CERT	Vietnam: VNCERT
Côte d'Ivoire: CI-CERT	Italy: IT-CERT	Myanmar: mmCERT	Slovenia: SI-CERT	Zambia: ZM CIRT
Croatia: HR-CERT	Japan: JPCERT/CC	Netherlands: NCSC-NL	South Africa: ECS-CSIRT	
Curacao: CARICERT	Japan: NISC	New Zealand: CERT NZ	Spain: CCN-CERT	

Reference:

<https://www.sei.cmu.edu/education-outreach/computer-security-incident-response-teams/national-csirts/index.cfm>

ENISA

European Union Agency for Network and Information Security (ENISA)

- Created in 2004 by EU Regulation No 460/2004
- Headquartered in Heraklion, Greece with offices in Athens
- Proposed for expanded mission in 2018

Mission of ENISA

- Act as an independent centre of expertise
- Promote cooperation and coordination with EU CSIRT/CERTS
- Support Member State capacity and capabilities
- Recommend policy and best practices
- Support global involvement and cooperation
- Assist policymaking at Union level
- Raise awareness among citizens and businesses
- EU cybersecurity certification framework



Coordinate the Cyber Europe exercises in 2010, 2012, 2014, and 2016

Reference:

<https://www.enisa.europa.eu>

ICS-CERT



The Department of Homeland Security (DHS) National Cybersecurity and Communications Integration Center (NCCIC) includes the Industrial Control Systems Cyber Emergency Response Team (ICS-CERT), whose mission is to guide a cohesive effort between government and industry to improve the cybersecurity posture of control systems within the nation's critical infrastructure. ICS-CERT assists control systems vendors and asset owners/operators with identifying security vulnerabilities and developing sound mitigation strategies to strengthen their cybersecurity posture and reduce risk.

ICS-CERT partners with members of the control systems community to help develop and vet recommended practices, provide guidance in support of the ICS-CERT incident response capability, and participate in leadership working groups to ensure the community's cybersecurity concerns are considered in our products and deliverables.

ICS-CERT also facilitates discussions between the federal government and the control systems vendor community, establishing relationships that foster a collaborative environment in which to address common control systems cybersecurity issues. ICS-CERT is also developing a suite of tools which, when complete, will provide asset owners and operators with the ability to measure the security posture of their control systems environments and to identify the appropriate cybersecurity mitigation measures they should implement.

Reference:

<https://ics-cert.us-cert.gov/>

US INFORMATION SHARING AND ANALYSIS CENTERS (ISACS)



ISACs are trusted entities established by Critical Infrastructure Key Resource (CIKR) owners and operators to provide comprehensive sector analysis, which is shared within the sector, with other sectors, and with governments. ISACs take an all-hazards approach and have a strong reach into their respective sectors, with many reaching more than 90 percent penetration. Services provided by ISACs include risk mitigation, incident response, and alert and information sharing. The goal is to provide users with accurate, actionable, and relevant information. Member benefits vary across the ISACs and can include access to a 24/7 security operations center, briefings, white papers, threat calls, webinars, and anonymous CIKR owner/operator reporting. (Source: *National Council of ISACs*)

Reference:

<https://www.nationalisacs.org/member-isacs>

DOE AND THE NATIONAL LABS

US DOE coordinates 17 national labs that perform research and technical support

Three of these do most of the cybersecurity research

- Idaho National Laboratory (INL)
- Sandia National Laboratories (SNL)
- Pacific Northwest National Laboratory (PNNL)

Very active - internet was analyzed here for reverse.

Here are two examples of DOE research results that are publicly available

- Cyber Security Vulnerability NSTB Program
- Vulnerability Analysis of Energy Delivery Control Systems

The Department of Energy has direct energy-sector control system security resources and a program designed to enhance the security of energy ICS. They also fund various ICS security-related projects and R&D. DOE has included security requirements in recent Smart Grid grants and efforts to modernize energy infrastructures.

DOE funded the National SCADA Test Bed Program (NSTB) starting in the 2002 time frame, which has evolved into a series of partnerships and projects to continue cyber-related risk management for ICSs that impact energy infrastructure.

DOE operates national labs that have a rich history of performing work and technical analysis in ICS-related security. National labs include Sandia National Laboratories (SNL), Pacific Northwest National Laboratory (PNNL), and Idaho National Laboratory (INL) to name a few. These labs have supported other agency ICS security programs such as the DHS CSSP, now simply named the ICS-CERT. The ICS-CERT includes a watch floor currently located at the INL and taps into expertise at different national labs to study technical problems and develop helpful resources for the community.

References:

US DOE National Labs – <https://goo.gl/ocSmBR>

Cyber Security Vulnerability NSTB Program – <https://goo.gl/JGup78>

Vulnerability Analysis of Energy Delivery Control Systems – <https://goo.gl/5UUhJ6>

IMPORTANCE OF BACKUPS FOR SECURITY

Eventually you WILL get hacked or infected ...

Imagine you had to choose between having a perimeter firewall and having a good enterprise-wide backup system. Which would be better for security?

You need to have current backups for

- Forensics analysis
- Performing audits against a baseline
- Disaster recovery
- Accidental data deletion
- Compliance with regulations

Commercial solutions from Veritas, IBM, Dell, Arcserve, Commvault and others exist to scale to enterprise levels

Having good backups is indispensable for doing post-attack and post-infection forensics, performing an audit against a prior baseline, surviving natural or deliberate disasters, restoring accidentally lost data, and staying in legal compliance. You also need to perform a backup of critical systems right before you apply a new Service Pack or patch so that you quickly can get back to square one if the update breaks something critical.

You have to assume that *eventually* your servers will be rooted, your workstations will get infected, and your databases data-diddled into garbage, so what is your plan for recovery? Every security system eventually fails, but a failure doesn't have to be a *Game-Over* catastrophe. You can live to fight another day, but only if you can recover your data. Your hardware, bandwidth, services, and reputation can (usually) be replaced, but it's your data that cannot be replaced—it can be restored only from backups. When backups are restored, ensure appropriate actions are performed prior to reintroduction: Asset name change, IP change, update in AD, license activation, changing log review of modifications to the device since the last backup, and verification of any compliance requirements prior to reintroduction to production.

What if you had to choose between having a perimeter firewall and having an enterprise-wide reliable backup system for all servers and workstations? Assume you can't have both. Which would you choose? Which would be more important for *security*?

"The single most important thing any company or individual can do to improve security is have a good backup strategy."
—Bruce Schneier, *CRYPTO-GRAM Newsletter* (July 15, 2008).

TAKEAWAYS AND RECOMMENDATIONS

Section takeaways

- Patching is one of our primary defenses on each endpoint
- Not expected to patch 100% of systems
- Due diligence is to document each missing patch
- OS patching in Level 3 is easy compared to third-party software on those systems

Recommendations to owner/operators

- Generate metrics of patched systems
- Also generate metrics of systems/devices/software not being tracked

Recommendations to vendors

- Provide solutions to gather patch statistics for all your products
- Create advisories for customers and a portal to accept reports

This page intentionally left blank.

Course Roadmap

Day 1: ICS Overview

Day 2: Field Devices and Controllers

Day 3: Supervisory Systems

Day 4: Workstations and Servers

Day 5: ICS Security Governance

1. Introduction
2. Patching ICS Systems
 - Patch Decision Tree
 - Vendors, CERTS, and Security Bulletins
3. Defending Microsoft Windows
 - Windows Services
 - Windows Security Policies and GPOs
 - **Exercise 4.1: Baselining with PowerShell**
4. Defending Unix and Linux
 - Differences with Windows
 - Daemons, SystemV, and SystemD
 - Lynis and Bastille
5. Endpoint Security Software
 - Antivirus and Whitelisting
 - Application Sandboxing and Containers
 - **Exercise 4.2: Configuring Host-Based Firewalls**
6. Event Logging and Analysis
 - Windows Event Logs and Audit Policies
 - Syslog and Logrotate
 - **Exercise 4.3: Windows Event Logs**
7. Connectivity to the Internet
 - Honeypots
 - Attacks on the Perimeter
 - **Exercise 4.4: Finding Remote Access**

This page intentionally left blank.

Defending Microsoft Windows

Applicable Standards:

- **NIST CSF v1.1:** PR.PT-3
- **ISA/IEC 62443-2-1:2009:** 4.3.3.5.1, 4.3.3.5.2, 4.3.3.5.3, 4.3.3.5.4, 4.3.3.5.5, 4.3.3.5.6, 4.3.3.5.7, 4.3.3.5.8, 4.3.3.6.1, 4.3.3.6.2, 4.3.3.6.3, 4.3.3.6.4, 4.3.3.6.5, 4.3.3.6.6, 4.3.3.6.7, 4.3.3.6.8, 4.3.3.6.9, 4.3.3.7.1, 4.3.3.7.2, 4.3.3.7.3, 4.3.3.7.4
- **ISA/IEC 62443-3-3:2013:** SR 1.1, SR 1.2, SR 1.3, SR 1.4, SR 1.5, SR 1.6, SR 1.7, SR 1.8, SR 1.9, SR 1.10, SR 1.11, SR 1.12, SR 1.13, SR 2.1, SR 2.2, SR 2.3, SR 2.4, SR 2.5, SR 2.6, SR 2.7
- **ISO/IEC 27001:2013:** A.9.1.2
- **NIST SP 800-53 Rev. 4:** AC-3, CM-7
- **CIS CSC:** 3, 11, 14
- **COBIT 5:** DSS05.02, DSS05.05, DSS06.06

This page intentionally left blank.

MICROSOFT CLIENT OPERATING SYSTEMS

Windows 2000 Workstation, sp1, sp2, sp3, sp4 (Feb 2000 – Jul 2010)

- NTFS permissions, RDP

Windows XP, sp1, sp2, sp3 (Oct 2001 – Apr 2014)

- Encrypted File System (EFS), domain-based PKI
- IPSec, Wi-Fi WPA1/2, Firewall, NAC/NAP
- Data Execution Prevention (DEP)

*seen of
Aired*

Windows Vista, sp1, sp2 / 7, sp1 (Oct 2009 – Jan 2020)

- BitLocker Drive Encryption, Advanced Firewall
- User Account Control (UAC), biometrics
- PowerShell, AppLocker, IPSec with DirectAccess

Windows 8, 8.1 (Oct 2012 – Jan 2023)

- Defender AV, PIN/Picture Auth, UEFI Secure Boot

Windows 10 (Jul 2015 – Oct 2025)

- Device Guard, Hello/Passport, Azure Rights Management, Edge browser
- The last Windows Client OS ... unless you count Threshold 1 and 2, and Redstone 1, 2, 3, and 4... ;-)

Microsoft has many client operating systems, extending all the way back to DOS, Windows for Workgroups, Windows 95, and so on (more than can fit on the slide, in fact). Microsoft often experiments with a new kernel or graphical interface by releasing it as a new OS version, and then fixes and polishes that experiment for the next OS release, which is usually more popular; for example, Windows Vista was a flop but cleaned up nicely. Windows 7 and Windows 8 flopped worse than Vista and then came Windows 10. So, where are we in that release cycle right now?

Each client OS is normally released in multiple editions, and it's important for security, manageability, and licensing to install the correct one. Microsoft makes this as confusing as possible. Even the names of the editions can change from one OS version to the next, so you have to research each new release.

In general, though, editions that include words such as Starter or Home are intended for personal use, have fewer features, cannot be joined to an Active Directory domain, and cost less. Editions that include words such as Business, Professional, or Enterprise are intended for business use, have more features, can be joined to a domain, and cost more. The Ultimate editions are aimed at enthusiasts, have more features than the Home editions, can be joined to a domain, and often cost the most.

For example, if you want features such as BitLocker, AppLocker, Kerberos (computer network authentication protocol utilizing tickets to verify identity), Group Policy, Encrypting File System, and Active Directory domain membership, you can't use any of the Home editions. In fact, sometimes there are features you'd expect to get in the Business or Professional editions, such as AppLocker, but get only in Enterprise or Ultimate. Not a nice surprise.

MICROSOFT SERVER OPERATING SYSTEMS

Windows NT 4.0 Server (Jul 1996 – Jun 2004)

- NTFS Secure File System, Terminal Server

Windows 2000 Server, sp1, sp2, sp3, sp4 (Feb 2000 – Jul 2010)

- IPSec, Smart Cards, Group Policy, Kerberos

Windows Server 2003, r2 (Apr 2003 – Jul 2015)

- Server Roles, .NET, LDAP to AD, DFS

Windows Server 2008, r2 (Feb 2008 – Jan 2020)

- Server Core Installs, PowerShell, NAC/NAP
- Read-only DCs

Windows Server 2012, r2 (Sep 2012 – Oct 2023)

- UEFI Secure Boot, Enhanced BitLocker, Early Launch Anti-Malware (ELAM), PowerShell History Viewer, Kerberos Armoring

Windows Server 2016 (2016 – Jan 2027)

- Nano Servers, Shielded Virtual Machines, Credential Guard, Device Guard, Privileged Access Identity

server on shelf

Windows Server is not intended for desktop, laptop, or tablet use, though there is nothing preventing this. Windows Server is normally installed in virtual machines running on computers that often have multiple CPUs, multiple storage devices, multiple network interface cards, a lot of memory, and possibly no monitor. These computers are often densely packed into shelving racks or "blade arrays" in special rooms or railroad-car-sized containers with their own air conditioning, dehumidifiers, and redundant power. However, a Windows Server appliance might be no bigger than a paperback novel with a system-on-a-chip (SOC) architecture and flash memory for storage. In either case, Windows Server is mainly intended for providing network services, such as DNS or HTTP, and not for running graphical applications with a directly attached monitor, keyboard, and mouse.

There are three primary editions of Windows Server: Datacenter, Enterprise, and Standard. The different editions have different scalability and fault-tolerance capabilities, such as for clustering and Network Load Balancing (NLB). Note that starting with Server 2012, the Enterprise edition no longer exists; there's only Standard and Datacenter, with the enhancements of Enterprise edition pushed into Standard by default. Datacenter supports the maximum possible CPUs, memory, clustering, and features. Standard supports the least. Unfortunately, you'll have to research the latest specifications included in each whenever a new OS is released.

There have been specialty editions, such as Small Business Server (SBS) and Windows Server Essentials, which are intended for small offices (less than 25 people), and Windows Storage Server, which is intended for OEM appliances, but these have the same security concerns as the major editions and can be ignored for this course. To better compete with VMware, though, note that Hyper-V Server edition is free, but it can be used only for hosting virtual machines. Despite being free, Hyper-V Server is not artificially hobbled; it supports hundreds of CPU cores, large cluster farms, live migration of VMs, terabytes of memory, and so on (at least for now).

END-OF-LIFE (EOL)

What does EOL really mean?

- Loss of security updates
- Loss of patches/service packs
- Loss of technical support
- Loss of third-party software support
- Loss of new hardware support



This also affects us in larger ways

- Compliance risk with regulations that require supported software
- Reduced options to restore or rebuild overtime
- Greatly diminishes our cybersecurity mitigation options
 - No AV/endpoint security software; limited to legacy hardening techniques

For attackers, this all increases value for vulnerabilities and exploits

In general, plan to be off a product before End of Mainstream Support and certainly before End of Extended Support (the dates in the slide are for End of Extended Support). Here are the essential terms to know:

- **End of Sales:** When the product is no longer sold to retailers or OEMs, but these resellers might stockpile licenses before End of Sales is reached.
- **End of Mainstream Support:** When warranties expire for the product, the product will no longer be improved, free incident support ends, and non-security hotfixes become unavailable unless specifically purchased during the (expensive) Extended Support phase.
- **End of Extended Support:** When security hotfixes and paid support can no longer be purchased, except in special cases with (expensive) Custom Support.

Microsoft has continued to provide notice on the planned end-of-life (EOL) for support of many of its software products. For instance, Windows XP was supported for 12 years and entered EOL on April 8, 2014. There are significant impacts for an organization that currently relies upon these systems. These impacts need to be understood and considered in planning for migrating to a supported operating system and in instituting security mitigations until an upgrade can be properly implemented. Microsoft has recommended that users take action and that they be on notice that EOL products should be considered "not protected," exposing users to the following risks:

1. No additional security updates and degraded protection from endpoint software
2. Compliance risk associated with regulations that require supported software
3. Lack of future independent software vendor support
4. New hardware will stop being supported (drivers and more)
5. Increased cost to maintain

Windows is a prevalent OS used for ICS/SCADA workstations and can be found as the OS behind many local HMI panels. The EOL will impact many ICS-reliant organizations. The implication of end-of-life should be understood and fully evaluated. Organizations should identify where non-supported OS exist, including components such as local HMIs and OS-driven devices. Organizations should mitigate existing implementations and plan for a future migration. Several third parties offer mitigation services for ICS implementations.

The automation and industrial technology community has been adjusting business and planning processes to account for the rate of change in the technology market. Traditional ICS planning strategies did not fully account for all the elements of an ICS and how they might be impacted by technology innovation and IT provider business models. The convergence of technology supporting industrial, business, and consumer applications results in faster life cycles for equipment and software. There are real business drivers to get the most of your capital project with a desire to operate for 10 to 18+ year time periods. Some regulated industries must get value for their investments over long periods of time or risk having stranded assets. The use of general computing technology in ICS has provided benefits in features, reliability, and cost but has carried with it the need to plan for shorter life cycles and the ability to update or modernize elements of a larger system more frequently.

The EOL of an OS results in significant implications for the security of impacted platforms and applications residing on those devices. Microsoft has recommended isolation of workstations running Windows XP to protect the platform and its installed applications. This guidance may not be appropriate when the workstation or device is a part of a larger integrated system (such as a DCS or SCADA EMS). ICS security teams should understand the security implications and assess how EoL impacts the security of their deployed technology. EOL has immediate impacts and longer-term impacts to consider:

Microsoft will no longer analyze disclosed vulnerabilities and develop security patches/updates for them. It is important to note that extended support (paying Microsoft to continue specialized support for your organization) is advertised as not including security patches/fixes.

Reuse of Microsoft XP code segments in existing and supported Microsoft OSs will likely result in vulnerabilities that are discovered in Windows 8/10, for example, that may apply to XP. Microsoft will provide a security patch/fix for the supported Windows 8/10, but not for the vulnerable XP.

The lack of future security updates/fixes will impact the future capability of antivirus solutions. Many independent antivirus solutions will continue to support Windows XP for some limited time period.

EOL will make existing non-disclosed and future discovered vulnerabilities for Windows XP valuable in the underground vulnerability/exploit market. Some have already reported a reduction in discovered but not reported vulnerabilities in anticipation of better pricing.

The loss of equipment due to natural causes or hazards will be more difficult to restore/rebuild as newer hardware won't be able to support Windows XP (lack of drivers, and so on). This issue will grow over time (3 years and beyond) as hardware manufacturers stay current with market demands and optimize for more current OS technology.

Microsoft makes a compelling case for migrating to a newer OS based on the cost to maintain and enhance security features found in Windows 10. According to Microsoft, a Windows 10 PC is 21 times less likely to be infected by malware than an XP machine because of security improvements engineered into the OS.

One way to think about the implications is to consider your "extended security team": Your extended security team includes your ICS suppliers and the vendors that make the OSs/RTOSs, applications, and hardware you rely upon to secure and operate your systems. If you continue to use Windows past its EOL date, you can consider that a significant reduction of your security team as Microsoft security and product experts will no longer be providing security patches/fixes.

NON-SECURITY EXAMPLE OF EOL

XP and 2003 have heavy use in existing ICS

- Java no longer supported in these products

ICS products that utilize Java

- Siemens – SIMOTION Controller
- Inductive Automation – Ignition
- GE Intelligent Platforms – Proficy Real-Time Information Portal
- GE Intelligent Platforms – Proficy Machine Edition
- Emerson Process Management – Ovation
- Trane – Tracer SC
- IDEC – Web server module
- Siemens – SIMATIC HMI sm@rtclient
- Hirschmann – Eagle

When an operating system is no longer supported, there is often a ripple effect across the industry. Once the OS is EOL, most application vendors will stop supporting products that run on that OS. This compounds the impact with applications like Java. Now, not only are all the OSs vulnerable, but all the Java applications are also exposed and vulnerable. Because many ICS products rely on Java, this creates a dilemma on how to handle it. In many traditional IT environments, upgrading to a supported version is the typical recommendation. In controls systems, this is not always possible, so other forms of mitigation are used, which include network segmentation, enhanced monitoring, or accelerated component upgrade.

RECOMMENDED MITIGATIONS

Identify and document all instances!!!

Evaluate hardware availability on EOL systems

- Consider virtualization

Consider less protected/trusted

- Place network defenses around the system
- Consult ICS vendors for mitigation guidance
- Update OS to final patch levels, if possible
- Conduct system hardening on the OS
- Perform security testing to verify defenses

Consult ICS vendors/suppliers and request their mitigation guidance for EOL Microsoft products. You should develop a mitigation plan and share that plan with the suppliers to evaluate any risks associated with implementing technical mitigations. Mitigations are necessary as your organization plans for the eventual migration to relying upon supported OSs. All mitigation plans begin with identifying where the non-supported OS exists in your ICS. This may not be a trivial task as many devices, like local panel HMIs, rely upon Windows XP as the underlying OS.

Microsoft has told users that EOL operating systems should now be considered "not protected." This means you should consider the workstations and devices running EOL operating systems as being less trusted (conceptually) and consider how to better protect these workstations and protect other ICS components. Microsoft recommends isolation of devices running EOL software, which is likely to be difficult if not impossible in ICS applications. Segmentation should be considered if possible. Mitigations can include a number of additional security practices, administrative controls, and technical controls.

Consulting your ICS vendors for system-specific mitigation guidance is recommended. Some vendors will have well-thought-out mitigation recommendations, whereas others will not. The vendor's guidance can help inform the development of your mitigation plan. Your final plan should also be shared with suppliers to ensure that any additional technical controls will not negatively impact the reliability of the system.

A longer-term business and security consideration is the availability of hardware spares that are able to support EOL software. This issue will become a greater risk over time as new hardware will not be designed to work with the non-supported OS. Your mitigation plan can include an initial evaluation of the health of your existing hardware running the non-supported OS. Plan for a visual inspection of existing hardware and consider maintenance steps designed to extend the life of the equipment. (This may include properly cleaning machines or providing environmental barriers for reducing exposure to dust and particles or extreme temperatures.)

Hardware that is already at risk might require replacement or upgrades to critical components such as backup power supplies. Organizations might even stock for additional parts and spares to support implementations that will need to function over longer periods of time before modernization can take place. There are some implementations that may never be replaced, such as computers in difficult-to-access places (at great depth under the ocean, for example).

Update your EOL software to the final patch level if possible, or to as current as you can, based on constraints or limitations. The final updates and fixes will still be able to be deployed using administrative tools, third-party patching systems, and Microsoft's own enterprise patch management solutions.

System hardening for workstations may include the following:

- Removing all unnecessary applications (e.g. development tools, games, unused programs) to reduce the available attack surface and simplify the required processes and executables to be monitored. Utilize tools like CIS benchmarking for establishing appropriate settings.
- Further restricting unused physical and logical ports and unnecessary services. Consulting your supplier can help you to develop the most restrictive configuration possible while still supporting system operations and maintenance.
- Disabling Autorun.
- Placing administrative controls on who may interact and how they interact with impacted systems.
- Implementing policies and controls to restrict the use of USB, the mobile media, and the ability to place software on the workstation.
- Implementing strong user authentication and restricting the roles and users that can log in to the workstations.
- Updating host-based security solutions and reviewing future support options.

Develop a testing capability to support problem identification/validation and test future mitigations. Your staff or support contractors may need to fill in the gaps left by Microsoft not supporting the OS. Being able to re-create identified vulnerabilities and assess the risks they pose to your organization may be necessary. You may even need to devise technical fixes and mitigations to address high-risk vulnerabilities or respond to known exploits.

Consider whether virtualization can help address concerns surrounding future hardware availability and can provide advantageous security features including monitoring the platform or enhancing your ability to back up and quickly restore.

WINDOWS SERVER UPDATE SERVICES (WSUS)

Free solution from Microsoft to do central patch management

- Signed patches from Microsoft for any system connected to Active Directory running Windows
- Since WSUS version 3, can deploy third-party patches (aka non-Microsoft software)
- Can scale up to to 30,000 systems per WSUS server

How does it work?

- WSUS server synchronizes with Microsoft (or a WSUS server in your Control DMZ) via HTTPS
- WSUS administrators test and approve patches for individual systems or groups of systems
- WSUS clients on each of your machines report into your WSUS server and check for approved patches
- Approved patches can be downloaded and installed manually per system
- On less sensitive systems such as workstations and laptops, patches can be automatically deployed in phases

For large remote sites with limited WAN bandwidth, patches can be carried into the site physically and deployed via a local WSUS server

Works well in ICS for systems connected to Active Directory in Purdue Level 3 and might be able to extend to Level 2 in some instances

provides recordability also.

Windows Server Update Services (WSUS) is a free solution from Microsoft that allows central patch management of systems connected to Active Directory. WSUS permits the grouping of computers into custom sets that can be managed and updated separately. For example, some machines, such as workstations, might download and install updates automatically, whereas others, such as servers, must be updated manually. A single WSUS server can handle up to 30,000 client computers, and WSUS servers can be load-balanced and chained together to form administrative arrays.

Administrators choose exactly which updates they want distributed inside the LAN to WSUS clients. Presumably, this is after the administrator has tested the patches in a lab. You select which hotfixes you want deployed in the WSUS console in Administrative Tools. Earlier versions of WSUS were managed through a special administrative website on the WSUS server itself, but not anymore.

The patch files from Microsoft can be downloaded to the local WSUS server and then downloaded by clients from there, or clients can be directed to download the approved updates directly from Microsoft. Caching the files locally on the WSUS server, of course, will spare internet bandwidth consumption. If you have a large number of clients, you can make your other WSUS servers download their files from your master WSUS server that downloads from Microsoft.

WSUS clients download updates using the Background Intelligent Transfer Service (BITS). BITS "drizzles" files down to the client in the background so that other applications are not interrupted and bandwidth is not monopolized by WSUS traffic.

One of the benefits of WSUS is the ability to create your own custom groups of computers and then approve or deny different sets of updates for each group of computers. For example, the patches and Service Packs approved for the "Executives" computer group will likely be different than the updates for the "External_IIS" computer group. These groups are defined and used within WSUS alone; they do not require or use Active Directory groups.

MICROSOFT SYSTEM CENTER CONFIGURATION MANAGER (SCCM)

WSUS on steroids

- Ability to create your own third-party software patches... ICS software
- System provisioning
- Inventory control for both hardware and software
- System performance monitoring

Reasons why you may choose WSUS over System Center

- Additional management overhead and complexity
- Some level of increased resource usage on each managed system
- Not a free solution like WSUS

Some third-party patch/config management solutions leverage WSUS and SCCM

References:

<https://www.microsoft.com/en-us/cloud-platform/system-center-configuration-manager>

SERVICES IN WINDOWS (SERVICES.MSC)

Name	Description	Status	Startup Type	Log On As
BranchCache	This service caches network content from peers on the local subnet.		Manual	Network Service
Cryptographic Services	Provides three management services: Catalog Database Service, which confirms the s...	Running	Automatic	Network Service
Distributed Transaction Coordinator	Coordinates transactions that span multiple resource managers, such as databases, ...	Running	Manual	Network Service
DNS Client	The DNS Client service (dnscache) caches Domain Name System (DNS) names and re...	Running	Automatic (Trigger Start)	Network Service
Downloaded Maps Manager	Windows service for application access to downloaded maps. This service is started o...		Automatic (Delayed Start)	Network Service
Fax	Enables you to send and receive faxes, utilizing fax resources available on this compu...		Manual	Network Service
IPsec Policy Agent	Internet Protocol security (IPsec) supports network-level peer authentication, data cr...		Manual (Trigger Start)	Network Service

First step in hardening a system

Identify and disable any services you don't need

- Decrease the attack surface
- Focus primarily on network services
- Sort by the last column to group in Windows service tool



This page intentionally left blank.

part-ir-pid list → netcat -nao (find 'LIST')

SECURITY POLICIES IN WINDOWS

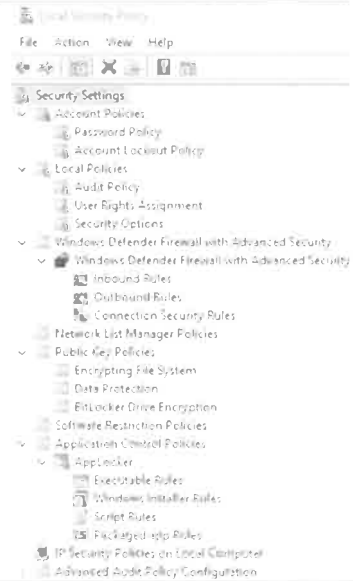
Second step in hardening: Enable better security settings

A template can store the following security settings

- Account authentication settings
- Audit policies for logging
- User rights around system administration
- Security options for internal Windows functionality
- Firewall settings
- Network lists for security-level definitions
- Public key cryptographic settings
- Software restriction policies (SRPs)
- AppLocker settings for program sandboxing
- IP security policies for host-to-host communications
- Advanced Audit Policies for additional logging

Templates can be applied like a rubber stamp to multiple machines

- Groups of machines via Group Policy Object (GPO) in AD
- Locally on non-AD systems with secpol.msc, scw.exe, or secedit.exe



SANS

ICS410 | ICS/SCADA Security Essentials 33

A security template is a plaintext configuration file that can store hundreds of security settings. A computer can be stamped with a template and reconfigured in one shot to match the settings in the template. Security templates are kept by default in %SystemRoot%\Security\Templates\ and also %SystemRoot%\Inf, and they end with the .INF filename extension. Templates can be edited with Notepad, but a much easier method is to use a Microsoft Management Console (MMC) snap-in named Security Templates.

On Windows 10, just use the search feature to open the Local Security Policy tool. In older versions of Windows, you will want to run secpol.msc on workstations or scw.exe on servers. In the Templates snap-in, double-click a template to open it up. Browse through its containers and double-click the policy icons to configure them. When you're done, right-click the edited template and save it. If you want, browse to that template's folder using Windows Explorer and open the template in Notepad.

Another option is the secedit.exe command-line tool for working with the security policies. You could use scripting across the network to quickly reconfigure hundreds of isolated machines not being managed by Active Directory! Alternatively, you could save the secedit.exe and security templates to a USB to manually update isolated and islanded machines.

Most changes take effect immediately, but to ensure that all settings have been reapplied, reboot the system after importing a template. Again, just as before, obtain a preconfigured template from an organization you trust; modify that template to match your preferences and then apply it using the SCA snap-in and/or Group Policy Objects.

Reference:

Administering Security Policy Settings – <https://goo.gl/oiNSs2>

*If local accounts exist, tune
it for logging in only from
console*

CREATING SECURITY TEMPLATES

Don't start from scratch!

- Use preconfigured and debugged security templates
- Test and modify as needed for each system type you harden
- Microsoft has two solutions to aid you in this process

Microsoft Security Compliance Manager

- For use with Windows 10, Windows Server 2012, and older
- The kit includes a variety of templates for different products
- Use the kit to export, compare, and otherwise manage templates

Microsoft Security Compliance Toolkit

- For use with Windows 10 1507, Windows Server 2012 R2, and newer
- Focuses on simplifying the process and providing flexibility to future systems

You don't have to create your own templates from scratch. Many players in the Windows security arena offer customized templates free for the download. It is highly recommended that you begin with someone else's templates instead of starting from scratch. The reason for this is that security is bad for usability. In general, the more security options you configure, the more applications you are likely to break. Templates from Microsoft, NIST, CIS, and others have been debugged and tested to improve security as much as possible while breaking as little as possible.

Remember, too, that you always can edit a template obtained from someone else in any way you want, so start with a template from someone you trust, then test and fine-tune it for your organization's needs.

Microsoft has a set of security templates and best practices for various applications and versions of Windows. These templates provide an excellent starting point for your internal testing. They should not be applied without compatibility testing first. To access the templates and their associated documentation, download one of their two free tools:

References:

Microsoft Security Compliance Manager (SCM) – <https://goo.gl/pCafYQ>

Security Compliance Toolkit (SCT) – <https://goo.gl/pnPAfX>

GROUP POLICY OBJECTS (GPO)

Security Policies can be deployed automatically via Active Directory

GPOs in Active Directory can be applied to organizational units

- GPOs overwrite Local Security Policies
- When applied to a machine object, it takes affects when no one is logged in
- When applied to a user object, that policy takes place upon login
- **Example:** Different security policies for engineers and technicians

GPOs are applied automatically

- At bootup
- At logon
- 90–120 minutes

Managing group policies

- Group Policy Management Console (gpmc.msc)
- Local Group Policy Editor (gpedit.msc)

Think of Group Policy Objects (GPOs) as special logon scripts that, when run, can reconfigure almost anything on the computer, including the user's desktop. When a computer boots up, it will download the GPOs assigned to it and execute them automatically. Every 90 to 120 minutes thereafter, the computer will check to see that none of its GPOs have been changed, and, if any have, then the computer will download the edited GPOs and run them automatically, too, even if the computer has not been rebooted.

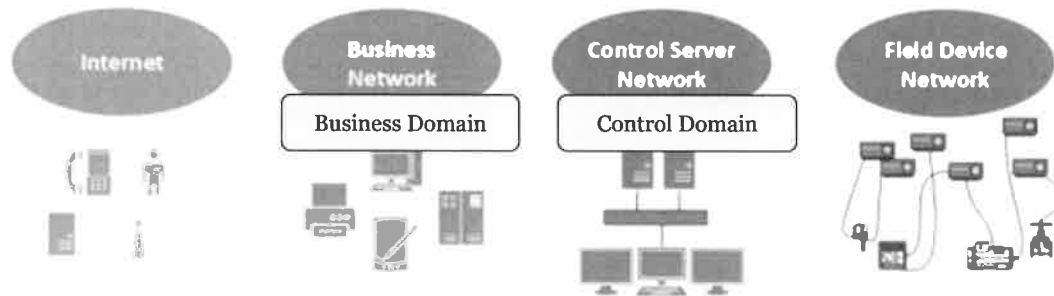
Similarly, when a user logs on, her computer will obtain the GPOs for that user and execute them automatically to reconfigure the user's desktop. Every 90 minutes, the computer will check for any newly edited GPOs and reapply them. Some settings do not take effect until after the user logs on again, but many settings apply immediately.

On an AD domain controller, you can see a domain-wide GPO by going to Start > Programs > Administrative Tools > Default Domain Policy. To see the entire GPO, open the "Active Directory Users and Computers" snap-in tool > right-click the name of your domain > Properties > Group Policy tab > highlight Default Domain Policy > Edit. This is the same "Default Domain Policy" GPO, but the snap-in shows all its properties. Best practice is to usually manage GPOs at the Organizational Unit level instead of the Domain level since few domains can have the same GPO on all systems.

If you have the Group Policy Management Console (GPMC) installed instead, double-click your forest > double-click your domain > right-click the Default Domain Policy icon > Edit. The GPMC is a free download from Microsoft and is built in to Windows Vista and later. The GPMC will be your primary tool for managing Group Policy when you become more familiar with GPOs.

DOMAIN CONTROLLERS IN ICS

If Active Directory is needed in ICS, a separate domain with no relationships with the business domain should be used



When deploying Active Directory in control networks, a separate domain with no relationships to the business network should be used. Any relationships or hierarchies created between business and control domains create additional risk to the control network. These relationships require too many holes in the control network perimeter, can possibly create control process interruptions due to unavailability, and create additional attack opportunities from the business network to the control network. Thus, the most secure and least impactful solution to control processes is the creation of a totally independent AD domain for the control network.

If AD is used in the control network, AD groups can be created for both users and workstations. This will allow you to restrict which individuals can log in to specific groups of machines and provide some level of control and audit who can access control interfaces. This is especially useful when the ICS software doesn't have sufficient authentication capabilities, thus allowing you to control access at the operating system level instead.

for file sharing between Business and Control Domain, there may be another Domain in DMZ. This Domain trusts Business or Control Domain, But one-way. It copies files from one site to another.

ADDITIONAL MICROSOFT SECURITY SOFTWARE

Privileged Access Workstations (PAWs)

- Run sensitive tasks in hardware plus virtualization security containers
- Limit lateral movement between processes on the same host

Local Administrator Password Solution (LAPS)

- Randomizes local passwords on each host and stores them in AD
- Limits lateral movement with local admin credentials
- Alternate PowerShell-based solution by Jason Fossen at <https://goo.gl/gEQ18z>



Enhanced Mitigation Experience Toolkit (EMET)

- Uses several mitigation techniques to limit RCE exploits
- EOL as of July 31, 2018
- Windows 10 and beyond have integrated many of its features with built-in Exploit Guard

Windows Defender Advanced Threat Protection

- Cloud-based solution to identify threats on Windows endpoints
- Built in to Windows 10, so no agent needed
- Not good for control networks, but might be a solution for remote laptops with internet connectivity

References:

PAWs – <https://goo.gl/gCAXaN>

LAPS – <https://goo.gl/WCp2wE>

Jason Fossen's PowerShell tool – <https://goo.gl/gEQ18z>

EMET – <https://goo.gl/D792TB>

Exploit Guard – <https://goo.gl/9Tkj19>

Windows Defender Advanced Threat Protections – <https://goo.gl/9EMPZU>

AUTOMATION

Everything that can be done with graphical tools can be done from the command line or a script

A vast amount of auditing data can be extracted with scripts

- Creating and comparing security baselines
- Verifying policy compliance
- Changing detection and analysis
- Scanning and reporting of Indicators of Compromise (IOC)

Scripting allows you to perform cybersecurity functions without the need to modify assets

- Native scripting languages are installed by default and thus agentless
- Scripts do not have to be copied to the hard drives of machines running them
- Perfect solution for ICS when we can't modify systems; just be careful with your scripts

When you have a breach, these scripting skills become **ESSENTIAL!**

- First thing asked of you by incident responders you call for help
- Don't let an incident be the first time you run a script across large numbers of assets
- Grow your experience and your management's comfort level now

Auditing and forensics are critical duties security administrators often need to perform. In ICS environments, it is effective to audit systems against security baselines and audit similar systems against each other to identify differences.

under the wire tech → Just to learn powershell. Web Site

MICROSOFT POWERSHELL

PowerShell for Windows Scripting

- It's a free command shell to replace the old CMD.EXE
- It's installed on Windows 7, Server 2008, and later by default
- It's available for Windows XP-SP2, 2003-SP1, and Vista

Thousands of cmdlets for almost every task

- <https://docs.microsoft.com/en-us/powershell/>
- Including for VMware, Amazon Web Services, and others

PowerShell Remoting

- Runs commands and scripts remotely
- Supports Kerberos and SSL/TLS

Microsoft PowerShell is the name of both a scripting language and a command shell in which that language can be used. PowerShell has all but replaced CMD.EXE as "the command shell" on Windows. The intention behind PowerShell is to have a better command shell on Windows than any shell found on any flavor of Unix or Linux, and they just might have done it!

PowerShell is built in to Windows Server 2008 and Windows 7 and later and can be downloaded and installed on Windows XP-SP2/2003-SP1/Vista-SP0 or later, too. To download PowerShell and a bunch of sample scripts for it, visit <https://docs.microsoft.com/en-us/powershell/>

A tool run within PowerShell is called a *cmdlet*. The most important feature of PowerShell is that cmdlets support the piping of text and objects, thus making the properties and methods of those objects available to subsequent cmdlets in a pipelined chain of commands. While Unix/Linux commands merely pipe text, PowerShell is thoroughly object-oriented and built on top of the .NET Framework.

In the CMD.EXE shell, you are typically "in" a directory, such as the C:\Windows directory. In PowerShell, you can be in a directory, too, but you can also be "in" the registry, the IIS Metabase, Active Directory, or "in" any other hierarchically formatted data structure that PowerShell recognizes. Hence, enumerating through all the objects in a folder, a registry key, or in an organizational unit is fundamentally the same operation.

It's important to know about PowerShell because it is the future of command-line administration on Windows. We can't discuss the PowerShell language in detail here, but here are a few sample commands. Remember that the pipe symbol ("|") in PowerShell is also for piping full .NET objects with their properties intact. You might want to create a shortcut on your desktop for "powershell.exe" and then right-click it and Run As Administrator.

Again, there is too little time to discuss PowerShell syntax today. The Windows track at SANS (SEC505) includes a full 1-day course on PowerShell, and you can freely download the sample scripts for this course from <https://cyber-defense.sans.org/blog/>. (All the scripts are in the public domain.)

WMIC.EXE

Older, non-PowerShell ways still exist

Can manage tons of configuration settings!

- Built in on Windows XP and later
- Works on local or remote systems

Show last Service Pack number applied

- `wmic.exe os get servicepackmajorversion`

Show shared folders on remote system named Server52

- `wmic.exe /node:Server52 share list brief`

Show commands which run at startup

- `wmic.exe startup list full`

Windows XP and later includes a command-line tool named WMIC.EXE (Windows Management Instrumentation Console) that can be used to get or set configuration data for a wide variety of settings. Think of this tool as the admin's Swiss Army knife.

For example, to find out the number of the last Service Pack applied, you could run this command:

```
wmic.exe os get ServicePackMajorVersion
```

Note that you can run the tool against remote systems, too, using the "/node" switch.

To get a list of the shared folders on a remote computer with IP address 10.4.2.2:

```
wmic.exe /node:10.4.2.2 SHARE list brief
```

Note that share folders utilize CIFS protocol and TCP port 445; this may be blocked.

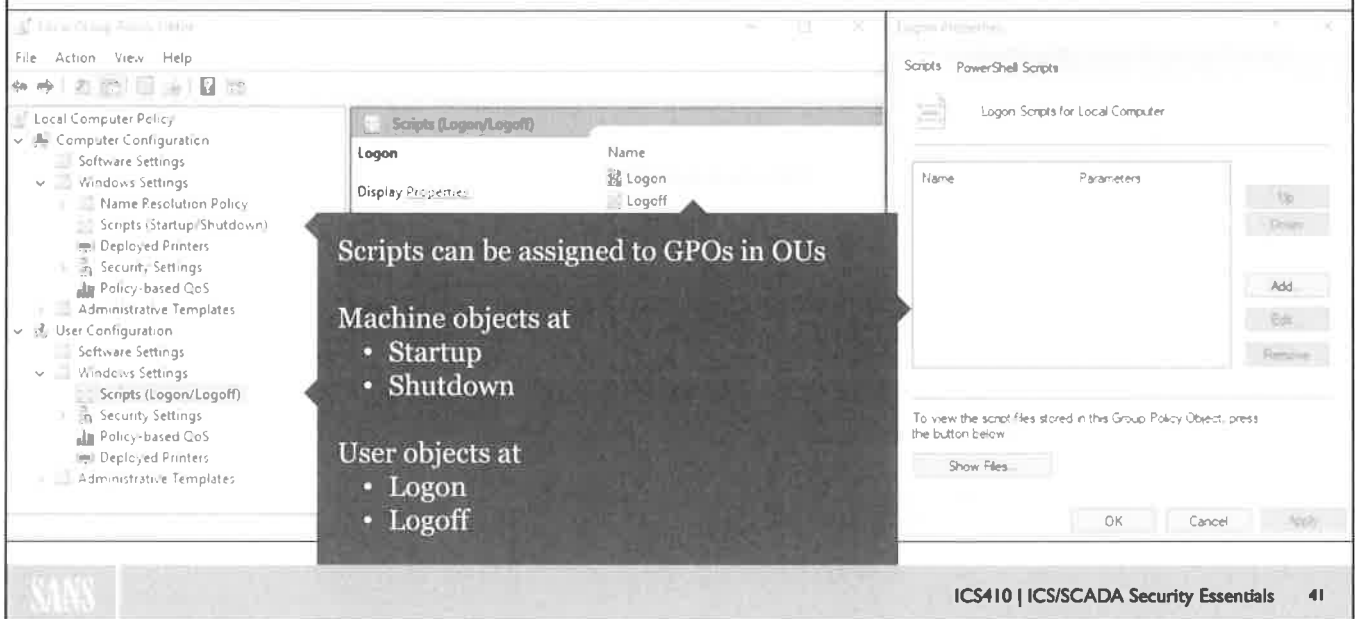
To dump a list of the programs that automatically execute at bootup (registry Run key):

```
wmic.exe /node:10.4.2.2 STARTUP list full
```

In these commands, you can get one particular value from a category of information (*os get ServicePackMajorVersion*) or do a brief/full listing of all the values from that category (*list brief*). Entering "`wmic.exe /?`" can display the categories of information available (OS, SHARE, STARTUP, etc.). Just run "`wmic.exe <categoryname> list full`" for each category shown when you run "`wmic.exe /?`"

Later in this section is a batch script (SNAPSHOT.BAT) that uses WMIC.EXE extensively to dump information. See this script for more examples of usage.

PUSH SCRIPTS WITH GROUP POLICY (GPMC.MSC OR GPEDIT.MSC)



You can assign scripts to be executed through a GPO as well. These options are found in both Computer and User Configuration > Windows Settings > Scripts. It bears repeating in this context that Group Policy can push out scripts to machines automatically. Because Group Policy Objects can be linked to individual Organizational Units, each OU can have its own custom set of scripts. These scripts can be executed at startup, shutdown, logon, or logoff. They can be written in any language for which the necessary interpreter is installed. Windows 2000 and later includes interpreters for batch files, JScript, and VBScript by default, but interpreters for Perl and Python can also be installed.

Scripts can be written in any language, as long as the necessary interpreter has been installed. By default, scripts can be written only in PowerShell, VBScript, Jscript, or as batch files. Startup/shutdown scripts run in System context, whereas logon/logoff scripts run in the context of the user. You can have as many of each type of script as you want, and you can mix and match your languages across your scripts as desired.

You can push out as many scripts as you want, and you can mix scripts written in different languages in a single category; for example, your logoff scripts might include one batch file, two PowerShell scripts, and three VBScripts. Domain controllers multi-master replicate scripts to each other automatically using the File Replication Service (FRS). Logon/logoff scripts execute in the context of the user, whereas startup/shutdown scripts execute in the context of the local System account.

Virtually every aspect of the operating system can be scripted, including users, groups, NTFS ACLs, shared folders, registry settings, IIS, Office application settings, and more.

example usage of script
Startup → create baseline system start up, get a baseline and send it to center.

TAKEAWAYS AND RECOMMENDATIONS

Section takeaways

- Patching is only the first part of protecting an endpoint
- Decrease your attack surface by disabling services and functions

Recommendations to owner/operators

- Start to harden new deployments before SAT testing is performed
- Revisit current systems based on risk and change windows

Recommendations to vendors

- Work with third-party patch management vendors for integration
- Provide Windows security policies for each product deployed there

This page intentionally left blank.

Course Roadmap

Day 1: ICS Overview

Day 2: Field Devices and Controllers

Day 3: Supervisory Systems

Day 4: Workstations and Servers

Day 5: ICS Security Governance

1. Introduction
2. Patching ICS Systems
 - Patch Decision Tree
 - Vendors, CERTS, and Security Bulletins
3. Defending Microsoft Windows
 - Windows Services
 - Windows Security Policies and GPOs
 - **Exercise 4.1: Baselining with PowerShell**
4. Defending Unix and Linux
 - Differences with Windows
 - Daemons, SystemV, and SystemD
 - Lynis and Bastille
5. Endpoint Security Software
 - Antivirus and Whitelisting
 - Application Sandboxing and Containers
 - **Exercise 4.2: Configuring Host-Based Firewalls**
6. Event Logging and Analysis
 - Windows Event Logs and Audit Policies
 - Syslog and Logrotate
 - **Exercise 4.3: Windows Event Logs**
7. Connectivity to the Internet
 - Honeypots
 - Attacks on the Perimeter
 - **Exercise 4.4: Finding Remote Access**

This page intentionally left blank.

EXERCISE 4.1: BASELINING WITH POWERSHELL

EXERCISE DURATION: 15 MINUTES

We are going to explore how to use PowerShell to create basic security baselines of a system. If you already have experience doing this, we recommend you think of a setting in Windows that you don't know how to collect with PowerShell and research it.

OBJECTIVES

Learn PowerShell basics

- Get-Help cmdlet
- Filtering output with Select-String

Learn how to collect a basic baseline

- List all installed services
- List all network connections
- List all local users
- List all installed software

PREPARATION

Start your Windows 10 virtual machine

This lab can also be performed on your host machine if you are running Windows 10

This page intentionally left blank.

EXERCISE 4.1: BASELINING WITH POWERSHELL

GETTING AROUND IN POWERSHELL

Start PowerShell using the search box next to the start menu.

Try the following cmd.exe commands in PowerShell

```
PS C:\Windows\system32> cd \Users\Public
PS C:\Users\Public> echo "My text file" > ICS410.txt
PS C:\Users\Public> type ICS410.txt
PS C:\Users\Public> copy ICS410.txt ICS410-2.txt; dir
PS C:\Users\Public> del ICS410-2.txt; dir
```

Now try the following psuedo-Linux commands in PowerShell ☺

```
PS C:\Users\Public> cd ~ (or you can type: cd \Users\student )
PS C:\Users\student> ls > file1.txt
PS C:\Users\student> ls > file2.txt
PS C:\Users\student> cat file1.txt, file2.txt
PS C:\Users\student> man diff
PS C:\Users\student> diff (gc file1.txt) (gc file2.txt)
PS C:\Users\student> ps
```

In many cases, PowerShell can be used just like cmd.exe; however, its output is usually slightly different looking than cmd.exe, and not all cmd.exe commands are available. Here is a breakdown of the commands above:

- Change directory to the Public user's home directory.
- Echo some text into a text file called ICS410.txt.
- Type the contents of ICS410.txt to the screen.
- Copy ICS410.txt to a second file with a similar name and list files in the directory.
- Delete the new file copy and list files in the directory.

Something that is initially exciting about PowerShell for Linux users is the aliases for basic shell commands. The commands above do the following:

- Change directory to your user's home directory; in most cases this will be "student").
- List the files in the directory and save as file1.txt.
- List the files in the directory again and save as file2.txt.
- Concatenate the contents of both files to the screen. (Note the needed comma.)
- Show the help page for the diff command (no, not normal diff syntax).
- Show the difference between the two files we created.
- Show the running processes.

Hey, for Linux users, that is a great start. Unfortunately, the aliases don't go much further than that, and the commands don't support all the same great options we are used to in Linux.

EXERCISE 4.1: BASELINING WITH POWERSHELL

GETTING HELP

Use the get-help command to use different cmdlets.

```
PS C:\Users\student> Get-Help
PS C:\Users\student> Get-Help *
PS C:\Users\student> Get-Help * | more
PS C:\Users\student> Get-Help Get*
PS C:\Users\student> help Set*
PS C:\Users\student> help Select*
```

Now let's use the Select-String cmdlet to "grep" our text file.

```
PS C:\Users\student> help Select-String
PS C:\Users\student> type file1.txt | sls Favorite
PS C:\Users\student> type file1.txt | sls d-r
```

However, the output is different if we pipe the dir command.

```
PS C:\Users\student> dir | sls Favorite
PS C:\Users\student> dir | sls d-r
```

You can use the Get-Help cmdlet to both search for available cmdlets and get help pages for those cmdlets. When you type **Get-Help** by itself, it shows you how to use the command. Notice there are two aliases available, man and help. The examples above switch midway through to show that you can use the full cmdlet name or alias. Try some of the ways we have listed above to search through the available cmdlets. Wildcards work just fine, both by themselves as well as before and after strings to find partial matches. There are only a few cmdlets that start with Select. Let's try to use one. If you pipe to more like we show above, use <spacebar> to scroll one page at a time, and the letter <q> to quit more if you don't want to scroll to the bottom.

The Select-String cmdlet works similar to the Linux grep command. Check out its help page for syntax and aliases. Now try piping text from the text file we created earlier to "grep" for keywords. Any line that contains your string will be printed to the screen. Lines that do not contain that string will not be printed.

Now let's try it with something other than the type command. If we use the dir command, it still works; however, the output is different than when we used the type command, even though the output of both type and dir are the same without piping to Select-String. This is because some cmdlets output everything as a text object, whereas other cmdlets expose only a few Properties that are not string selectable. We'll see this in the next task.

Note that the "dir | sls d-r" in the last command above filters out all files and shows only directories by looking for the file type "d."

EXERCISE 4.1: BASELINING WITH POWERSHELL

LISTING SERVICES

Use the get-help command to learn how to find cmdlets for services.

```
PS C:\Users\student> help *service*
PS C:\Users\student> help Get-Service
```

Notice the Select-String approach fails on the Get-Service cmdlet.

```
PS C:\Users\student> gsv
PS C:\Users\student> gsv | sls "running"
```

Now let's use the Where-Object cmdlet to only show running services.

```
PS C:\Users\student> help Get-Service
PS C:\Users\student> help Get-Member
PS C:\Users\student> gsv | gm
PS C:\Users\student> help Where-Object
PS C:\Users\student> gsv | where Status -like "Running"
```

Search the cmdlets for something to list all the services running on this machine. You should quickly find Get-Service as an option, which has an alias of gsv.

Try running Get-Service or its alias by itself. Notice it shows all services, both running and stopped. Let's try using sls to show only running services like we did on the last page. You'll quickly see that it returns nothing—not exactly what we wanted.

So, let's dig a little more into the Get-Service output. Notice how its help page doesn't list "-Property" as valid syntax. So they don't have a simple way to filter this output. This means we need to use Where-Object to filter based on the output field (called an object property). To get a list of properties for Get-Service, we first have to use the Get-Member cmdlet. Pipe the output of Get-Service to Get-Member to list all the Get-Service.Methods() and Get-Service.Properties available for Get-Service. Now, let's pipe the output of Get-Service to Where-Object; so specify which Get-Service.Property we want to use. For the Where-Object syntax, we need to put the comparison in braces and use the -eq to do an "equals" comparison. If you get the syntax correct, you should get a list of running services. Try modifying that to show only stopped services. Now try it by searching for service names.

EXERCISE 4.1: BASELINING WITH POWERSHELL

LISTING NETWORK CONNECTIONS

Use the `Get-NetTCPConnection` cmdlet to list all network connections.

```
PS C:\Users\student> help Get-NetTCPConnection
PS C:\Users\student> Get-NetTCPConnection
PS C:\Users\student> Get-NetTCPConnection -State Listen
PS C:\Users\student> Get-NetTCPConnection -LocalPort 135
```

Listing network connections is easier.

That cmdlet has a built-in syntax to do filtering of the table. Check out the `Get-NetTCPConnection` help page to see the syntax options. Notice how they have "-Property" as an option? Also, did you notice that this is the first cmdlet we've used that doesn't have an alias? To get the Property names, you can either use the `Get-Member` cmdlet like we did on the last page, or you can guess the Property names are the column titles, which they usually are. Go ahead and try filtering for `State` and `LocalPort`, and then try other fields that you might be interested in. Note: You can use multiple `-Property` options at once to do more complex filters. Try combining "`-State Listen -LocalPort 135`" together to see how that works.

EXERCISE 4.1: BASELINING WITH POWERSHELL

WORKING WITH WMI

Use the Get-WmiObject cmdlet to list all users on the machine.

```
PS C:\Users\student> help Get-WmiObject
PS C:\Users\student> gwmi -List | more
```

Use the Select-String cmdlet to filter the -List output for User classes.

```
PS C:\Users\student> gwmi -List | sls User
PS C:\Users\student> gwmi -Class Win32_UserAccount
```

Use the Get-WmiObject cmdlet to list all installed software.

```
PS C:\Users\student> gwmi -List | sls Product
PS C:\Users\student> gwmi -Class Win32_Product
PS C:\Users\student> gwmi -Class Cim_Product
```

Sometimes, it is just too difficult to get the data you want using raw PowerShell objects. A list of available users on the machine is one such case. If you know how to do it in wmi.exe, you can leverage that in PowerShell by using the Get-WmiObject cmdlet. Check out its help page. Notice the -List option? Try it, but if you don't pipe to more, it will take a while. If you DO pipe to more, just use <spacebar> to scroll one page at a time, and the letter <q> to quit more.

While -List is a great option, the output is too long to read. Luckily, the Select-String cmdlet works here. Let's do this to find all WMI Classes that work with Users. There are a few options, but if you test them all, the Win32_UserAccount is the best one to show a short list of available users.

Now let's try finding a WMI Class that works with installed software. You can try searching for Classes with Software in their name, but Microsoft likes using the term Product for installed software. You should easily spot the Win32_Product Class, which works great. However, note there is a Cim_Product Class as well. If you try it, you might get a few more 64-bit-only applications appearing in the list.

EXERCISE 4.1: BASELINING WITH POWERSHELL

TAKEAWAYS AND RECOMMENDATIONS

Section takeaways

- PowerShell is one of the most useful tools Microsoft has ever created for Windows
- Use it, but also limit who can use it, and audit its use

Recommendations to owner/operators

- Continue learning how to use PowerShell – <https://docs.microsoft.com/en-us/powershell>
- We covered collecting baselines, exporting to files, and comparing files; now put it together
- Explore all the PowerShell scripts out there on GitHub and across the internet

This page intentionally left blank.

Course Roadmap

Day 1: ICS Overview

Day 2: Field Devices and Controllers

Day 3: Supervisory Systems

Day 4: Workstations and Servers

Day 5: ICS Security Governance

1. Introduction
2. Patching ICS Systems
 - Patch Decision Tree
 - Vendors, CERTS, and Security Bulletins
3. Defending Microsoft Windows
 - Windows Services
 - Windows Security Policies and GPOs
 - **Exercise 4.1: Baselining with PowerShell**
4. Defending Unix and Linux
 - Differences with Windows
 - Daemons, SystemV, and SystemD
 - Lynis and Bastille
5. Endpoint Security Software
 - Antivirus and Whitelisting
 - Application Sandboxing and Containers
 - **Exercise 4.2: Configuring Host-Based Firewalls**
6. Event Logging and Analysis
 - Windows Event Logs and Audit Policies
 - Syslog and Logrotate
 - **Exercise 4.3: Windows Event Logs**
7. Connectivity to the Internet
 - Honeypots
 - Attacks on the Perimeter
 - **Exercise 4.4: Finding Remote Access**

This page intentionally left blank.

Defending Unix and Linux

Applicable Standards:

- **NIST CSF v1.1:** PR.PT-3
- **ISA/IEC 62443-2-1:2009:** 4.3.3.5.1, 4.3.3.5.2, 4.3.3.5.3, 4.3.3.5.4, 4.3.3.5.5, 4.3.3.5.6, 4.3.3.5.7, 4.3.3.5.8, 4.3.3.6.1, 4.3.3.6.2, 4.3.3.6.3, 4.3.3.6.4, 4.3.3.6.5, 4.3.3.6.6, 4.3.3.6.7, 4.3.3.6.8, 4.3.3.6.9, 4.3.3.7.1, 4.3.3.7.2, 4.3.3.7.3, 4.3.3.7.4
- **ISA/IEC 62443-3-3:2013:** SR 1.1, SR 1.2, SR 1.3, SR 1.4, SR 1.5, SR 1.6, SR 1.7, SR 1.8, SR 1.9, SR 1.10, SR 1.11, SR 1.12, SR 1.13, SR 2.1, SR 2.2, SR 2.3, SR 2.4, SR 2.5, SR 2.6, SR 2.7
- **ISO/IEC 27001:2013:** A.9.1.2
- **NIST SP 800-53 Rev. 4:** AC-3, CM-7
- **CIS CSC:** 3, 11, 14
- **COBIT 5:** DSS05.02, DSS05.05, DSS06.06

This page intentionally left blank.

UNIX/LINUX IN THE ICS WORLD

Workstations: Limited use in ICS

Servers: Common for high-load servers

- Unix can occasionally be found
 - Older server may have SunOS or SCO UnixWare
 - Newer servers may include Solaris, IBM's AIX, HP-UX, and on occasion one of the BSDs
- Linux is common
 - Red Hat Enterprise Linux most common, with CentOS where you don't need support
 - Novell SUSE, Oracle, and Ubuntu have some presence

Embedded devices: Common for many vendor black-box solutions

- Many technician handhelds have started to migrate to Linux
 - Often using homegrown versions of Linux or specialized variants like Wind River's Linux
 - Some are starting to use modified versions of Android (Linux)
- Network and field devices not needing RTOS

Unix/Linux servers are substantially different than their Windows counterparts

It is still uncommon to see a Linux-based workstation in the ICS world, as the laptop and desktop are still substantially controlled by the Microsoft Windows-based operating systems. The largest exception of this is the up-and-coming use of mobile devices, which will require additional management steps and possibly the implementation of a Mobile Device Management (MDM) and endpoint protection system. These are mostly run on Android (Linux-based) or Apple iOS (Unix-based), but mobile devices are still rare to see in actual production due to a higher security risk.

On the server side, it is different. Unix can be found on older servers all over the ICS server world, mostly deployed several years or even decades ago. On the latest and greatest servers, we often see Linux being required or even automatically installed on "black-box" solutions from manufacturers and vendors, especially on high-load servers. The most commonly deployed Linux server is probably Red Hat Enterprise Linux (RHEL), but both Ubuntu and Novell's SUSE have some presence in the ICS market. These Unix/Linux servers are substantially different than their Windows counterparts, and we will spend the rest of the day comparing and contrasting Windows and Linux to help you understand this difference.

LINUX/UNIX DESIGN PHILOSOPHY

Windows started as a desktop OS

- Single-user platform
- Grew into a server OS (multi-process and multi-users)

Unix started as a server OS

- Multi-process and multi-user
- Grew into a desktop platform

Linux grew out of Unix

- Built with same design philosophy
- Kernel is the primary difference between the two
- User interface and software is mostly the same

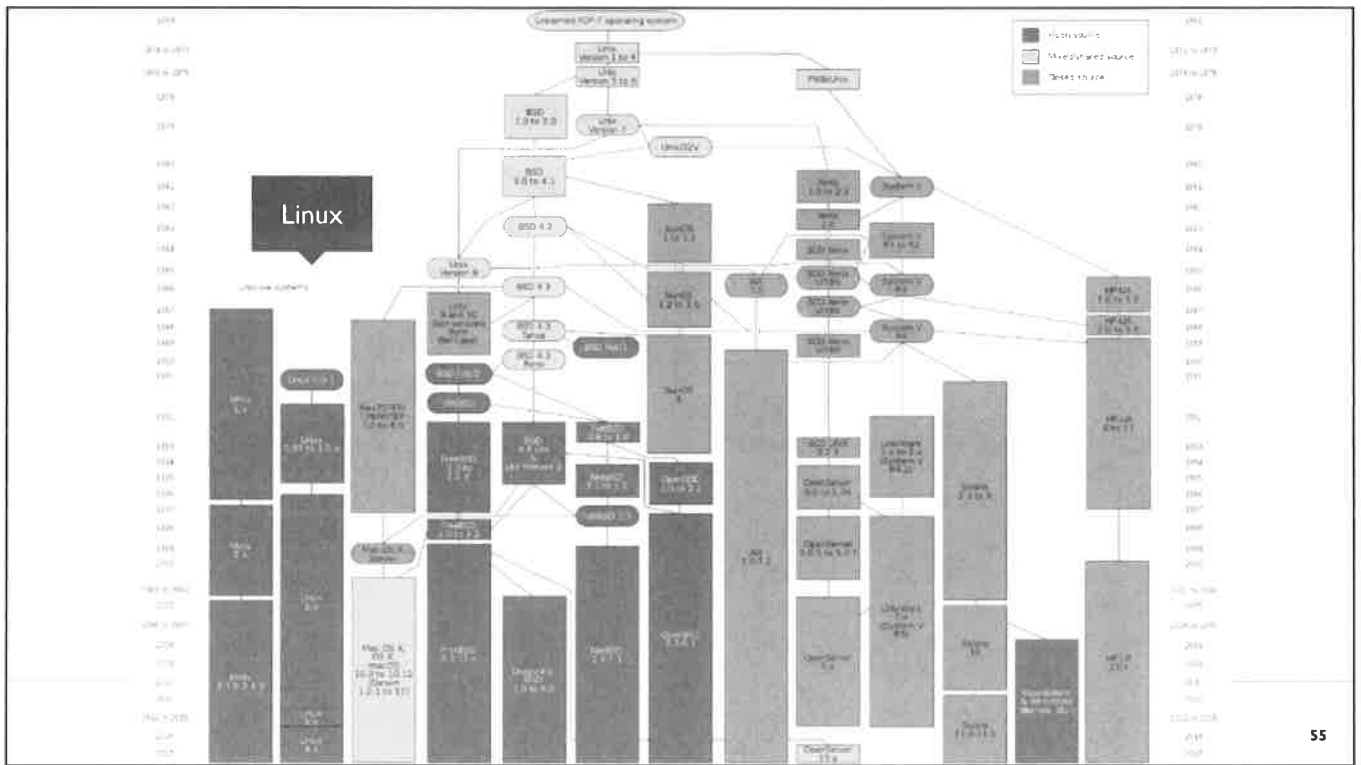
Although it's vital to be comfortable with the most common operating system (OS) in use today, there are substantial (and some say vital) reasons to know about the others as well. There's always room for people to argue over which OS is "better" than another, but the bottom line is they all have advantages, limitations, and liabilities. Because they all get together on big, fast networks, they represent threats to each other that security-minded people need to know about and protect against. Which OS you choose to use depends on what you want to do with a computer, and that's another reason to be familiar with everything that is available. Although Windows has morphed into an OS that one might just as likely find on a server as on a person's desktop computer, it wasn't always that way.

Windows started its commercial life as a desktop OS only, something used by one person to do (basically) one thing at a time until the job was done. Think of a person in an office writing a report or a letter for a project. As things progressed, there came a need for several people to work on a project together, and that required them to all work on (or at least access remotely) the same computer running the same OS. For many reasons, including the need for people to collaborate and to keep all their files together, Microsoft was prompted to develop a server version of Windows that would perform multiple tasks for multiple users simultaneously. So, all of a sudden, you could run a desktop version of Windows on your desktop computer and a server version of Windows on a server, which juggled and synchronized everyone's requests.

Interestingly, Unix developed along the same lines but running in the opposite direction. Unix was conceived in 1969 as an OS that was going to run on a server with many users all doing different things at the same time. The OS needed to manage not only these users and their processes but also all the background programs that made these processes available to the users in the first place. It became stable and dependable, and people started businesses that did nothing but write and maintain a version of Unix as their main source of income. One guy even thought of writing his own version and then shared it with the rest of the world expecting no money in return.

References:

<http://en.wikipedia.org/wiki/Unix>
<http://www.faqs.org/docs/artu/ch02s01.html>




Consider the above chart from Wikipedia. I understand that you probably can't read this chart in your books, so please visit the URL in the references below to see the full-sized original.

Many things combined to give us the Unix landscape we see today. Legal considerations, advances in technology, mass collaboration among thousands of programmers, and the changing needs of users and administrators all shape what is put in software and how it is all put together. By looking at the chart, it is interesting to see the evolution of Unix over the last 40+ years.

Reference:
 Wikipedia Unix Family Tree – <https://goo.gl/Xf8fGg>

ALL SYSTEMS HAVE KERNELS



Manages all other operating system components

- Hardware via installable drivers for hardware interoperability
- Memory management for other processes
- CPU processing time for other processes
- Communication stacks like TCP/IP, USB, HDMI, etc...
- Filesystems access for long-term data storage

Any operating system overview begins with a 50,000-foot view of the key elements. The three key elements of Unix and Linux systems are:

Hardware: A collection of components that house our data and provide means to communicate.

Kernel: A memory resident part of the operating system.

Shell: The portion of the operating system with which users and processes interact directly, which may or may not have a GUI running on top of it. This is similar to the first version of Windows that ran on top of DOS; however, Windows has used an integrated GUI on top of the kernel since Windows 2000 and XP, unless you count servers running Windows Core, which is a new evolution for Microsoft.

The kernel is the most important part of the Unix operating system or any operating system, for that matter. Without the kernel, operating systems do not function. The kernel is responsible for order, instructions to all (both hardware and software), and complete interaction of all system elements. Destroy the kernel and the system becomes useless. Destroy a user or application and only limited functions are affected, yet the rest of the operating system functions.

The kernel of any operating system is generally loaded into memory at bootup. This allows faster interaction between all components, and yet it must have a dedicated space in memory to reside without shifting. What I mean by this is, when loaded, it is normally the first thing loaded and therefore may be predictable. Attackers and software developers alike know. Here is where a common "line in the sand" is drawn, and where the focus on software development security is targeted.

FILESYSTEM LAYOUT

Windows	Linux / Unix
C:\	/ - root folder or root filesystem
\Windows	/boot - kernel /bin - core shell command for normal users /sbin - core shell commands for super users (root)
\Program Files \Program Files (x86)	/usr - higher level OS and third-party apps /var - writable storage for apps /tmp - temporary storage for apps
\Users \Documents and Settings	/root - root's home directory (*nix's administrator) /home - all normal users' home directories
Windows Registry	/etc - configuration directory
Windows APIs	/dev - interface with devices /proc - interface running processes /sys - interface with kernel
D:\ E:\ F:\	All storage devices mounted somewhere inside / Most Linux distributions mount removable media in subfolders of /media

There used to be a lot of variance in where different files were located in different Unix "flavors," but modern Unix variants have settled down to using the same general filesystem layout. The top of the filesystem tree is the "root directory," /. Below this directory are various important subdirectory trees.

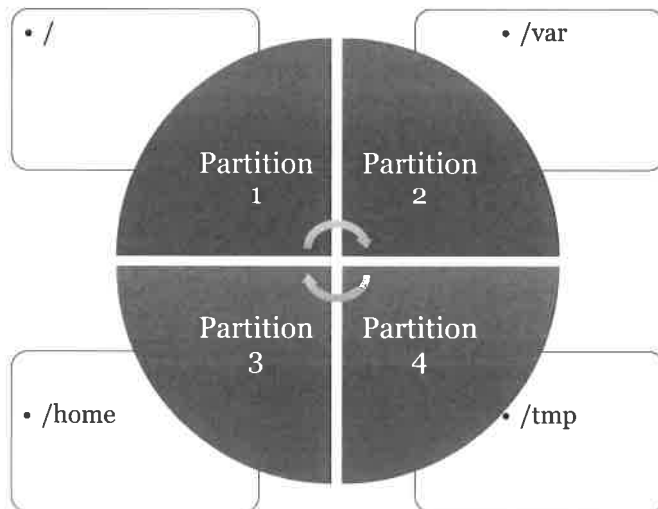
The /dev directory contains the special "device files" that programs running on the system use to communicate with the physical hardware devices controlled by the operating system kernel. Unix systems derived from the AT&T "System V" (SYSV) Unix standard, such as Solaris and HP-UX, often will put these device files into the /devices directory but also usually maintain a /dev directory for compatibility with other systems.

/usr is where most of the critical components of the operating system live, including system binaries, programming libraries and tools, and online documentation. This directory structure can be thought of as being "read-only": After the operating system is loaded, not many changes under /usr should occur unless the operating system is upgraded or patches are installed. /var is where the system keeps frequently-changing data, such as log files and temporary queues for system services like email and printing.

While programs provided with the operating system end up in directories like /usr/bin and /usr/sbin, other programs can be found scattered throughout the system. A standard convention is to put third-party software obtained from the internet into the /usr/local directory. SYSV-derived machines systems like Solaris and HP-UX often will put third-party software (particularly commercial software) into /opt. Different sites may choose to use a different directory-naming scheme for third-party software, however, such as /pkg or /sw.

User home directories often are found under /home. For large Unix networks, though, /home often is an NFS-mounted directory, and /export/home is the directory in which the files physically reside on the file server. Again, many other local conventions exist. Some sites prefer directory names like /users or /u1, /u2, ..., and such.

FILESYSTEMS AND PARTITIONS



Physical drives can be divided into partitions

- Common in Unix/Linux
- Mounted as different folders

Provides enhanced security

- Prevents one from affecting another when out of space
- Can set different security options on each

Configured in `/etc/fstab`

Commands to manage this

- `mount`
- `df` (disk free)

While the Unix filesystem appears to be a single logical entity to the users on the system, it actually is made up of several pieces (called "partitions") that correspond to physical sections of the machine's disk drives. Partitions generally are assigned to critical pieces of the logical directory structure. For example, the system might have one partition for the root filesystem, another for `/usr`, another for `/var`, and so on. During the boot process, all these partitions are "mounted" into their proper place to make the filesystem appear contiguous.

Disk partitions, their mount points, and the options applied to each partition generally are described in a file called `/etc/fstab`. Solaris uses the name `/etc/vfstab` apparently just to be perverse. These filesystem options in the `/etc/fstab` file allow the administrator to specify different security settings for various partitions.

If system log files were to fill up `/var` with spurious log messages, then no more logging could be done until some space was freed up by deleting old log files or removing other data from the `/var` partition. Not being able to log new messages is bad but not as bad as the entire system becoming unusable because the logs consumed all available disk space. By mounting different folders, especially those writable by normal users and processes like `/home`, `/tmp`, and `/var` you can prevent most disk space exhaustion issues interfering with kernel and core OS.

A second reason for partitioning is to make backups easier. Some partitions like `/usr` rarely need to get backed up because they change so infrequently. However, user home directories might need to get backed up every night.

Third, and most important from a security perspective, splitting the Unix filesystem up into different partitions allows the administrator to set different security options on different parts of the filesystem. While the exact partitioning scheme for a machine will vary from system to system and site to site, it is important that the root filesystem, `/usr`, and `/var` be separate partitions to apply appropriate security measures to various OS directories. Non-OS data (user data and home directories, plus application data) and third-party applications not supplied by the OS vendor generally should be put into their own partitions so that they do not "pollute" the OS directories.

THE DF (DISK FREE) COMMAND

```
$ df -h
```

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/sda5	922M	314M	546M	37%	/
/dev/sda6	4.5G	2.3G	2.0G	55%	/usr
/dev/sda9	453M	2.3M	423M	1%	/tmp
/dev/sda7	1.9G	29M	1.7G	2%	/home
/dev/sda8	922M	106M	753M	13%	/var
/dev/sda1	922M	63M	797M	8%	/boot
udev	10M	0	10M	0%	/dev

Administrators can display the currently mounted partitions with the `df` command. `df` stands for "disk free" because the command was actually created to help administrators find partitions who were running out of disk space. The `-h` is an option to view the output in "human" readable format, meaning convert total KB to the most appropriate unit such as MB, GB, TB, and such.

The first output column shows the disk device associated with the partition, whereas the last column shows the piece of the logical filesystem to which the partition corresponds.

The middle columns in the `df` output show the total size of the partition, the amount of disk space currently used, and the amount of free disk space, all reported in 1-kilobyte chunks. The attentive reader will note that the "bytes used" figure plus the "bytes available" amount is not equal to the "total size" number. For performance reasons, standard Unix filesystems typically will reserve 5–10% of the total filesystem for free space. When a partition gets to 90% full, no user except root is allowed to write data into the filesystem. Note that the "percentage of capacity" figure in the second-to-last column of the `df` output takes into account this reserved space. Thus, `df` shows 100% full when the partition actually is only 90–95% full.

Note that `df` shows only currently mounted filesystems. Typically, the machine also will have a partition devoted to "swap space" for the virtual memory system. This is a "raw disk" partition that is not mounted into the logical Unix filesystem. However, this swap partition will have an entry in the `/etc/fstab` file. Information about swap partitions currently in use can also be obtained using the `swapon -s` command.

DOS/WINDOWS-TO-LINUX COMMAND CHEATSHEET

Windows	Linux/Unix	Description
dir /w	ls	Lists contents of current working directory
dir	ls -l	Long listing which shows attributes/permissions
dir /ah	ls -al	Lists hidden and regular files
cd	cd <dir>	Changes the working directory
rename	mv <file1> <file2>	Renames file1 to file2, works for directories, too
	cat <file>	Prints a text file to the terminal (short for concatenate)
type	head <file>	Prints first few lines (defaults to 10)
	tail <file>	Prints last few lines (defaults to 10)
	tac <file>	Prints file backwards
more	less <file>	Prints a text file to the terminal one page at a time
attrib	chmod <attributes> <file>	Changes file attributes/permissions
md	mkdir <newdir>	Makes a new directory
rd	rmdir <empty dir>	Removes a directory; only works if empty
del	rm <file>	Deletes files (use -r to delete a whole directory with files)
copy	cp <file1> <file2>	Copies file1 to file2
ipconfig	ifconfig	Shows current IP address for each interface
	ip addr	
>	<command> > <file>	Saves output of command to a file
>>	<command> >> <file>	Appends output of a command to a file
	<command1> <command2>	Sends output of a command to the input of another command

This slide provides a correlation between Windows and Unix commands for those who may already have familiarity with DOS command-line functions as used in cmd.exe on Windows.

OTHER LINUX/UNIX COMMANDS YOU SHOULD KNOW

Command	Description
<code>pwd</code>	Print working directory; useful if you get lost in the terminal
<code>ps</code>	Show current user's running processes
<code>ps aux</code>	Show all running processes on any Unix/Linux system
<code>ps -Al</code>	Show all running processes on Linux systems
<code>kill <ID></code>	Terminate a process by ID
<code>killall <name></code>	Terminate all processes with the name specified
<code>man <command></code>	Show help manual for command specified
<code>find / -name <file></code>	Search entire hard drive for a file
<code>locate <file></code>	Can use this on most Linux systems to quickly find files from an index
<code>sudo updatedb</code>	Use to update the indexed files (usually runs as a scheduled job)
<code>grep <string> <file></code>	Search through a file for a specific text string
<code>chown <file></code>	Change the owner of a file or directory
<code>sudo <command></code>	Run command as root
<code>sudo -s</code>	Log in as root user
<code>su <user></code>	Log in as different user
<code>sudo route -n</code>	Show local routing table
<code>sudo netstat -ltunp</code>	Show listening TCP and UDP ports in numerical form with process name

This slide contains some of the key Unix commands. Some of the more predominant commands in analyzing incidents and locating items are the "grep" and "find" commands. An example command that combines some of the commands listed above: `>> ps -ef | grep (search term)`. This command will view current processes and perform a search based on your search term. A modified version of the su command that is utilized often is sudo, which enables a user to execute a command as super user.

UPDATING LINUX AND UNIX OS FILES

All Linux distributions provide software repositories

- Includes thousands of third-party software packages
- Most of this is not installed by default but can be added

Each distribution has its own local package management tools

- Red Hat/CentOS: rpm, yum, Add/Remove Software
- Debian/Ubuntu: dpkg, apt-get, aptitude, Ubuntu Software Center
- SUSE/OpenSUSE: rpm, zypper, Yet another Setup Tool (YaST2)
- Mostly a manual process on each machine, often via SSH

Centralized management solutions exist with functionality similar to Microsoft's SCCM

- Each distribution has a for-pay solution, often limited to their distribution
- Open source solutions exist like Puppet, Chef, and Ansible that have commercial support
- Commercial solutions exist for combined Microsoft and Linux patch management
 - IBM BigFix, Ivanti Patch, Symantec Client Management Suite, and others

References:

Red Hat Satellite – <https://goo.gl/c8TsrQ>
SUSE Manager – <https://www.suse.com/products/suse-manager/>
Ubuntu Landscape – <https://landscape.canonical.com>
Puppet – <https://puppet.com>
Chef – <https://www.chef.io>
IBM BigFix for Patch Management – <https://goo.gl/MmPjVz>
Ivanti Patch – <https://goo.gl/tS9pDE>
Symantec Client Management Suite – <https://goo.gl/M5ePWc>

SERVICES... ERRRRR... DAEMONS

All Unix and Linux services are traditionally called daemons

- From Maxwell's Demon, an imaginary being working in the background sorting molecules
- Pronounced as /day-mns/ or /dee-mns/

Most Unix systems (minus the BSDs) use the **System V** (or SysV) initialization scheme

- Most Linux systems prior to 2014 used this as well
- Based on AT&T's Unix released in 1983
- Everything starts with init script (short for initialization)
- The init script defines startup and shutdown process
- Primary configuration file: /etc/inittab

Most Linux systems since 2014 have migrated to **systemd** initialization system

- Modern software suite for management and configuration
- Much more complex than System V
- Uses compiled binaries instead of scripts
- Provides backwards compatibility for System V initialized daemons

First thing to learn about services in Unix and Linux: They aren't called services. They are traditionally called daemons, which can be pronounced as /day-mns/ or /dee-mns/. The name was chosen from William Thompson (Lord Kelvin)'s description of Clerk Maxwell's invisible being that sorted molecules in his thought experiment, which he created to help understand the second law of thermodynamics. Here is an excerpt from Maxwell's book *Theory of Heat*, published in 1874:

"... if we conceive of a being whose faculties are so sharpened that he can follow every molecule in its course, such a being, whose attributes are as essentially finite as our own, would be able to do what is impossible to us. For we have seen that molecules in a vessel full of air at uniform temperature are moving with velocities by no means uniform, though the mean velocity of any great number of them, arbitrarily selected, is almost exactly uniform. Now let us suppose that such a vessel is divided into two portions, A and B, by a division in which there is a small hole, and that a being, who can see the individual molecules, opens and closes this hole, so as to allow only the swifter molecules to pass from A to B, and only the slower molecules to pass from B to A. He will thus, without expenditure of work, raise the temperature of B and lower that of A, in contradiction to the second law of thermodynamics."

After the Linux kernel finishes loading, it immediately starts looking for what "init" processes it needs to start. Think of an init process as being started when the computer boots and continues to run until the system is shut down – a program that provides the most fundamental layer of goodness that goes between the kernel and the user. Init is short for "initialization" and it comes in two styles: SysV, as in Debian/Ubuntu and Red Hat/Fedora, and BSD, which you'll find in FreeBSD and other BSD-type distributions. In the case of newer Linux distributions, System V has been replaced with systemd where compiled services initialize the system instead of bash scripts.

References:

- [https://en.wikipedia.org/wiki/Daemon_\(computing\)](https://en.wikipedia.org/wiki/Daemon_(computing))
- https://en.wikipedia.org/wiki/Maxwell%27s_demon
- https://en.wikipedia.org/wiki/Unix_System_V
- <https://en.wikipedia.org/wiki/Systemd>

SYSTEM V

System V's init starts daemons based on runlevel

- **0:** Shutdown or Halt
- **1:** Single user mode or recovery mode
- **2–5:** Varies by distribution (commonly 5 graphical mode, 2 or 3 for multi-user terminal)
- **6:** Reboot

Daemons are set to start or stop for each runlevel

- Most Unix and Linux systems have graphical tools to configure this
- Unfortunately, few share the same terminal commands

Example of Red Hat/CentOS terminal commands for System V

- **chkconfig:** Used to list and change default states for each daemon on each level
- **service:** Used to start, stop, and restart services immediately
- **runlevel:** Used to see previous and current runlevel for the system
- **telinit:** Used to change runlevels
- See notes below for examples of how to use these

Here are some examples of how to use the System V terminal commands. Note that not all Unix and Linux systems using System V use these commands, so you may need to do some Google searching for your specific system. In the case of older versions of Ubuntu, the `chkconfig` is not installed by default but can be installed with `apt-get`.

```
sudo runlevel                (shows previous and current runlevels, N=none)
sudo telinit 3               (change to runlevel 3)
sudo chkconfig --list        (show list of all daemon states for current runlevel)
sudo chkconfig syslogd on    (enable syslogd to start at current runlevel)
sudo chkconfig syslogd off   (disable syslogd from starting at current runlevel)
sudo chkconfig --level 3 syslogd on (enable syslogd to start at current runlevel)
sudo chkconfig --level 45 syslogd off (disable syslogd from starting at runlevel 4 and 5)
sudo service syslogd start   (immediately start syslogd)
sudo service syslogd stop    (immediately stop syslogd)
sudo service syslogd restart (immediately restart syslogd)
sudo service syslogd status  (show current status of syslogd)
```

systemd

systemd starts daemons based on system state's "target" instead of runlevel

- **poweroff.target**: Shutdown or Halt
- **emergency.target**: Emergency mode for system recovery
- **rescue.target**: Single user mode (root)
- **multi-user.target**: Multi-user terminal mode
- **graphical.target**: Graphical mode
- **reboot.target**: Reboot

Daemons are set to start or stop for each target level

- Most Unix and Linux systems have graphical tools to configure this
- All Linux systems share the same terminal commands :-)

Terminal commands for systemd

- **systemctl**: List current states for all daemons
- **systemctl <command> <daemon>**: Used to start, stop, restart, enable, disable, and get current status
- **systemctl isolate <target>**: Used to immediately change target mode
- **systemctl set-default <target>**: Used to change target mode for all future boots

Enable and disable default boot mode for each daemon/service

Here are some examples of how to use the systemd terminal commands. Note that only modern Linux systems use these commands, basically from 2014-ish on. If you don't have a systemctl command, then you don't have systemd installed. Your Control Things Platform runs systemd, so open it up and try some of these commands. Besides the status command below, you can verify that Apache has stopped by trying to connect to localhost with the browser in Control Things Platform.

systemctl	(show list of all daemon current and default states)
systemctl grep apache	(search the list for any service with apache in name)
sudo systemctl stop apache2	(stop apache2 service)
sudo systemctl disable apache2	(enable syslogd to start at current runlevel)
sudo systemctl status apache2	(check status of apache2 to verify it is stopped and disabled)
sudo systemctl start apache2	(start apache2 service)
sudo systemctl enable apache2	(enable syslogd to start at current runlevel)
sudo systemctl status apache2	(check status of apache2 to verify it is started and enabled)

Reference:

Man page for systemctl – <https://goo.gl/a1xDfZ>

HARDENING LINUX AND UNIX

Linux and Unix don't have unified security profiles like Windows

- Some distributions may have security configuration tools unique to them
- Most settings are handled in configurations of each service or application
- Most distributions provide their own security documentation, but quality varies
- Two auditing tools exist for hardening Unix/Linux: Both are open source!

Lynis

- Works on most *nix systems: AIX, HP-UX, Solaris, Mac OS, all BSD Unixes, and others
- Well maintained and regularly updated, **with no need to install!**
- Audits a system, identifying weak settings with links to descriptions on their website
- Options for commercial support, which include a central management system

Bastille Linux

- Works on Linux and Unix based on System V, but best on Linux as its name implies
- Has not been updated since 2008, but still effective on Linux pre-systemd
- Audits, hardens via a wizard (and un-hardens), and educates you in a single tool

*tell you why.
Does it for you.*

References:

Lynis Homepage – <https://cisofy.com/lynis>

Lynis GitHub Repo – <https://github.com/CISOfy/lynis>

Bastille Linux – <http://bastille-linux.sourceforge.net>

CENTER FOR INTERNET SECURITY (CIS)

How do I harden my other platforms, services, and software?

Consensus guides with steps to make system secure

Free resources after registration

- Hardening Benchmarks (aka guides)
- Top 20 Critical Controls

Member-only content

- Configuration Assessment Tool (CAT)
- Windows security templates
- Hardened virtual images

Description	Items		Score	
	Passed	Failed	Actual	Max
1 Patches, Packages and Initial Lockdown	2	1	7.407	11.111
2 Minimize xinetd network services	6	2	8.333	11.111
3 Minimize boot services	11	10	5.820	11.111
4 Kernel Tuning/Network Parameter Modifications	0	2	0.000	11.111
5 Logging	2	2	5.556	11.111
6 File/Directory Permissions/Access	2	7	2.469	11.111
7 System Access, Authentication, and Authorization	2	9	2.020	11.111
8 User Accounts and Environment	5	7	4.630	11.111
9 Warning Banners	0	3	0.000	11.111
9.1 Reboot	0	0	0.000	0.000
10 Anti-Virus Consideration	0	0	0.000	0.000
11 Remove Backup Files	0	0	0.000	0.000
Overall Score:	30	43	36.240	

SANS

ICS410 | ICS/SCADA Security Essentials

67

Reference:

Center for Internet Security – <https://www.cisecurity.org/>

TAKEAWAYS AND RECOMMENDATIONS

Section takeaways

- Unix and Linux are substantially different than Windows
- Built from the ground up to be managed by a command prompt
- Lends themselves for remote administration via OpenSSH
- Solutions exist to manage some aspects of Unix, Linux, and Windows together
- Significant actions still must be performed on each system type

Recommendations to owner/operators

- Identify all Unix and Linux systems in your infrastructure; they hide well
- Find solutions and methods to unify management of each system type

Recommendations to vendors

- Provide hardening guides for customers managing host Unix/Linux systems
- Harden all black-box solutions that include Unix/Linux
- Track software dependencies and patches of software in your black-box solutions
- **Please stop using outdated versions of Linux in your embedded devices!!!**

This page intentionally left blank.

Course Roadmap

Day 1: ICS Overview

Day 2: Field Devices and Controllers

Day 3: Supervisory Systems

Day 4: Workstations and Servers

Day 5: ICS Security Governance

1. Introduction
2. Patching ICS Systems
 - Patch Decision Tree
 - Vendors, CERTS, and Security Bulletins
3. Defending Microsoft Windows
 - Windows Services
 - Windows Security Policies and GPOs
 - **Exercise 4.1: Baselining with PowerShell**
4. Defending Unix and Linux
 - Differences with Windows
 - Daemons, SystemV, and SystemD
 - Lynis and Bastille
5. Endpoint Security Software
 - Antivirus and Whitelisting
 - Application Sandboxing and Containers
 - **Exercise 4.2: Configuring Host-Based Firewalls**
6. Event Logging and Analysis
 - Windows Event Logs and Audit Policies
 - Syslog and Logrotate
 - **Exercise 4.3: Windows Event Logs**
7. Connectivity to the Internet
 - Honey pots
 - Attacks on the Perimeter
 - **Exercise 4.4: Finding Remote Access**

This page intentionally left blank.

Endpoint Security Software

Applicable Standards:

- **NIST CSF v1.1:** DE.CM-4
- **ISA/IEC 62443-2-1:2009:** 4.3.4.3.8
- **ISA/IEC 62443-3-3:2013:** SR 3.2
- **ISO/IEC 27001:2013:** A.12.2.1
- **NIST SP 800-53 Rev. 4:** SI-3, SI-8
- **CIS CSC:** 4, 7, 8, 12
- **COBIT 5:** DSS05.01

This page intentionally left blank.

TYPES OF ENDPOINT PROTECTION FEATURES

Endpoint protection software goes by many names

- Antivirus software
- Application whitelisting
- Application sandboxes
- File integrity and configuration management
- Host-based firewalls

Many modern-day solutions include many of these functions

- Sometimes as different solutions
- Sometimes as different licenses to a single solution
- Rarely include all those functions

Comparison of different solutions is difficult and time-consuming

- Features are named and described differently
- Marketing efforts obscure what these tools actually do; it is the silver bullet you need
- There is no one-tool-fits-all solution; you must select and trial your short list

Here is a list of vendors offering endpoint protection software. Note: These vendors constantly change, so a few might already be defunct.

- Bitdefender
- Cisco
- Comodo
- CrowdStrike
- Cylance
- Endgame
- ESET
- F-Secure
- FireEye
- Fortinet
- Kaspersky Lab
- Malwarebytes
- McAfee
- Microsoft
- Palo Alto Networks
- Panda Security
- SentinelOne
- Sophos
- Symantec
- Trend Micro

ANTIVIRUS

Antivirus software works like a blacklist for programs

- Contains signatures of known malicious software
- Users can run anything that doesn't match the signature

Antivirus is effective against most malware spreading on the internet

- Once an AV vendor finds a sample, they create a signature
- However, efficiency decreases as signature count increases

Antivirus is useless against targeted attackers

- Zero-day malware by definition hasn't been found by AV vendors
- Existing malware can be modified just enough to break signature match

AV vendors attempt to get around these limitations

- Attempt to write signatures at a broader level to minimize count
- Heuristic detection methods are added to detect common behaviors by malware
- However, attackers always finds ways around these defenses
- You can't create a list of infinite badness...

John Strand and his team at Black Hills Information Security have done several webcasts, blog posts, and presentations around bypassing AV. They are worth the time and effort to review.

<https://www.blackhillinfosec.com/tag/anti-virus>

There are hundreds of other presentations and examples of this on the internet. Try some Google searches for your current endpoint product name with the word “bypass” to see what you can find and do some experimentation on your own systems. It is ALWAYS a rewarding experience to bypass your own security defenses—the hard part is trying to get the resources to address that bypass.

APPLICATION WHITELISTING

Unlike AV, application whitelisting enumerates all the goodness

- Cryptographically signed and/or hashed binaries for each application
- Each are verified before execution
- Allows for whitelists of identified applications

Whitelisting has stronger defenses against attackers

- Helps prevent file-based malware from starting on systems
- Helps to limit malware from establishing permanence through reboots

Whitelisting still isn't a silver bullet

- Doesn't help prevent shellcode from being inserted into a running process
- Struggles at preventing malware written in scripting languages
- Whitelists must be re-created with each patch or program update, often taking shortcuts like W-L folders

Whitelisting is becoming more popular and easier to deploy in Windows

- Windows comes with native capability in the form of **AppLocker**
- Many third-party solutions provide enhanced management features
- Whitelisting is even easier in static environments like control networks, especially Levels 1 and 2...
- Whitelisting is greatly lacking in Unix/Linux and often not available or effective

References:

Microsoft Software Restriction Policies – <http://technet.microsoft.com/en-us/library/bb457006.aspx>

Microsoft AppLocker – <http://technet.microsoft.com/en-us/library/hh831440.aspx>

APPLICATION SANDBOXING

Considered an advanced Mandatory Access Control (MAC) feature

Restrict a running process to certain OS interactions

- Users and services
- Files and folders
- Network input/output
- System calls
- Registry and/or other configuration sources
- Think permissions for applications...

Software Restriction Policies (SRP) and **AppLocker** provide very basic features in Windows

Linux has two of the most powerful solutions built in, depending on the distribution

- **SELinux** (most powerful and flexible, but the most complex)
- **AppArmor** (good flexibility and good ease of use)

FreeBSD Unix and Mac OS include **TrustedBSD**, which performs this function as well

References:

Microsoft Software Restriction Policies – <http://technet.microsoft.com/en-us/library/bb457006.aspx>

Microsoft AppLocker – <http://technet.microsoft.com/en-us/library/hh831440.aspx>

Linux Kernel Security – <http://www.cyberciti.biz/tips/selinux-vs-apparmor-vs-grsecurity.html>

FILE AND CONFIGURATION INTEGRITY MONITORING

Intrusion detection through integrity checking

- Creates a secure database of file checksums
- Verifies current checksums with previous checksums in database

Open source solutions

- Tripwire (Linux only for OSS version)
- OSSEC
- PowerShell script by Jason Fossen (<https://goo.gl/EF24vE>)

Tripwire's enterprise offering includes central management server and configuration monitoring

Some commercial endpoint protection solutions include these features

- Look for features around configuration management

Integrity checkers are valuable tools that allow administrators to take a check-and-balance look at their systems and be presented with a snapshot of any modifications that may have occurred.

Tripwire is both a commercial and open source integrity tool. It creates a digital snapshot of files and/or directories and places this in a portable database. The database should be maintained off-site and kept secure. Should an incident occur, or if an organization would like to audit systems, Tripwire will provide a quick glance at any changes made to systems since the snapshot was taken. This can allow responders to immediately highlight any significant changes made to a system.

References:

For more suggestions on its use – <https://goo.gl/pAXn9y>

PowerShell script by Jason Fossen (<https://goo.gl/EF24vE>)

CONTAINERS

Linux and Unix have provided `chroot()` or `jail()` to isolate individual processes to a single folder tree

- If attackers compromise an application, they have only limited access to the filesystem
- Many network services can be run in `chroot()` or `jail()`
- Some network services have built-in `chroot()`: TFTP, FTP, BIND, SSH

This grew into what we know today as container technologies

- Processes run in a virtual container that emulates subsystems of the OS
- Think virtualization of everything except the hardware and kernel
- Can isolate process tree, network, users, memory limits, processor limits, even root isolation
- Containers help minimize lateral movement between processes on the same host

There are various implementations on the market today

- Docker, Solaris Containers, and LXC/LXD (Linux Containers) are popular examples
- Microsoft uses similar technologies in Windows 10 to protect Edge browser and other software
- Microsoft Server 2016 added their own container implementation

The Unix `chroot()` restriction is an application isolation feature that must be enabled on an application-by-application basis.

Application developers may choose to have their program invoke `chroot()`, typically early in the setup and configuration phase of the application before the application starts accepting requests from external users. When an application calls `chroot()`, it is basically specifying to the OS kernel a directory in the filesystem where the application wants to isolate itself. From that point forward, the kernel enforces this isolation and doesn't let the application "see" any of the filesystems outside of the directory that the application specified when making the `chroot()` call. From the application's perspective, it is as if this directory and all the subdirectories and files below it are now the only files and directories on the system. (The system call is named `chroot()` because it effectively "changes the root" of the filesystem from the application's perspective.)

The advantage from a security perspective is that if a vulnerability is discovered and exploited in the application, then the attacker can access only the parts of the filesystem where the application has `chroot()`ed itself. This prevents the attacker from modifying other configuration files and binaries in the "main" operating system directories and compromising the rest of the system. In fact, the attacker's exploit may fail completely because it relies on other operating system binaries that are not present in the `chroot()` directory used by the application. This is particularly useful for applications like BIND and SSH, which are often internet-facing applications and are complex enough that it's likely there are going to be more remote exploits discovered over time. Running these applications with the `chroot()` restriction helps mitigate the potential damage from future security events.

Container technologies from Docker, Solaris, Microsoft, and Linux's LXC/LXD take this to the next level by providing the ability to isolate a process not only to its own filesystem, but pretty much everything above the kernel, including in some implementations its own network stack, process tree, and users.

HOST-BASED FIREWALLS

Most systems have built-in firewalls (aka host-based)

- Allow each system to control its own inbound and outbound traffic
- Can be used to decrease attack surface for services and clients on the system
- Level 0/1 devices can use small transparent firewalls from Tofino, Ultra 3eTI, or others
- Great solution for mitigating systems you cannot patch

Client use cases

- Only allow access to management services like RDP and SSH from IT/ICS admin subnets
- Prevent hosts that should not connect to internet from sending to non-private IPs

Server use cases

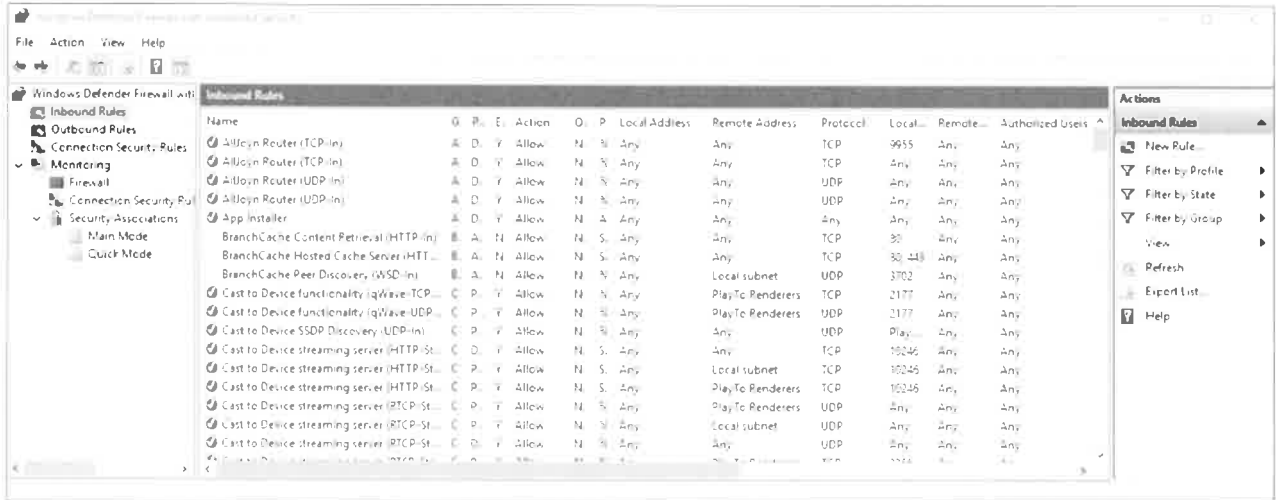
- Only allow certain subnets to access a service
- Prevent normal user subnets from accessing management ports

Control device (Level 0/1/2) use cases

- Only allow communication to PLC from HMI or other specific supervisory system
- Only allow management traffic to vendor X devices from vendor X designated subnets

This page intentionally left blank.

WINDOWS FIREWALL (WF.MSC)



You can access the Windows firewall tool in your Windows configurations or just run wf.msc from the run menu or from a command prompt.

WINDOWS FIREWALL

Windows firewalls can create three types of rules

- Inbound: Traffic entering the machine
- Outbound: Traffic leaving the machine
- Connection Security: Authentication, encryption, and tunneling of any in/out traffic

Inbound and Outbound rules can specify

- Program: Application initiating or receiving the connection
- Protocol and Ports: Such as source and destination TCP and UDP ports
- Scope: Source and destination IP addresses
- Action: Block, allow, or allow if secure (as in a connection security rule)
- Profile: When your computer is connected to a domain, private network, or public network

Firewall rules can be obtained with PowerShell with a few different cmdlets

```
Get-NetFirewallRule | Format-Table -Property Displayname,Direction,Action  
Get-NetFirewallRule -DisplayName 'Remote Desktop - User Mode (TCP-In)'  
Get-NetFirewallRule -DisplayName '*Remote Desktop*' | Get-NetFirewallPortFilter
```

References:

List of PowerShell cmdlets for Windows Firewall – <https://goo.gl/zx1dJo>
Microsoft Windows Firewall Best Practices – <https://goo.gl/SRxicZ>

UNIX AND LINUX FIREWALLS

Unix and Linux firewalls are usually tied to the kernel they run

Unix host-based firewalls

- ipfilter (ipf): AIX, HP-UX, Solaris, and most other Unixes
- ipfw: FreeBSD, Mac <= 10.6
- pf: OpenBSD, Mac >= 10.7, Solaris >= 11.3

Linux host-based firewalls

- ipchains: Rewrite of original ipfwadm in 1998 with kernel 2.2
- iptables: Has become the de facto standard for Linux, providing many benefits over ipchains
- nftables: Improved syntax and unified fw layers; is also backward compatible with iptables

Graphical interfaces for all Unix and Linux firewalls exist

- All are based on the subsystems listed above
- Many are unique for each distribution
- Firewall Builder can build and install configurations for all Linux, Unix, and Cisco firewalls

References:

Unix

ipfilter (ipf) in FreeBSD Handbook – <https://goo.gl/nSp9EW>

ipfw in FreeBSD Handbook – <https://goo.gl/63cxRK>

OpenBSD pf Handbook – <https://goo.gl/9Gu7nz>

Linux

ipfwadm Project – <https://goo.gl/24PSdP>

ipchains Project – <https://goo.gl/bPhY26>

iptables Project – <https://goo.gl/ssm77c>

nftables Project – <https://goo.gl/9PgwiV>


IPTABLES

Built-in host-based firewall for Linux

- Powerful and highly customizable
- Many commercial firewalls have been based on it
- Stateful firewall with NAT capability
- Does not have the ability to specify related processes like Windows firewall

Iptables groups all of its rules into one of three pre-defined chains

- INPUT: All traffic destined to the machine
- OUTPUT: All traffic being sent from the machine
- FORWARD: All traffic being routed through the machine (multi-homed)
- A packet never goes through more than one chain



SANS

ICS410 | ICS/SCADA Security Essentials 81

Because iptables is a complete packet inspection system, you can implement very powerful firewall rulesets. The problem is that all that powerful functionality comes with a lot of complexity. Learning how to write iptables filters to do what you want requires not only learning a lot about how firewalls are configured, but also overcoming the sometimes bizarre and confusing iptables rule language itself.

The rules in chains are what make iptables work. All packets are checked against the rules in the appropriate chain until they find a matching rule that specifies an action to take. If the packet doesn't match a rule, then it will take whatever default policy action is specified for that chain. If no default is specified by the administrator, then it will be ACCEPTED.

References:

Man page for iptables – <https://goo.gl/pLT1Gt>

Official documentation for iptables in multiple languages – <https://goo.gl/8NKy9B>

Oskar Andreasson's iptables tutorial – <https://goo.gl/oMyjhZ>

EXAMPLES OF IPTABLES COMMANDS

Example: sudo iptables -L

-L List all iptable configs

Example: sudo iptables -F

-F Flush... or remove all iptable configs

Example: sudo iptables -P INPUT DROP

-P INPUT Set INPUT chain's default policy to DROP
(also can be ACCEPT but not REJECT)

Example: sudo iptables -A INPUT -s 10.33.66.0/24 -p tcp --dport ssh -j ACCEPT

-A INPUT Append rule to the input chain... or incoming traffic
-s 10.33.66.0/24 If source IP is from the 10.33.66.0/24 subnet (also can be a hostname)
-p tcp If protocol is TCP (also can be udp, icmp, etc...)
--dport ssh And if destination port is 22 (names are mapped to numbers)
-j ACCEPT Then jump to the ACCEPT target (also can be DROP or REJECT)

Basic iptables rules options

-A: Append this rule to a rule chain. Valid chains for what we're doing are INPUT, FORWARD, and OUTPUT, but we mostly deal with INPUT in this tutorial, which affects only incoming traffic.

-L: List the current filter rules.

-m state: Allow filter rules to match, based on connection state. Permits the use of the --state option.

--state: Define the list of states for the rule to match on. Valid states are:

NEW: The connection has not yet been seen.

RELATED: The connection is new but is related to another connection already permitted.

ESTABLISHED: The connection is already established.

INVALID: The traffic couldn't be identified for some reason.

-m limit: Require the rule to match only a limited number of times. Allows the use of the --limit option. Useful for limiting logging rules.

--limit: The maximum matching rate, given as a number followed by "/second," "/minute," "/hour," or "/day" depending on how often you want the rule to match. If this option is not used and -m limit is used, the default is "3/hour."

-p: The connection protocol used.

--dport: The destination port(s) required for this rule. A single port may be given, or a range may be given as start:end, which will match all ports from start to end.

BASIC IPTABLES SCRIPT FOR ENDPOINTS

```
# Flush any existing rules and protect any already established network connections
iptables -F
iptables -A INPUT -p tcp -m state --state ESTABLISHED -j ACCEPT

# Don't interfere with loopback traffic
iptables -A INPUT -i lo -j ACCEPT
iptables -A OUTPUT -i lo -j ACCEPT

# Inbound Rules: Allow Pings and Admin traffic to SSH and HTTPS
iptables -P INPUT DROP
iptables -A INPUT -s 10.33.66.0/24 -p tcp --dport 22 -j ACCEPT
iptables -A INPUT -s 10.33.66.0/24 -p tcp --dport 443 -j ACCEPT
iptables -A INPUT -p icmp -j ACCEPT

# Forward Rules: Allow nothing to route between interfaces
iptables -P FORWARD DROP

# Outbound Rules: Allow Any TCP, UDP, ICMP
iptables -P OUTPUT ACCEPT
```

If you wanted to create firewall rules to control IPv6 traffic, you would need to create a second configuration similar to this using `ip6tables`. There are also `arptables` (for ARP) and `ebtables` (for bridging) commands that are part of the `iptables` project.

You might use these rules on a simply configured system, such as a personal laptop, but for any real application, your firewall rulesets would have to be much more complicated than those shown. You can see how the command-line syntax would tend to discourage users and administrators from making full use of `iptables`.

The good news is that simple GUIs have been developed to allow even novice users to take advantage of `iptables`' functionality. For example, Red Hat-based systems have a simple GUI widget that you can access via the "Applications... System Settings... Security Level..." menu choice. This widget allows you to turn the firewall on the system on and off, and select from a few different services that you want to allow outsiders to access on your machine.

Just to give you a quick impression of the `iptables` command-line interface for building firewall rules, this slide shows a series of `iptables` commands for creating a simple policy that allows incoming basic management via SSH and HTTPS from the admin's subnet and allows all outgoing TCP, UDP, and ICMP traffic.

NFTABLES

Benefits of nftables over iptables

- Simplified syntax inspired by BPF in tcpdump
- Increased scalability, performance, and code maintenance
- Single command for all OSI Layer 3 protocols (iptables, ip6tables, arptables, etables)
- Can define your own table and chain names
- Can handle IPv4 and IPv6 in the same table/chain
- Can perform multiple actions per rule
- Can define dictionaries, maps, and intervals for increased performance and simpler rulesets

Examples of nftables usage

<code>sudo nft list tables</code>	Show a list of all tables
<code>sudo nft list table inet firewall</code>	Show rules for the firewall table
<code>sudo nft flush table inet firewall</code>	Flush all rules in the firewall table
<code>sudo nft flush ruleset</code>	Flush all tables on machine
<code>sudo nft -f <file></code>	Load a firewall configuration file

Reference:

NetFilter wiki for nftables – <https://goo.gl/hqibyT>

BASIC /etc/nftables.conf FOR LINUX ENDPOINTS

```
flush ruleset
table inet firewall {
    chain input {
        type filter hook input priority 0; policy drop;
        ct state established, related accept comment "allow already established connections"
        ct state invalid drop comment "drop invalid packets"
        iifname lo accept comment "allow loopback traffic"
        icmp type echo-request accept comment "Allow pings"
        icmpv6 type {echo-request, nd-neighbor-solicit} accept comment "allow ping and ipv6's ARP"
        ip saddr 10.33.66.0/24 tcp dport {ssh, https} accept comment "allow admin traffic"
        counter comment "count packets dropped by default policy"
    }
    chain output {
        type filter hook output priority 0; policy accept;
        counter comment "count packets accepted by default policy"
    }
    chain forward {
        type filter hook forward priority 0; policy drop;
        counter comment "count packets dropped by default policy"
    }
}
```

table type "inet" includes all IPv4 and IPv6 packets, a huge benefit

Most of the lines above have inline comments, so you should be able to understand most of the lines there. However, one thing you may not be familiar with is ICMPv6's Neighbor Discovery protocol. This serves the same function as ARP but in IPv6. In the one ICMPv6 rule above, you can see we are allowing nd-neighbor-solicit messages. This allows a device to discover the MAC address of a machine on the local network for which it is trying to reach via an IPv6 address.

References:

Wikipedia's article on Neighbor Discovery Protocol – <https://goo.gl/oVuP5f>
YouTube video on Neighbor Discovery Protocol – <https://goo.gl/ZHZsRA>

TAKEAWAYS AND RECOMMENDATIONS

Section takeaways

- Endpoint security software protects and increases attacker detection
- Many solutions exist, including commercial, open source, and OS integrated
- Only as effective as your ability to manage the solution

Recommendations to owner/operators

- Install one or more of these on the majority of your Level 3 systems, especially clients
- Find a workable solution for new Level 0/1/2 deployments before SAT testing
- Look closely at OS integrated solutions for existing Level 0/1/2

Recommendations to vendors

- Ensure new products function correctly with OS integrated and major third-party solutions
- Provide recommended hardening guides, at least for OS integrated solutions
- Add host-based firewall functionality on new embedded products

This page intentionally left blank.

Course Roadmap

Day 1: ICS Overview

Day 2: Field Devices and Controllers

Day 3: Supervisory Systems

Day 4: Workstations and Servers

Day 5: ICS Security Governance

1. Introduction
2. Patching ICS Systems
 - Patch Decision Tree
 - Vendors, CERTS, and Security Bulletins
3. Defending Microsoft Windows
 - Windows Services
 - Windows Security Policies and GPOs
 - **Exercise 4.1: Baselining with PowerShell**
4. Defending Unix and Linux
 - Differences with Windows
 - Daemons, SystemV, and SystemD
 - Lynis and Bastille
5. Endpoint Security Software
 - Antivirus and Whitelisting
 - Application Sandboxing and Containers
 - **Exercise 4.2: Configuring Host-Based Firewalls**
6. Event Logging and Analysis
 - Windows Event Logs and Audit Policies
 - Syslog and Logrotate
 - **Exercise 4.3: Windows Event Logs**
7. Connectivity to the Internet
 - Honey pots
 - Attacks on the Perimeter
 - **Exercise 4.4: Finding Remote Access**

This page intentionally left blank.

EXERCISE 4.2: CONFIGURING HOST-BASED FIREWALLS

EXERCISE DURATION: 15 MINUTES

We are going to play with the firewall on our Control Things Platform virtual machine and use nmap to verify our changes.

OBJECTIVES

- Discover currently open ports
- Use iptables to reject traffic to services
- Test effectiveness of iptables
- Review and modify iptables rules
- Flush all iptables rules

PREPARATION

Start your Control Things virtual machine

This page intentionally left blank.

EXERCISE 4.2: CONFIGURING HOST-BASED FIREWALLS

DISCOVER OPEN PORTS WITH NMAP

```
control@ctp:~$ sudo netstat -ltpn
Active Internet connections (only servers)
Proto Recv-Q Send-Q Local Address      Foreign Address    State       PID/Program name
tcp        0      0 127.0.0.1:3306     0.0.0.0:*          LISTEN      942/mysqld
tcp        0      0 127.0.0.1:631     0.0.0.0:*          LISTEN      1949/cupsd
tcp6       0      0 :::80             :::*               LISTEN      2180/apache2
tcp6       0      0 :::1:631          :::*               LISTEN      1949/cupsd
control@ctp:~$ sudo nmap localhost --reason
<some output removed to fit slide>
Other addresses for localhost (not scanned): ::1
Not shown: 996 closed ports
Reason: 996 resets
PORT      STATE SERVICE REASON
80/tcp    open  http   syn-ack ttl 64
631/tcp   open  ipp    syn-ack ttl 64
3306/tcp  open  mysql  syn-ack ttl 64
```

Looks like both netstat and nmap tell us we have 3 TCP ports open. Notice that the netstat output shows duplicate port 631 entries.

Before we start playing with `iptables`, first we need to find out what services are open on our Control Things virtual machine. There are two ways to do this: One is to simply ask the operating system what it has open with the `netstat` command, and the other way to do this is a basic Nmap scan of our local host.

Type:

```
control@ctp:~$ sudo netstat -ltpn
control@ctp:~$ sudo nmap localhost --reason
```

Remember, the password is "things". The `netstat` command uses four options, which are:

- l Show only network services in a LISTEN state
- t Show only TCP ports
- n Show numerical ports instead of the IANA assigned port names
- p Show process ID (PID) and process name for each port

The `--reason` option for Nmap tells us why it thinks each port is open or closed. Here, we see that we have four ports open. Your virtual machine may have a different number of open ports.

EXERCISE 4.2: CONFIGURING HOST-BASED FIREWALLS

REJECT TRAFFIC WITH IPTABLES

Use the following command to open GEdit to create an iptables script

```
control@ctp:~$ gedit iptables.sh
```

Type the following lines into the GEdit window to create your script

```
iptables -F
iptables -A INPUT -p tcp --dport 22 -j DROP
iptables -A INPUT -p tcp --dport smtp -j DROP
iptables -A INPUT -p tcp --dport 80 -j ACCEPT -s 127.0.0.1
iptables -A INPUT -p tcp --dport https -j ACCEPT -s 127.0.0.1
iptables -A INPUT -p tcp --dport http -j REJECT
iptables -A INPUT -p tcp --dport 443 -j REJECT
```

Save your file and close the GEdit window, which should drop you back to the terminal

Now run the script to start the firewall

```
control@ctp:~$ sudo bash iptables.sh
```

Follow the instructions in the slide above. Notice we are creating firewalls for services we don't even have running.

The first two rules use the DROP action, which should result in no response to requests. The last two rules use the REJECT action, which should result in an ICMP message informing the requestor that those ports are not reachable. There are reasons why you may want to DROP over REJECT; however, that is beyond the scope of this class.

The two ACCEPT rules must be entered before the DROP rules for the same port. Whichever rule for port 80 and port 443 is matched first will result in the specified action. So, traffic from the localhost to those ports will be accepted because that is the first rule that it would match. However, traffic from another computer to those ports wouldn't match until it gets to the REJECT rules.

Also, note that it doesn't matter if you use port numbers or registered service names.

EXERCISE 4.2: CONFIGURING HOST-BASED FIREWALLS

TEST IPTABLES RULES

```
control@ctp:~$ sudo nmap localhost --reason
```

```
Starting Nmap 7.60 ( https://nmap.org ) at 2018-02-04 22:34 MST
Nmap scan report for localhost (127.0.0.1)
Host is up, received localhost-response (0.0000060s latency).
Other addresses for localhost (not scanned): ::1
Not shown: 995 closed ports
```

```
Reason: 995 resets
```

PORT	STATE	SERVICE	REASON
22/tcp	filtered	ssh	no-response
25/tcp	filtered	smtp	no-response
80/tcp	open	http	syn-ack ttl 64
631/tcp	open	ipp	syn-ack ttl 64
3306/tcp	open	mysql	syn-ack ttl 64

Because this traffic is coming from localhost, we expect only port 22 and 25 to be blocked. Port 80 we permitted from local host, and we didn't add rules for 631 and 3306.

```
Nmap done: 1 IP address (1 host up) scanned in 1.25 seconds
```

SANS

ICS410 | ICS/SCADA Security Essentials 91

Now, type:

```
control@ctp:~$ sudo nmap localhost --reason
```

Here, we are running the same Nmap command we ran the first time. Now we see that our access to TCP port 22 and 25 are being dropped, as indicated by the "no-response" indicator. We can still access ports 80 and 443 because Nmap is using the localhost IP address (127.0.0.1) to do the scanning, which our iptables rules permit. Remember, your virtual machine may have a different number of open ports.

EXERCISE 4.2: CONFIGURING HOST-BASED FIREWALLS

TEST WITH SPOOFED ADDRESS

```
control@ctp:~$ sudo nmap localhost --reason -S 127.1.1.1 -e lo
WARNING: If -S is being used to fake your source address...
Starting Nmap 7.60 ( https://nmap.org ) at 2018-02-04 22:37 MST
Nmap scan report for localhost (127.0.0.1)
Host is up, received localhost-response (0.0000040s latency).
Other addresses for localhost (not scanned): ::1
Not shown: 994 closed ports
Reason: 994 resets
PORT      STATE    SERVICE REASON
22/tcp    filtered ssh     no-response
25/tcp    filtered smtp    no-response
80/tcp    filtered http    port-unreach ttl 64
443/tcp   filtered https   port-unreach ttl 64
631/tcp   open     ipp      syn-ack ttl 64
3306/tcp  open     mysql    syn-ack ttl 64

Nmap done: 1 IP address (1 host up) scanned in 1.26 seconds
```

Now that we are spoofing the source IP address, we can see that `iptables` is rejecting our access to HTTP/HTTPS.

SANS

ICS410 | ICS/SCADA Security Essentials 92

With a slight modification of our nmap command, we can tell nmap to spoof our IP address to a different source IP address. Type:

```
control@ctp:~$ sudo nmap localhost --reason -S 127.1.1.1 -e lo
```

Even though 127.1.1.1 is technically a loopback address (127.0.0.0 – 127.255.255.255 all are loopback addresses), our `iptables` rules explicitly permitted only 127.0.0.1. Because of this, we can now see that `iptables` is rejecting our access to ports 80 and 443 with an ICMP "Port Unreachable" response. Also, note that we are getting a filtered response for port 443 (HTTPS), even though there isn't a service listening on that port because we created a rule for it.

EXERCISE 4.2: CONFIGURING HOST-BASED FIREWALLS

REVIEW AND MODIFY IPTABLES

```
control@ctp:~$ sudo iptables -L
Chain INPUT (policy ACCEPT)
target     prot opt source                destination
DROP      tcp  -- anywhere             anywhere             tcp dpt:ssh
DROP      tcp  -- anywhere             anywhere             tcp dpt:smtp
ACCEPT    tcp  -- localhost            anywhere             tcp dpt:http
ACCEPT    tcp  -- localhost            anywhere             tcp dpt:https
REJECT    tcp  -- anywhere             anywhere             tcp dpt:http
REJECT    tcp  -- anywhere             anywhere             tcp dpt:https

Chain FORWARD (policy ACCEPT)
target     prot opt source                destination

Chain OUTPUT (policy ACCEPT)
target     prot opt source                destination
control@ctp:~$ sudo iptables -F
control@ctp:~$ sudo iptables -L
```

Here are all the rules we created in our firewall script.

reject-with icmp-port-unreachable
reject-with icmp-port-unreachable

Don't forget to flush your firewall rules and verify it worked.

(output not shown in slide, but should show all three chains are empty)

We can list all the rules in iptables with the `-L` or `--list` options. We can also remove one rule at a time by typing the same command to enter a rule but changing the `-A` to a `-D` for delete. Finally, if you want to remove all rules, you can use the `-F` to remove all rules for the specified chain. Because we are using only the INPUT chain, we need only one flush rule. If your `iptables -L` command doesn't show any rules, then your firewall is basically turned off and doing nothing.

Type:

```
control@ctp:~$ sudo iptables -L
control@ctp:~$ sudo iptables -F
control@ctp:~$ sudo iptables -L
```

EXERCISE 4.2: CONFIGURING HOST-BASED FIREWALLS

TAKEAWAYS AND RECOMMENDATIONS

Section takeaways

- Make sure you know what services you have running and what their traffic needs
- iptables can reject, drop, or accept traffic depending on the order you enter the rules
- Always test effectiveness of your iptables rules after configuration
- You can remove individual iptables rules with -D or flush them all with -F

This page intentionally left blank.

Course Roadmap

Day 1: ICS Overview

Day 2: Field Devices and Controllers

Day 3: Supervisory Systems

Day 4: Workstations and Servers

Day 5: ICS Security Governance

1. Introduction
2. Patching ICS Systems
 - Patch Decision Tree
 - Vendors, CERTS, and Security Bulletins
3. Defending Microsoft Windows
 - Windows Services
 - Windows Security Policies and GPOs
 - **Exercise 4.1: Baselining with PowerShell**
4. Defending Unix and Linux
 - Differences with Windows
 - Daemons, SystemV, and SystemD
 - Lynis and Bastille
5. Endpoint Security Software
 - Antivirus and Whitelisting
 - Application Sandboxing and Containers
 - **Exercise 4.2: Configuring Host-Based Firewalls**
6. Event Logging and Analysis
 - Windows Event Logs and Audit Policies
 - Syslog and Logrotate
 - **Exercise 4.3: Windows Event Logs**
7. Connectivity to the Internet
 - Honey pots
 - Attacks on the Perimeter
 - **Exercise 4.4: Finding Remote Access**

This page intentionally left blank.

Event Logging and Analysis

Applicable Standards:

- **NIST CSF v1.1:** DE-AE
- **ISA/IEC 62443-3-3:2013:** SR 6.1
- **ISO/IEC 27001:2013:** A.12.4.1, A.16.1.7
- **NIST SP 800-53 Rev. 4:** AU-6, CA-7, IR-4, IR-5, IR-8, SI-4
- **CIS CSC:** 1, 3, 4, 5, 6, 7, 8, 11, 12, 13, 14, 15, 16
- **COBIT 5:** BAI08.02

This page intentionally left blank.

CONFIGURE WINDOWS LOGGING

All windows machines have at least three log files

- **Application:** Logs from applications sending Windows events
- **Security:** Logs related to security events
- **System:** Logs from internal Windows components

Additional events exist on servers depending on role

Research event ID numbers in Google/Bing

- Microsoft's event descriptions are sometimes... lacking
- <http://www.ultimatewindowssecurity.com/securitylog/encyclopedia>

All Windows systems have at least three event logs named Application, Security, and System. You can view the contents of these logs with the Event Viewer snap-in: Administrative Tools > Event Viewer. If you are a domain controller, then you also will have logs named Directory Service and File Replication Service. If you have DNS installed, you'll also have a DNS Server log.

Often, the best way to research an entry, though, is to do a search in Google/Bing on the words *Windows Event ID XXXX*, where XXXX is the event ID number in question. Also try:
<http://www.ultimatewindowssecurity.com/securitylog/encyclopedia/>

The System log mainly is for OS troubleshooting. Here, you'll find entries related to service start/stop/failures, device driver issues, status reports of background maintenance operations, and more. The Application log is where applications and third-party developers can write anything they want. We have little control over what goes here; it is entirely up to the original programmers of the software in question.

WINDOWS EVENT VIEWER (EVENTVWR.MSC)

The screenshot displays the Windows Event Viewer interface. On the left, a tree view shows the hierarchy of log files under 'Event Viewer (Local)'. A callout box labeled 'List of Windows log files' points to this tree. The main pane shows a 'Summary of Administrative Events' table with columns for Event Type, Event ID, Source, Log, and counts for Last hour, 24 hours, and 7 days. A callout box labeled 'Options for selected log file' points to the top right of the main pane. Below the summary is a 'Recently Viewed Events' table with columns for Name, Description, Modified, and Created. A callout box labeled 'Quick summary view by time period and event criticality' points to this table. On the right, an 'Actions' pane lists various options like 'Open Selected Log' and 'Create Custom View...'. The SANS logo is in the bottom left, and 'ICS410 | ICS/SCADA Security Essentials 98' is in the bottom right.

Event Type	Event ID	Source	Log	Last hour	24 hours	7 days
Critical	79	AppModel Runtime	Microsoft-Windows-AppModel-Runtime-Admin	0	0	0
Error	32	AppModel Runtime	Microsoft-Windows-AppModel-Runtime-Admin	0	0	0
Warning				4	23	115
Information				4	243	3,152
Audit Success				2	150	3,051
Audit Failure				0	11	5

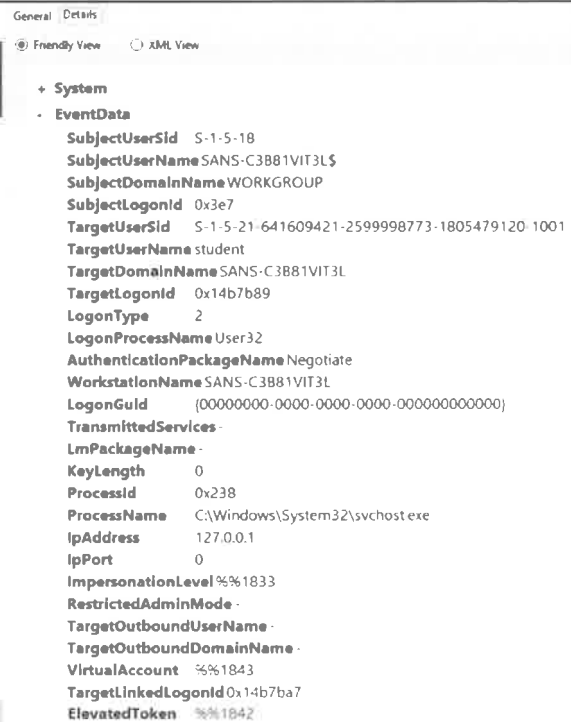
Name	Description	Modified	Created
Windows Event Application	None	2/5/2018 3:29:48 AM	5/21/2018 7:02:02 PM
Windows Event Security	None	5/5/2018 3:30:48 AM	5/21/2018 7:02:02 PM
Windows Event System	None	2/5/2018 3:30:48 AM	5/21/2018 7:02:02 PM
Windows Event Windows	None		
Windows Event Setup	None	11/7/2018 3:30:48 AM	5/21/2018 7:02:02 PM
Custom View: Alerts	Global: All	None	None
Application and Services	None	5/21/2018 7:02:02 PM	5/21/2018 7:02:02 PM

This page intentionally left blank.

WINDOWS EVENT DETAILS

Windows events each contain numerous details

- This even shows the user "Student" logging in
- Depending on the event, you may like the General or Detailed output more



The screenshot shows the 'Details' tab of a Windows Event Viewer window. The event is categorized under 'System'. The 'EventData' section is expanded, showing a list of properties for a system event. The properties include SubjectUserSid, SubjectUserName, SubjectDomainName, SubjectLogonId, TargetUserSid, TargetUserName, TargetDomainName, TargetLogonId, LogonType, LogonProcessName, AuthenticationPackageName, WorkstationName, LogonGuid, TransmittedServices, LmPackageName, KeyLength, ProcessId, ProcessName, IpAddress, IpPort, ImpersonationLevel, RestrictedAdminMode, TargetOutboundUserName, TargetOutboundDomainName, VirtualAccount, TargetLinkedLogonId, and ElevatedToken.

```
General Details
+ System
- EventData
  SubjectUserSid S-1-5-18
  SubjectUserName SANS-C3B81VIT3L$
  SubjectDomainName WORKGROUP
  SubjectLogonId 0x3e7
  TargetUserSid S-1-5-21-641609421-2599998773-1805479120-1001
  TargetUserName student
  TargetDomainName SANS-C3B81VIT3L
  TargetLogonId 0x14b7b89
  LogonType 2
  LogonProcessName User32
  AuthenticationPackageName Negotiate
  WorkstationName SANS-C3B81VIT3L
  LogonGuid {00000000-0000-0000-0000-000000000000}
  TransmittedServices -
  LmPackageName -
  KeyLength 0
  ProcessId 0x238
  ProcessName C:\Windows\System32\svchost.exe
  IpAddress 127.0.0.1
  IpPort 0
  ImpersonationLevel %%1833
  RestrictedAdminMode -
  TargetOutboundUserName -
  TargetOutboundDomainName -
  VirtualAccount %%1843
  TargetLinkedLogonId 0x14b7ba7
  ElevatedToken %%1842
```

Double-click an event icon in one of the logs to bring up its property sheet. Here, you can see the event's date, time, source, user, computer, category, and event ID number. The description field in the middle may also include a few lines or paragraphs of text, including a hyperlink on which you may click to obtain more information from Microsoft about that type of event.

CHANGE WHAT IS LOGGED IN AUDIT POLICY

Audit Policy says what is logged

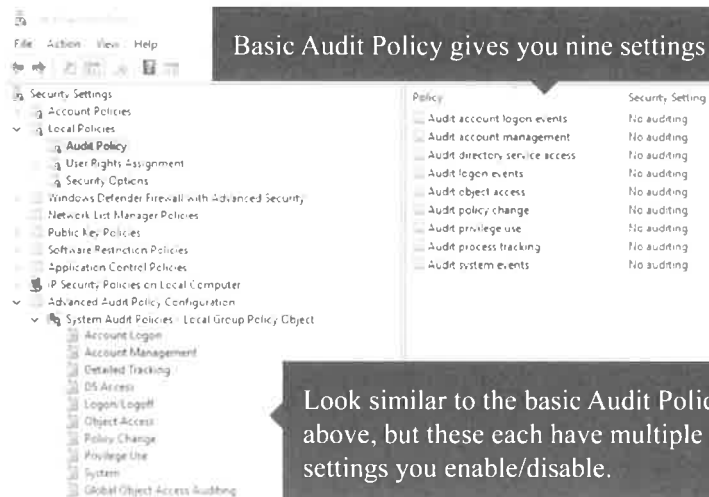
- Managed in LSP or GPO
- Included in hardened templates

Two different Audit Policies

- Basic and Advanced
- Basic policies should not be used in systems after XP

Advanced Audit Policies

- Introduced with Vista/7
- Provide many more settings
- Should override basic settings, but may cause unpredictable issues if both used



To enable logging to the Security log, go to Administrative Tools > Local Security Settings > Advanced Audit Policy Configuration (or in Windows XP, Local Policies > Audit Policy). Here, you can choose to log successful and/or failed events of certain types.

Almost all Audit Policy settings on both servers and workstations can be remotely configured through Group Policy. Audit Policy is set under Computer Configuration > Windows Settings > Security Settings > Local Policies > Audit Policy. Hence, when you need to configure audit settings throughout your domain, Group Policy is the way to do it.

If you'd rather manage Audit Policy from the command line, use the AUDITPOL.EXE tool from the Windows Server *Resource Kit*. It's also built in to Windows Vista/2008/7 and later by default. Starting with Vista, by the way, you can enable or disable special subcategories of audit policies using AUDITPOL.EXE, too. To see your currently activated subcategories of audit policies on Vista/2008/7 or later, run "auditpol.exe /get /category:*."

WHAT OBJECTS SHOULD BE LOGGED?

What are your regulatory requirements?

- These will list specific events you must collect
- If that is your primary reason for logging, start here until you mature your processes

Is intrusion detection your primary driver?

- Microsoft provides good recommendations, but don't set them by default
- Microsoft recommends a minimum set of events to monitor and a suspect list for IR
- Minimum set – <https://goo.gl/y7nreg>
- If cybersecurity is your driver, then start with minimum set until you mature your processes

Other events you can consider adding to the minimum list later down the road

- Windows PowerShell logging
- Groups assigned to local logon
- Logon session creation for network sessions
- User Mode Driver Framework "Driver Loaded" event

Deciding what to log and what not to log is a problem every company faces when they start to collect their logs to a centralized location. In truth, there isn't a fixed answer that everyone can follow. Our desired archive times continually fight against our available disk space. And once we find the right balance there, our desire for more details and artifacts starts to challenge our ability to process those events and scale our architectures to larger sizes. Identifying what is more important to your organization is key to finding your way through these decisions.

What are your regulatory requirements?

Start with any regulatory requirements you have. Many regulations like PCI will require events A, B, and C for X numbers of days on each machine, Y number of months immediately available for search in your central storage, and Z number of year archived that can be retrieved if needed. If your primary purpose for collecting logs is to meet compliance, literally start there, with only the exact events specified in your regulatory requirements. This will allow you to get the solution up and running, letting you measure the effectiveness of your design, the size of your storage, and the ability of your storage server to keep up with the required events.

If intrusion detection is your primary driver?

If you do not have any regulatory drivers, then I'd highly recommend starting with Microsoft's minimum audit policy recommendations for Windows 10. Microsoft's recommended audit policy for Windows 10 is strong compared to their recommendations for previous versions of Windows. Their recommendations for Windows 10 took a stronger emphasis on security than the system administration emphasis used in previous versions. For the first time in workstation history, you can see this in Microsoft's documentation for Windows Event Forwarding for Windows 10 at the following bookmarked sections:

References:

Audit Settings – <https://goo.gl/y7nreg>

Subscription Settings – <https://goo.gl/b3h1V1>

Malware Actadology → Cheat Sheet
about MS hardening
and sec issues

Whichever you use as a starting point, my strong recommendation is not to increase your bare minimum requirements until the solution is fully built and operational, with all planned endpoints in your initial deployment successfully logging to the central server.

Once you are up and running, start tuning!

Tuning your log files seems like an endless process, and it varies from analyst to analyst. The general law of tuning is that you always want more, until the more gets in the way. During an event, your analyst will always think to themselves, "If I only had event X enabled, we could better track the movements of this attacker." The problem is that enabling event X is never-ending, increases your central log volumes each time, and in some cases, may generate hundreds of times more events than you'd ever use. All the low-hanging fruit events that we have recommended in the past to add to XP, Windows 7, and Windows 8, have almost entirely been added to the Windows 10 baseline profile. For Windows 10, there is nothing else I would recommend adding to that baseline that I would say is a hard requirement for all companies to add. However, if you have been running that baseline across all your machines and determine that your solution can withstand a greater number of events, here are the first I'd consider.

- Windows PowerShell logging
- Groups assigned to local logon
- Logon session creation for network sessions
- User Mode Driver Framework "Driver Loaded" event

These events will help identify the first actions attackers usually take once they have compromised a system. However, if you are following the recommendations in the Windows Event Forwarding document for Windows 10 referenced above, those events are already being collected on each machine; they just aren't being sent to the central server.

Use the Suspect Subscription!

Microsoft's recommendation for creating a second event subscription named Suspect is a great solution, allowing you to collect these events as needed on an individual machine basis during an event. This almost eliminates the need to increase your baseline security events. All four of the events I suggested above, plus many more, are included in the Suspect Subscription. And for other versions of Windows, both server and workstation, Microsoft created a list of recommended baselines and a second list called "stronger recommendation" which you should consider as part of a "Suspect" subscription for each of those Windows versions. You can find this document here:

<https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/plan/security-best-practices/audit-policy-recommendations>

Final Thoughts

1. Start with a minimum event collection
2. Verify your solution with a full deployment of minimum events
3. Determine your ability to increase event collection
4. Grow it one event at a time
5. Use a Suspect Subscription as part of your incident response

And don't forget, some of the best ways to determine which events you'd like to add is to go through a few incidents. And if you don't want to wait for them to occur, simulate one yourself by trying to detect the actions of the next penetration test you go through, or organize a Red Team vs. Blue Team exercise between your log analysts and penetration testers. After all, it fits into step one of your incident response process of preparation.

LOG SIZE AND WRAPPING OPTIONS

Full Name:	Security
Log path:	%SystemRoot%\System32\Winevt\Logs\Security.evtx
Log size:	20.00 MB (20,975,616 bytes)
Created:	Wednesday, August 19, 2009 7:21:35 PM
Modified:	Monday, May 10, 2010 5:32:48 PM
Accessed:	Wednesday, August 19, 2009 7:21:35 PM
<input checked="" type="checkbox"/> Enable logging	
Maximum log size (KB):	18014398509481984
When maximum event log size is reached:	
<input checked="" type="radio"/> Overwrite events as needed (oldest events first)	
<input type="radio"/> Archive the log when full, do not overwrite events	
<input type="radio"/> Do not overwrite events (Clear logs manually)	

Each log is finite and can be sized separately from the other logs

Appropriate log size will be determined by the rate of new events and your wrapping options

On Vista/2008/7 and later, logs are compressed XML and can grow to the size of the volume

On XP/2003, don't make the logs larger than 300MB total because of known problems

A security template also can be used to set the maximum size and wrapping options for your event logs. On Windows XP/2003, the logs are stored in the %SystemRoot%\System32\Config folder, but on Vista and later, look in %SystemRoot%\System32\Winevt\Logs\. The names of the log files are self-explanatory. Event logs are assigned maximum sizes beyond which they are not permitted to grow. What happens, then, when they fill up?

To change your Event log size and wrapping (rotation) options without applying a template, open the Event Viewer from the Administrative Tools folder > right-click a log > Properties > General tab. On Windows XP/2003, each log can be increased in size to a maximum of 4,194,240 K (about 4.2 GB), but don't try it. On those older Windows systems, the combined size of all log files should never be set larger than 300 MB because there is a well-known problem with them spontaneously dropping event messages when the logs grow larger than 300 MB in total (KB183097). On Windows Vista/2008/7 and later, there's no such problem because the logs are compressed XML, and they can be set to the maximum size of the volume if desired.

When setting the maximum size, Windows does not reserve the necessary free space. Rather, the log will grow on the filesystem as necessary up to the maximum size or when the drive runs out of free space (whichever comes first). 1 MB of log space holds approximately 7,500 events. How big should the logs be? Well, set them each initially to at least 50 MB, but the size really is determined by your wrapping requirements. You can also investigate archiving the log files to remove the files from the system. There are numerous ways to perform this task with offline storage vaults, streaming to an archiving server with write once / read many media, relying on the backup media for older log file access, and long retention availability for log file aggregators.

UNIX AND LINUX LOGS WITH SYSLOG

Unix and Linux have both standardized on syslog for events

- Records similar types of events as Windows
- Has less detail per event – both a good and bad thing
- Syslog is also the de facto standard network protocol for passing events

There are several different syslog daemons out there

- syslogd, rsyslog, syslog-ng, etc...
- Each varies slightly in configuration and feature set

Each Unix/Linux system uses different configurations, but three most common log files are:

- `/var/log/messages` Valuable messages without debug and non-critical
- `/var/log/syslog` Includes most log files generated, overlaps with messages file
- `/var/log/auth.log` Encrypted log of success/failed authentications
- Many other files are usually configured for different needs of each system

Most Unix/Linux systems also use the logrotate daemon to automatically manage log files

- Unlike windows, can rotate based on size and time frame

The messages log (syslog) can record major events that take place on the system. Some major events that the administrator would want to be aware of would be an attempt to switch the user's privileges (su) to root as well as failed login attempts. As these messages are aggregated and sent out to monitoring and alerting systems, consider obscuring the system names and event types as the information in some of these messages could be useful to an adversary.

MESSAGES

There are five fields present on each line of a syslog-generated file

- 1** Date
- 2** Time
- 3** Hostname of Originating Machine
- 4** Origination Program Name
- 5** Message sent to the Syslogd (follows ":")

```
Jan 25 16:53:38 mjackson PAM_pwdb[446] (login) session opened  
for user root by (uid=0)  
Jan 25 16:53:38 mjackson login[446]: ROOT LOGIN ON tty1
```

PAM = Pluggable Authentication Module

Messages logged by syslog include these five fields: Date, time, hostname, originating program, and the message with instructions for the syslogd.

SYSLOG.CONF

Log routing format

facility.severity target

Example of a syslog.conf

```
auth,authpriv.* /var/log/auth.log
*.*;auth,authpriv.none /var/log/syslog
*.=info;*.=notice;*.=warn; /var/log/messages
#cron.* /var/log/cron.log
daemon.* /var/log/daemon.log
kern.* /var/log/kern.log
mail.* /var/log/mail.log
user.* /var/log/user.log
*.emerg /dev/console
*.alert,crit,err root
auth,authpriv.* @siem.localdomain
```

- all authentication events
- all events except auth
- all higher-level severity
- disabled log entry for cron
- log all daemon events
- log all kernel events
- log all mail events
- log all user events
- emergencies to all logged in
- major to root if logged in
- sent auth to siem server

Many Unix systems use syslogd to control log files. Syslogd uses a configuration file located at /etc/syslog.conf. The syslog.conf file contains information on logs and locations of the logs on the system. The default syslog UDP port is 514.

Syslogd follows the configuration settings provided in syslog.conf, and based on these settings will log events from a variety of system facilities such as mail, auth, and cron. Events occurring in those system facilities will be logged to the location specified in the configuration file. Syslogd can log to a number of locations; examples of these locations include the system console, a local log file, and a remote host.

This screen shows an example of a syslog.conf file. You'll notice two main fields: The selectors in the column on the left and the actions in the column on the right. The selector field has two parts, the **facilities** and the **level** of priority, for each action.

Facilities

The **facility** merely specifies how the message was produced in the syslog file. On the previous page, several logs were generated by the mail programs such as imap, pop, or smtp. Those lines may be tagged with the "mail" keyword.

Levels

The second part of the selector is the **priority level**. The levels, listed here in descending order of priority, indicate the level of alert of the log message.

Actions

The actions column specifies how specific messages should be handled. Low-priority informational messages might be stored to a file, whereas critical messages might be sent to a terminal or output to a printer. The administrator has control over specifying how the messaging information should be handled.

CENTRALIZED LOGGING

Protects against log wiping and simplifies ability to correlate and analyze

- Start with: Firewall, Proxy (may not be in ICS), DNS, DHCP, Netflow
- Active Directory (Domain Admin create/delete, crashed process, failed login, remote logins)

If you have space to log data. also if you have analyze ability

SIEM (Security Information and Event Management) Solutions

- Two of the leading commercial solutions, Splunk and IBM QRadar, offer free low-volume tiers
- LogRhythm and McAfee and many others have great solutions
- Elastic Stack (Elastic Search, Log Stash, and Kibana) is a highly scalable free solution

Log aggregation is often based on a variation of the syslog protocol

- Native support for syslog is on most systems, including ICS devices
- However Microsoft uses an agent on each Windows host or on a Windows Event Collector
- Each SIEM solution has its own agents that can be used when native support is not available

SIEM solutions are not just about storing the logs

- Real-time correlation and alerting of events from various sources
- Analyst interface for threat hunting, security incidents, and forensics

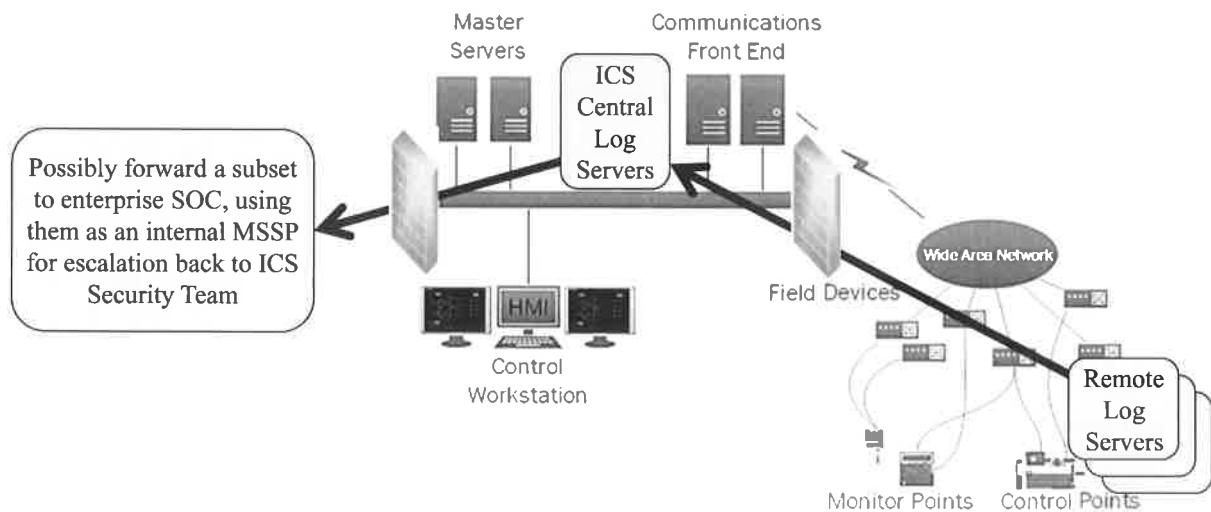
The main advantage to centralized logging is that it makes it difficult for a remote attack to wipe or otherwise corrupt the system logs. Any logs generated from an attack will be sent immediately to another machine, which will store the data. Assuming the syslog server does not, in turn, get hacked, the information will remain there to be discovered by the system administrator and can be used in the recovery process.

The only protection against a syslog client that is sending a lot of information is to have a syslog server that has a lot of disks to store that information. Logs can be rotated quicker if needed, and several partitions can be used for logging to segment the risk of bringing down all the logging in an environment. Syslog does have some built-in mechanisms to aggregate repeated messages. (That is, syslog messages will show strings like "last message repeated X times.") In the end, though, the only protection against a malicious user causing logs to fill up is to use a firewall to block the attacker from reaching the environment, rotate logs quicker, or have enough disks to weather the attack.

The consequence of having centralized system logs is that all the logs are now residing in one place. Many logs will contain information that may provide valuable clues as to what is important in an environment. Some logs may even contain privileged information. With all this information in one place, that machine now becomes a valuable machine to target. The central syslog server becomes a machine that is critical to protecting. It should run only the system logging service to receive messages and have an SSH daemon that will accept only a limited number of logins from a small set of IP addresses belonging to administrators.

With all the system logs being in one place, it makes it easy to have various log-alerting programs (logwatch, logsurfer, and swatch) to run only in one place and send only one set of alerts for all the events in an environment. For instance, some log-scanning programs will scan log files and send daily reports of the notable events on a system. Such a program would do so for each machine it is installed on to give a picture of the entire environment. Running the program on a central log server, however, will generate only one single report for the entire environment and all the systems. To track down problems, only one machine needs to be accessed to find error logs, even if the problem spans multiple machines (such as following an email as it travels through an environment).

LOG AGGREGATION FOR ICS



SANS

ICS410 | ICS/SCADA Security Essentials 108

In ICS environments, we should collect our logs in a central location, first in our central control centers, and possibly in a location outside of the control network. This allows control engineers to access the needed data from inside their control environments, but also allows for other security analysis tools and personnel to access this data from outside of the control environment. However, these log servers outside of the control network should be placed in high-security areas in the business network to limit their odds of compromise.

Just like in IT environments, you should develop a baseline of normal activity in log files and monitor for abnormal activity. However, in control networks, it is often just as beneficial to look for an absence of normal log events as to look for anomalous log events. Attacks in control networks that affect control processes often change the "normal" of the logs, or possibly even stop a string of normal events from occurring. Traditional ICS systems should be predictable from a log file perspective, which differs from some of the less predictable IT environments.

You could even consider working with an engineer to identify and send a small subset of data from your historian to your SIEM solution for additional correlation options.

TAKEAWAYS AND RECOMMENDATIONS

Section takeaways

- Log files are a key method of detecting compromise
- ICS networks are less successful at log collection in general
- Lack of log collection is a major reason ICS doesn't detect compromise for 1.5+ years..

Recommendations to owner/operators

- Start collecting logs now to build experience and business process
- Start small, adding machines and event types over time
- Leverage the Elastic Stack and collect your enforcement firewall logs

Recommendations to vendors

- Support all three major RFCs for syslog in your products
- RFC5424 – The Syslog Protocol
- RFC5425 – Transport Layer Security (TLS) Transport Mapping for Syslog
- RFC5426 – Transmission of Syslog Messages over UDP

This page intentionally left blank.

Course Roadmap

Day 1: ICS Overview

Day 2: Field Devices and Controllers

Day 3: Supervisory Systems

Day 4: Workstations and Servers

Day 5: ICS Security Governance

1. Introduction
2. Patching ICS Systems
 - Patch Decision Tree
 - Vendors, CERTS, and Security Bulletins
3. Defending Microsoft Windows
 - Windows Services
 - Windows Security Policies and GPOs
 - **Exercise 4.1: Baselining with PowerShell**
4. Defending Unix and Linux
 - Differences with Windows
 - Daemons, SystemV, and SystemD
 - Lynis and Bastille
5. Endpoint Security Software
 - Antivirus and Whitelisting
 - Application Sandboxing and Containers
 - **Exercise 4.2: Configuring Host-Based Firewalls**
6. Event Logging and Analysis
 - Windows Event Logs and Audit Policies
 - Syslog and Logrotate
 - **Exercise 4.3: Windows Event Logs**
7. Connectivity to the Internet
 - Honeypots
 - Attacks on the Perimeter
 - **Exercise 4.4: Finding Remote Access**

This page intentionally left blank.

EXERCISE 4.3: WINDOWS EVENT LOGS

We are going to explore the Windows Event Viewer and the events generated by Windows.

OBJECTIVES

- Access the Windows Event Viewer
- Track event logs for logons and logoffs
- Editing log sizes and storage behaviors
- Find where the event logs are stored

PREPARATION

Start your Windows 10 virtual machine

This page intentionally left blank.

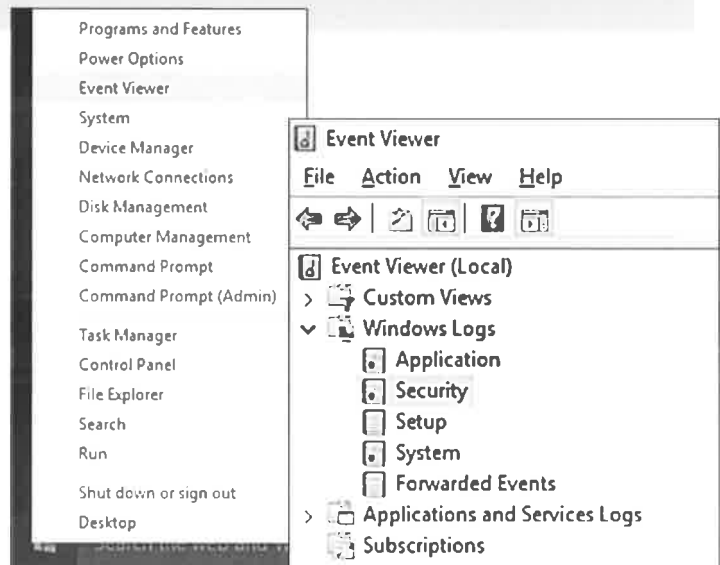
EXERCISE 4.3: WINDOWS EVENT LOGS

OPENING EVENT VIEWER

Right-click the Windows menu icon to access the admin shortcuts

Select Event Viewer from the list

Expand the many folders on the left side of Event Viewer and examine the types of logs each contain, then open the Security event log



There are a couple of different ways you can open Event Viewer on Windows. By far the easiest method in Windows 10 is to right-click the main menu and select "Event Viewer" from the list of admin shortcuts.

When Event Viewer is open, the left side of the window will contain various folders, each holding different types of event files or event views. Expand each of them and examine the types of log files each of them contains. For instance, the "Windows Logs" folder contains the three primary event files for Windows. Of particular interest to us is the Security event file; however, the Application and System event files can contain just as important information for security professions trying to identify an intrusion or recover forensics evidence after a break. The "Application and Service Logs" folder contains a large number of log files for the various applications and services installed on this machine. For ICS, your vendor applications may have their own event files located here that you may want to collect and monitor.

EXERCISE 4.3: WINDOWS EVENT LOGS

TRACKING LOGON EVENTS

- Note the time on your Windows 10 VM in the bottom right of the desktop
- Sign out of your Windows 10 VM, mistype your password, and then log back in to create some events
- Open Event Viewer again, identify the new events based on the event time, and examine those events to answer the questions below (your logs may differ from the screenshot)

Keywords	Date and Time	Source	Event ID	Task Category
Audit Success	10/12/2015 10:19:48 PM	Microsoft Windows security auditing.	4799	Security Group Management
Audit Success	10/12/2015 10:19:48 PM	Microsoft Windows security auditing.	4672	Special Logon
Audit Success	10/12/2015 10:19:48 PM	Microsoft Windows security auditing.	4624	Logon
Audit Success	10/12/2015 10:19:48 PM	Microsoft Windows security auditing.	4624	Logon
Audit Success	10/12/2015 10:19:48 PM	Microsoft Windows security auditing.	4648	Logon
Audit Success	10/12/2015 10:19:44 PM	Microsoft Windows security auditing.	4798	User Account Management
Audit Success	10/12/2015 10:19:44 PM	Microsoft Windows security auditing.	4798	User Account Management
Audit Success	10/12/2015 10:19:44 PM	Microsoft Windows security auditing.	4634	Logoff
Audit Success	10/12/2015 10:19:44 PM	Microsoft Windows security auditing.	4634	Logoff
Audit Success	10/12/2015 10:19:44 PM	Microsoft Windows security auditing.	4672	Special Logon
Audit Success	10/12/2015 10:19:44 PM	Microsoft Windows security auditing.	4672	Special Logon
Audit Success	10/12/2015 10:19:44 PM	Microsoft Windows security auditing.	4624	Logon
Audit Success	10/12/2015 10:19:44 PM	Microsoft Windows security auditing.	4624	Logon
Audit Success	10/12/2015 10:19:44 PM	Microsoft Windows security auditing.	4648	Logon
Audit Success	10/12/2015 10:19:43 PM	Microsoft Windows security auditing.	4647	Logoff

ICS410 | ICS/SCADA Security Essentials 113

The best way to learn how to analyze log files is to perform actions yourself and identify the events generated by your actions. Let's start by simply signing out and signing back in. Before doing this, take note of the time in your Windows 10 VM (which might be different than your host machine) so we know where to start looking in the event logs. After you have logged back in, open Event Viewer and answer the following questions:

What is the difference between event ID 4647 and 4634?

What is the difference between event ID 4648 and 4624?

Who logged on between the time you initiated logoff and completed logoff?

What is the second 4624 logon event for each user?

What do the other events represent, such as 4672 (Special Logon), 4798 (User Account Management), and 4799 (Security Group Management)?

Answers: When you log off, Windows will issue one event for your logoff attempt (4647) and a second event (4634) when the logoff has completed. When logging in, Windows logs events for initial attempts and if credentials are correct, successful logons. When all users are logged off, Windows needs to run the Display Window Manager process to allow people to log on. To do this securely, Windows uses special users like DWM-1 and DMW-2, thus the mysterious extra logon and logoff events. When a user is logged on, Windows will remove "Elevated Token" (the second 4624 event) and then adjust necessary permissions as needed (4798 and 4799 events).

EXERCISE 4.3: WINDOWS EVENT LOGS

CHANGING LOG FILE SETTINGS

While viewing the Security event, click the Properties option on the right side of the window

Examine the options available to manage the size and rollover behavior options

The screenshot shows the 'Log Properties - Security (Type: Administrative)' dialog box. The 'General' tab is selected. The 'Full Name' is 'Security'. The 'Log path' is '%SystemRoot%\System32\Winevt\Logs\Security.evtx'. The 'Log size' is '1.07 MB(1,118,208 bytes)'. The 'Created' date is 'Thursday, September 10, 2015 12:43:29 AM'. The 'Modified' date is 'Monday, October 12, 2015 9:58:02 PM'. The 'Accessed' date is 'Thursday, September 10, 2015 12:43:29 AM'. The 'Enable logging' checkbox is checked. The 'Maximum log size (KB)' is set to '20480'. Under 'When maximum event log size is reached', the 'Overwrite events as needed (oldest events first)' radio button is selected.

Property	Value
Full Name	Security
Log path	%SystemRoot%\System32\Winevt\Logs\Security.evtx
Log size	1.07 MB(1,118,208 bytes)
Created	Thursday, September 10, 2015 12:43:29 AM
Modified	Monday, October 12, 2015 9:58:02 PM
Accessed	Thursday, September 10, 2015 12:43:29 AM
Enable logging	<input checked="" type="checkbox"/>
Maximum log size (KB)	20480
When maximum event log size is reached	<input checked="" type="radio"/> Overwrite events as needed (oldest events first) <input type="radio"/> Archive the log when full, do not overwrite events <input type="radio"/> Do not overwrite events (Clear logs manually)

Note that Windows does not log failed logon events by default, nor does it log when you open a command prompt as admin. This can (and SHOULD) be changed in your Local Security Policy.

EXERCISE 4.3: WINDOWS EVENT LOGS

LOGS FOLDER

Open a command prompt as Admin and type:

```
PS C:\Windows\system32>cd Winevt\Logs
```

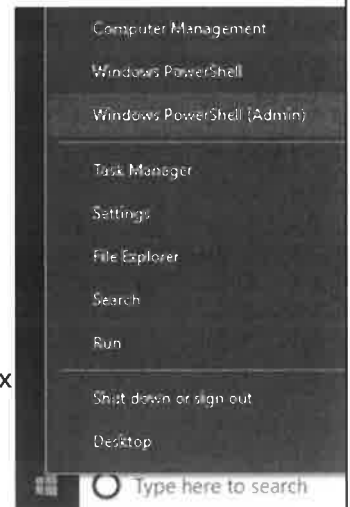
```
PS C:\Windows\System32\winevt\Logs>dir
```

```
Volume in drive C has no label.
```

```
Volume Serial Number is 3C86-B98D
```

```
Directory of C:\Windows\System32\winevt\Logs
```

```
10/12/2015  10:04 PM    <DIR>      .
10/12/2015  10:04 PM    <DIR>      ..
10/12/2015  09:59 PM    1,118,208 Application.evtx
09/10/2015  12:44 AM           69,632 HardwareEvents.evtx
09/10/2015  12:44 AM           69,632 Internet Explorer.evtx
09/10/2015  12:44 AM           69,632 Key Management Service.evtx
...list continued...
```



SANS

ICS410 | ICS/SCADA Security Essentials 115

Now let's find where these event files are physically stored on the filesystem. This is helpful to know because these files should be included in your backup and audited for proper permissions.

To do this, you need to use a command prompt running as Administrator. Right-click the Start menu to access the admin shortcuts and select "Windows PowerShell (Admin)" from the options. When that is running, we need to cd into the C:\Windows\system32\Winevt\Logs folder and run dir to see all the event files.

If you are surprised at how many event files are in this directory, stop and remember all options under the "Applications and Services Logs" folder in Event Viewer. Scroll up and down, and you should find the three primary event files: Application.evtx, Security.evtx, and System.evtx.

EXERCISE 4.3: WINDOWS EVENT LOGS

TAKEAWAYS AND RECOMMENDATIONS

Section takeaways

- For a great resource for understanding which actions create each Windows event and suggested audit rule changes, see <https://www.ultimatewindowssecurity.com/securitylog/default.aspx>

Suggestions for all attendees

- Try a few more actions to see if those actions create events in any of the event files
- Try editing your Advanced Audit Policies to change what is being logged

This page intentionally left blank.

Course Roadmap

Day 1: ICS Overview

Day 2: Field Devices and Controllers

Day 3: Supervisory Systems

Day 4: Workstations and Servers

Day 5: ICS Security Governance

1. Introduction
2. Patching ICS Systems
 - Patch Decision Tree
 - Vendors, CERTS, and Security Bulletins
3. Defending Microsoft Windows
 - Windows Services
 - Windows Security Policies and GPOs
 - **Exercise 4.1: Baselining with PowerShell**
4. Defending Unix and Linux
 - Differences with Windows
 - Daemons, SystemV, and SystemD
 - Lynis and Bastille
5. Endpoint Security Software
 - Antivirus and Whitelisting
 - Application Sandboxing and Containers
 - **Exercise 4.2: Configuring Host-Based Firewalls**
6. Event Logging and Analysis
 - Windows Event Logs and Audit Policies
 - Syslog and Logrotate
 - **Exercise 4.3: Windows Event Logs**
7. Connectivity to the Internet
 - Honey pots
 - Attacks on the Perimeter
 - **Exercise 4.4: Finding Remote Access**

This page intentionally left blank.

Honeypots

Applicable Standards:

- **NIST CSF v1.1:** ID.RA-3
- **ISA/IEC 62443-2-1:2009:** 4.2.3, 4.2.3.9, 4.2.3.12
- **ISO/IEC 27001:2013:** Clause 6.1.2
- **NIST SP 800-53 Rev. 4:** RA-3, SI-5, PM-12, PM-16
- **CIS CSC:** 4
- **COBIT 5:** APO12.01, APO12.02, APO12.03, APO12.04

This page intentionally left blank.

WHAT IS A HONEYPOT?

A tempting treat for bears... and attackers to come snack on

Main purposes

- Early detections of attacks
- Draw in attackers to understand how they break into a system
- To better determine what is attack traffic so defense measures can be improved

There is no authorized activity on a honeypot

- Interaction with the honeypot is accidental or hostile in nature

Advanced technique that is usually deployed after other security measures have been implemented

The ultimate goal of security is to reduce or eliminate risks to an organization's critical assets. Ideally, we prefer to do this by preventing attacks, but one of the key mottos of information security is, "Prevention is ideal, but detection is a must." We must realize that an organization's key resources will be attacked, and we have to be ready to detect the attack as early in the cycle as possible and take advantage of this when it does occur. One way of doing this is with honey-x technology, such as honeypots.

The functionality of honeypots is so diverse that it has been a challenge to define exactly what a honeypot is: Honeypots serve many different purposes for different organizations. Generally, a honeypot is an information system resource whose value lies in the unauthorized or illicit use of that resource. In fact, its value lies in its being misused. The information system resource might be:

- A dedicated server
- A simulated system or state machine
- A service on a selected host
- A virtual server
- A single file with special attributes that is sometimes called a honeytoken

The value in a honeypot is derived from the lack of any authorized activity to the resource. A honeypot resource is never meant for legitimate use; therefore, any use of the honeypot resource is illegitimate and accidental, or hostile, in nature.

When most people hear the term honeypot, they think of a system that you unpatch and put on the internet and that you hope gets broken into. Although this works well for pure research where a site does not have critical systems, it does not scale to a typical DMZ. You do not want your DMZ to be attacked or compromised. If you have critical systems on your DMZ, you need to keep an attacker away. You do not want to draw them in with an unpatched system.

In this case, you would use a honeypot to better understand what is happening on your key systems. A typical web server can get millions of hits a day. Attempting to identify the difference between legitimate connections and attackers is impossible. This is the case unless you have an easy way to discern attack traffic; thus, you have the second use of a

honeypot. In this case, your honeypot is as secure as your production web server and is put on the same network segment. Now, when worms and attackers hit, they attack both your honeypot and your legitimate web server. Because your honeypot has no legitimate uses, you can quickly identify the attack traffic and use that information to build better defenses.

Liability implies you could be sued if your honeypot is used to harm others. For example, if it is used to attack other systems or resources, the owners of those may sue. Liability is not a criminal issue, but civil. The argument for the attacker is that if you the defender had taken proper precautions to keep your systems secure, the attacker would not have been able to harm my systems, so the defender shares the fault for any damage occurred to the defender during the attack.

As with any technology, there is no perfect solution. A honeypot can provide value to an organization if it is deployed correctly. However, it can also cause a decrease in an organization's security by being more attractive to worms or attacks. Therefore, an organization must clearly define the risks it wants to reduce with a honeypot and the requirements for accomplishing this. Then, any deployment can be tested to make sure it benefits the organization.

LOW-INTERACTION HONEYPOTS

A well-placed honeypot in a remote field station or in the control center can be a simple but effective early warning system

- Acts as a canary; any unicast traffic sent to the device could be suspicious as no one should be talking to it

Low-interaction honeypots simply listening for traffic on control protocol ports that log and reject connections can

- Detect unauthorized scanning and probing
- Detect attempted breaches into secured networks
- Provide alerts of detected activity

Honeyports are great examples of low-interaction honeypots

- <http://sourceforge.net/projects/honeyports/>

A well-placed honeypot in a remote field station or in the control center can be a simple but effective early warning system. These can act as a canary by sounding an alarm when any traffic is sent to the device. This traffic should be considered suspicious as no one should be talking to it.

Low-interaction honeypots simply listening for traffic on control protocol ports that log and reject connections can detect unauthorized scanning and probing, detect attempted breaches into secured networks, and provide alerts of detected activity.

Honeyports are great examples of low-interaction honeypots that listen on any port you choose. When someone tries to connect to one of these ports, honeypots simply send out the configured alert or log message without trying to provide any response to the requester. Honeyports can also be configured to perform defensive actions such as a local firewall rule to block the offending IP address. This feature is provided in case you want to install honeypots on a production system to provide automated responses to honeypot alerts.

Reference:

The honeypots project – <http://sourceforge.net/projects/honeyports/>

HIGH-INTERACTION HONEYPOTS

Examples and projects are ongoing for ICS-related honeypots

- SCADA HoneyNet Project
- Control system honeynet code named conpot
- honeynet.org

Analysis of ICS honeypots referenced in notes

- Trend Micro paper
- Black Hat presentation by Kyle Wilhoit

High-interaction honeypots might make sense for researchers and vendors, less so for asset owners

References:

ICS-focused honeynet or honeypots have been in proof of concept or in development for years.

Examples include:

SCADA HoneyNet project – <http://scadahoneynet.sourceforge.net/>

The Control system honeynet code named conpot – <http://www.honeynet.org/node/1047>

Findings from ICS-focused honeynets have been detailed in a 2013 Trend Micro paper:

<http://www.trendmicro.com/cloud-content/us/pdfs/security-intelligence/white-papers/wp-whos-really-attacking-your-ics-equipment.pdf>

And in a Black Hat 2013 presentation by Kyle Wilhoit from Trend Micro at: <https://media.blackhat.com/us-13/US-13-Wilhoit-The-SCADA-That-Didnt-Cry-Wolf-Whos-Really-Attacking-Your-ICS-Devices-Slides.pdf>

TREND MICRO FINDINGS

Findings

- In the second study, it took only 18 hours for the first attack to occur
- Over a 28-day period, there were a total of 39 attacks from 14 different countries
- Twelve of the attacks were unique and could be classified as targeted



From the two reports, Trend Micro first established three honeynet environments and provided data and analysis on the findings. In the second report, it further expanded the honeynet environments to be located in multiple countries and customized the environments to appear genuine in each country's language. In the second project, it also went to great lengths to attempt to determine the source of the attack through the use of some security framework tools that allowed it to obtain information from the attacker's machine that would provide specific location details.

The attacks seen included: Attempts to modify system files; unauthorized access; attempts to modify CPU fan speed on the water pumps; Modbus traffic modification; spear phishing attempts; and malware exploitation. In some of the more interactive attacks, they witnessed exfiltration of data, new exploits, and manipulation of control system values.

Reference:

To read an article discussing the findings from the Trend project, visit: <http://www.controleng.com/single-article/cyber-security-experiment-reveals-threats-to-industrial-systems/e599b7b9842ba6b859176a23486ca017.html>

ADVANTAGES OF HONEYPOTS

All honeypots can reduce challenges of false alarms and data collection

- Early warning alarms
- Help determine true attack traffic

High-interaction honeypots can provide insight into the tactics, motives, and tools of attackers

Can provide additional Defense-in-Depth for organizations

If deployed properly, honeypots play a critical role in the network security arsenal. There is no silver bullet or perfect solution when it comes to network security. Therefore, it is important to understand the advantages of a technology to make sure you deploy it correctly. Some of the advantages of honeypots are:

- Provide insight into the tactics, motives, and tools of attackers
- Reduce challenges of false positives, false negatives, and data collection by determining true attack traffic
- Provide additional Defense-in-Depth for organizations

For organizations and researchers, one of the primary purposes of deploying honeypots is to learn about the tactics and motives of attackers. By utilizing honeypot technology and by watching how attackers compromise systems and what they do after the system is compromised, we can identify the tools they use, their skill levels, and their motives for attacking systems. In addition, on large-volume networks, honeypots can help us focus on the attack traffic, providing a straightforward way to isolate the legitimate traffic.

The question of motive is often useful to organizations on a case-by-case basis. When we detect attempted (or successful) system compromises, we usually have little opportunity to determine whether the attack was the result of random selection or if it was specifically targeted at the victim's organization.

Using a honeypot, researchers can watch the tactics of attackers after they compromise a system. Does the attacker start to scan for more systems to compromise? If so, does he use the same exploit that compromised the honeypot, or does he direct his analysis toward more valuable internal resources, such as database systems?

We mentioned one of the critical factors of honeypots is that the honeypot is deployed without authorized uses. By default, this makes any use of the honeypot accidental or hostile, but always unauthorized. In intrusion detection, we speak of the challenges of false positives (alerts on benign activity) and false negatives (the lack of alerts for hostile activity). When we deal with honeypots, we eliminate almost all the risks of false positives and

false negatives. All data associated with a honeypot (whether it is network traffic, application utilization, or use of the honeypot resource) is logged. Because we can log the data associated with the honeypot, we have a significant advantage over traditional signature-based IDS techniques. Signature-based IDSs generate alerts only for known hostile activity, whereas the honeypot system can capture and identify hostile activity from exploits that are not currently known.

There is no question that honeypots have value and can be used to reduce risks in an organization. As with any technology, if it is deployed correctly, it has value. However, if it is deployed incorrectly, it can cause more problems than it solves. It is also important to deploy a honeypot with other technologies to maximize the benefit.

→ attackers can use honeypots to fool gov and hide what they are doing in fact.

DISADVANTAGES OF HONEYPOTS

Issues with all honeypots

- Fingerprinted honeypots can be used against an organization to cause confusion during attacks
- Honeypots see only traffic sent to it; they do not help identify other compromised systems

Issues particularly with high-interaction honeypots

- Can be a resource burden (not set and forget)
- Improper deployment of honeypots can lead to increased risk of attack
- Legal liability and honeypots

One of the most significant reasons organizations should not deploy honeypots is the risk of misconfigured honeypots, which increases the threat to other production networks. The last thing you want to give attackers is a platform from which to extend their attack into the rest of your network. If attackers can compromise a honeypot system and the rest of the production system isn't sufficiently protected from the perspective of the honeypot, it is likely that the network is in significant jeopardy. An organization wants to keep attackers off of **its network, not give the attackers a foothold behind the firewall.**

Another disadvantage of honeypot technology is the risk of honeypot fingerprinting. In general, you don't want the attacker to identify the target system as a honeypot. Although you might hope this would make an attacker move on to a different network to avoid being detected by another honeypot he wasn't able to identify, this is seldom the case. Attackers use the identity of a honeypot to throw off administrators by spoofing traffic from a legitimate system to the honeypot system, or worse, the attacker feeds the honeypot incorrect information about tactics and motives. While administrators struggle to make sense of the attacks against the honeypot, the attacker **might try to leverage the confusion generated and attack other production systems instead.**

As a method of detecting attack activity against a network, a honeypot is useful only if it is scanned and exploited before an attacker discovers other vulnerable systems. The honeypot sees only traffic sent directly to it; it does not identify other systems that might be compromised before the attacker reached the honeypot system. The honeypot is not an IDS (intrusion detection system) that sees all traffic. It sees traffic going only to that specific system. Just because a honeypot does not see attacker traffic does not mean the network is properly protected.

Honeypots can also be a resource burden for organizations. Honeypot technology is not a set-and-forget option; the deployment of honeypots requires constant monitoring, swift responses, and detailed analysis of attacks on compromised honeypot systems. If your organization is already overburdened with other security measures, deploying honeypot technology will not help reduce that burden. Deploying limited-interaction honeypots will limit the amount of resource burden that is required to maintain the honeypot, but it will still require resources for honeypot monitoring.

One of the more complex issues surrounding honeypot technology is the variety of legal issues that can affect organizations. Serious legal consequences can arise from the use of a honeypot. Organizations are encouraged to consult with legal counsel before deploying honeypots.

In addition, if a honeypot is used by an intruder to attack other systems downstream, the operator of the honeypot might find himself embroiled in litigation by the downstream victims for facilitating the attack and failing to take steps to prevent the use of the system. An operator also might face liability if he learns of attacks against others to whom the operator owes a duty of care but fails to notify the other victims. The honeypot operator also may find himself in the precarious position of having "stolen" information or other contraband (such as child pornography) stored on the system.

Although you may be tempted to do so (especially if the attack is ongoing), do not try to "hack back" to the intruder or attacker. Generally, doing so is illegal. There is no self-defense provision in attacker statutes; this is not a duel.

If you run a honeypot and want to use the collected data legally, pay attention to system clocks on all your systems and follow a strict chain of custody on the data. Honeypots are an excellent place for consent-to-monitor banners, although this technically is infeasible for many ports. One of the many, and best, exceptions to the USA's Wiretap Act is consent. Because such banners are routine on some ports, it can actually make the honeypot appear more like a real production system.

HONEYPOT CHECKLIST/SUMMARY

Honeypot Checklist

- Make sure you have the resources needed to analyze the results
- Validate whether a honeypot is the best solution
- Determine whether a honeypot is increasing your risk
- Identify what information you are trying to obtain from the honeypot

Advanced technique—do everything else first

Honeypots are classified by interaction level and purpose

- High interaction versus low interaction
- Research versus production

Capture and identify unknown threats

Honeypots reduce complications with false positives, false negatives, and data collection

Risk of having attackers use the honeypot if they break the controls

To ensure an organization deploys honeypots effectively, you need to address the following key items:

- Make sure you have the resources needed to analyze the results.
- Validate whether a honeypot is the best solution.
- Determine whether a honeypot is increasing your risk.
- Identify what information you are trying to obtain from the honeypot.

Honeypots are not a fire-and-forget technology. The value of a honeypot is that someone can analyze the data captured by the firewall and use it to increase security. Gathering a lot of information does not typically provide any value if no one is reviewing it; it just uses up valuable resources and increases the risk of compromise, which is not a good business decision. In addition to making sure that you have resources to analyze the results of a honeypot, make sure that it is the best use of that individual's time. For example, if your firewall rulesets have not been validated and your systems have not been hardened, it does not make sense to have administrators review honeypots; their time is better spent on higher-risk tasks and critical security tasks.

Honeypots have a high "cool" factor, and many people want to deploy them. However, you need to make sure a honeypot is the best solution for your problem.

For any high-risk situation, first identify a list of possible solutions, perform a brief cost-benefit analysis, and then choose the most appropriate solution. For example, if a new worm propagates on the internet, it is better to block it at the firewall than allow the traffic into the network so that it can be captured by a honeypot.

The goal of security is to decrease or eliminate risks; you do not want to increase risks. Honeypots are meant to be scanned, connected, and compromised by an attacker, and a compromised system can potentially cause more problems than it solves. Therefore, to ensure its value, it is important that the honeypot is carefully designed and deployed. A honeypot without monitoring or analysis can quickly turn into a liability, especially if an attacker inflicts harm without your knowledge.

Before a solution is deployed, the goals of the solution should be documented. Then the goals should be mapped against a given risk. If you cannot map a solution to a risk, the solution should not be deployed. After the goals have been validated, the deployed system should be tracked against the goals to ensure it accomplishes the goals. If it does not meet the goals, the solution should be modified or eliminated.

The use of honeypots involves advanced techniques. Use honeypots after you have applied other security techniques. Researchers who want to capture new worms or other malware for analysis use honeypots. For the rest of us, honeypots provide a means to get a more detailed insight into attacks attempted against our systems. Normally, firewalls prevent this type of detailed insight, which only becomes available through logs of successful intrusions. Honeypots enable a visibility that comes with penetration, without compromising a production system.

TAKEAWAYS AND RECOMMENDATIONS

Section takeaways

- Honeypots are an advanced technique

Recommendations to owner/operators

- Consider low-interaction honeypots at unmanned outstations
- High-interaction honeypots don't really belong in production

Recommendations to vendors

- Consider high-interaction honeypots to gain threat intel
- Actual products in a safe and highly monitored subnet are the best honeypot
- Have actual attackers find and show you your vulnerabilities...
- Would probably take 3–4 full-time staff to do it right

*Fake user id.
a user who is not used at all.
If you see a log with that
id, there is
sth. strange.*

This page intentionally left blank.

Course Roadmap

Day 1: ICS Overview

Day 2: Field Devices and Controllers

Day 3: Supervisory Systems

Day 4: Workstations and Servers

Day 5: ICS Security Governance

1. Introduction
2. Patching ICS Systems
 - Patch Decision Tree
 - Vendors, CERTS, and Security Bulletins
3. Defending Microsoft Windows
 - Windows Services
 - Windows Security Policies and GPOs
 - **Exercise 4.1: Baselining with PowerShell**
4. Defending Unix and Linux
 - Differences with Windows
 - Daemons, SystemV, and SystemD
 - Lynis and Bastille
5. Endpoint Security Software
 - Antivirus and Whitelisting
 - Application Sandboxing and Containers
 - **Exercise 4.2: Configuring Host-Based Firewalls**
6. Event Logging and Analysis
 - Windows Event Logs and Audit Policies
 - Syslog and Logrotate
 - **Exercise 4.3: Windows Event Logs**
7. Connectivity to the Internet
 - Honeypots
 - Attacks on the Perimeter
 - **Exercise 4.4: Finding Remote Access**

This page intentionally left blank.

Attacks on the Perimeter

Applicable Standards:

- **NIST CSF v1.1:** PR.AC-3
- **ISA/IEC 62443-2-1:2009:** 4.3.3.6.6
- **ISA/IEC 62443-3-3:2013:** SR 1.13, SR 2.6
- **ISO/IEC 27001:2013:** A.6.2.1, A.6.2.2, A.11.2.6, A.13.1.1, A.13.2.1
- **NIST SP 800-53 Rev. 4:** AC-1, AC-17, AC-19, AC-20, SC-15
- **CIS CSC:** 12
- **COBIT 5:** APO13.01, DSS01.04, DSS05.03

This page intentionally left blank.

INFORMATION LEAKAGE LEADS TO ATTACK

Sometimes control systems are exposed to the public internet

- Often by ignorance
- Sometimes by mistake
- Occasionally by valid business reason

These systems can be found through a number of different methods

- Port scanning
- Google searches (aka Google hacking)
- Shodan searches

Sometimes control systems are exposed to the public internet. Often this is from the ignorance of companies and control engineers who don't understand the risk of such configurations. Sometimes control engineers or IT staff accidentally expose control systems to the internet by mistakes, such as making network configuration accidents, incorrect firewall rules, or network WAN link upgrades. And occasionally, businesses identify a valid business need to expose a control system to the internet for remote access for employees or business partner use. In such cases, it is hoped that multiple security defenses are put in place to prevent attacker exploitation. However, sometimes a single security defense such as a password is used, or worse, there are no security defenses at all.

Attackers can find these systems through a number of different methods. The most common is port scanning, by which the attacker uses a tool such as Nmap to look for running ICS services, but this can expose the attacker's source IP to discovery. Another method that attackers can use on internet-facing resources is Google searches. Any control system with a web interface available on the internet will be indexed by Google Googlebot and will be discoverable with simple Google searches. This is also known as Google hacking when attackers use Google to find vulnerable servers. Although Google works only with web-based interfaces on control systems, a similar search engine called Shodan works with many additional services and it allows attackers to search service banners and HTTP headers to identify control systems.

OTHER INFORMATION LEAKED ONLINE

Organization, officers, key personnel

- Positions, email conventions, groups, contacts, and relationships

Third parties/partners

- Engineering firms, industry groups and associations, regulators, and customers

Control systems

- Technologies deployed, databases, HMIs, architectures, IPs, and safety systems

Technology

- R&D, pilots and demonstrations, processes, and IP

Assets

- Facility names, location info, layouts, equipment, accidents, and associated utilities

Supply chain

- Technology suppliers, integrators, maintenance agreements, and contractors

It is important to consider how an adversary may be able to view your organization and its attack surfaces. Publicly available information has supported the targeting process for many attacks. A good adversary can assemble a trove of information and highlight options for delivering attacks or developing exploitation tools. Information found on the internet should serve as a basis for what is initially available to an attacker (i.e. what they are assumed to know about attacking your systems). Defenders should evaluate what information is available and explore your information attack surface. A baseline evaluation sets the stage, and ongoing monitoring of what information becomes available with critical projects and company activities can help determine your current risk. Many risk assessments assume an attacker would have to "know things" (so those scenarios are considered or classified as limited to insiders). You might find that several attack scenarios don't require insider knowledge as the information might already be available. Several studies have uncovered rich datasets that could help an attacker with planning, delivery, and setting priorities for the development of exploitation tools.

Company/facility information, press releases, and websites (including third parties that performed work) can:

- Yield IP addresses of the company's corporate web server, which can be targeted and is often connected (even if through multiple zones) to the ICS
- Include history and accompanying information for better phishing emails
- Discuss previous safety events, lawsuits, or industrial accidents that reveal weak spots or process-specific information
- Include email addresses and names of key personnel

Job descriptions and skill requirements can reveal:

- Detailed information about the exact type of systems an ICS uses
- Product versions for both hardware and software
- Vendor choices and interdependencies
- Positions that are currently unfilled or teams that are growing
- Team structures and positions with trusted access
- Job titles and descriptions as well as an avenue to submit phishing emails or malware-embedded resumes

Regulated industries are often required to provide information to allow for fair market competition (e.g. bulk power systems and the ability of Independent Power Producers to plan, justify, and build plants and connect to the system), or meet requirements for safety, emissions, and reliability. The information, which includes studies, can provide a tremendous amount of detail about the performance or assets, and in some cases, the efforts taken to mitigate certain types of risks. At a strategic level, there are some very high-level observations about an industry that come from studying enforcement information and determining which requirements are proving difficult for compliance.

Suppliers and integrators are proud of their solutions and services, and they tend to discuss their successful implementations. The desire to publish information about implementations will certainly increase if your facility has used new features or components and is an early adopter of new product lines. Working with your corporate communication department and procurement organizations can add a helpful review to remove information that is unnecessary to publish while supporting the desire to announce new achievements. As a security or engineering professional, it is important to present concerns while trying to support the overall goal of the public communication effort.

The information that is publicly available may help a would-be attacker to determine whether your facility is an attractive target for them. Understanding what technology might be in place and starting the process of identifying potential attack surfaces will allow an attacker to determine whether they have the necessary "capability" or to define gaps. Defenders should understand their information attack surface and manage it over time.

GOOGLE HACKING

site:doe.gov login

About 17,100 results

EIA Single Sign On Login Screen

<https://signin.eia.doe.gov/> - Cached

Welcome to the EIA Single Sign On Login System. All Intern... will be unavailable the first and third weekends of the month

FTMS Login

<https://ftms.doe.gov/> - Cached - Similar

Disclaimer. This is a Federal computer system and is the pro... Government. It is for authorized use only. Users (authorized

ISERNet Login - U.S. Department of Energy

<https://www.oe.net.doe.gov/ISERNET/login.aspx> - Cached

ISERNet Login. Welcome to the secure Web site of the Infra... Restoration (ISER) Division. Please log in below.

Creative search queries to find information leakage

- Strings from ICS web interfaces
- Error messages indicating vulns
- Keywords such as "login"
- Restrict searches with "site:"

www.exploit-db.com/google-hacking-database

- Home of the Google Hacking Database
- 1000+ search examples to find vulnerabilities and data leakage



When it comes to internet-exposed web interfaces for control systems or even customer portals that may be leveraged to traverse to a control system, Google hacking is the quickest and most silent method for an attacker to leverage. As the largest search engine on the web, Google has a wealth of information about the various websites exposed to the internet. It has already done all the work for an attacker by indexing all the available web interfaces that are exposed to the internet. All an attacker has to do is make a couple of carefully crafted search terms to find the systems they are looking for. Even by simply searching specific asset owner domains for the word "login", an attacker can manually inspect each search result to see if any resemble control systems.

References:

Google directives can be further researched at:

http://www.googleguide.com/advanced_operators_reference.html

A good source of Google hacking approaches can also be found at:

<https://www.exploit-db.com/google-hacking-database/>

FINDING ICS DEVICES: PORT SCANNING

Easy to do for anybody with an internet connection

Nmap is the best tool and is freely available

Port scanning is the most accurate technique; however, it is also the most obvious and time-consuming

Can cause failures and unpredictable behavior in ICS environments



```
$ sudo nmap -p- -A 127.5.2.3
Starting Nmap 5.21 ( http://nmap.org ) ...
Nmap scan report for 127.5.2.3
Host is up (0.00010s latency).
Not shown: 65529 closed ports
PORT      STATE SERVICE VERSION
22/tcp    open  ssh      OpenSSH 5.9p1 Debian
***
| ssh-hostkey: 1024 bb:16:71...50:37:4c (DSA)
|_ 2048 60:91:1d...ac:6f:dc (RSA)
80/tcp    open  http     Apache httpd 2.2.22
|_html-title: Site doesn't have a title
443/tcp   open  http     Apache httpd 2.2.22
((Ubuntu))
|_html-title: 404 Not Found
502/tcp   open  asa-appl- proto?
5001/tcp  open  apc-agen APC PowerChute agent
8834/tcp  open  ssl/unknown
Network Distance: 0 hops
Service Info: OS: Linux; Device: power-device
```

One of the best tools available for port scanning is Nmap, a very powerful open source port scanner available on all major operating systems. Nmap can be used to identify all services running on a single IP, one service on any group of IPs, or all services on all IPs. Nmap can retrieve service banners from any service offering them and can use their Nmap Scripting Engine (NSE) scripts to do advanced analysis and enumeration of service data. There are several NSE scripts that are ICS-specific, and more are expected to be added in the future.

Port scanning is by far the most accurate method available for attackers to find and identify vulnerable services. The downside for attackers is that it exposes their source IP to discovery and can be very time-consuming if working with large numbers of IPs and ports. However, using Nmap on specific companies can be relatively quick to reveal interesting services for the attacker.

Caution: Tools like Nmap can cause failures and unpredictable behavior in ICS environments. SANS offers additional courses, such as the "Assessing and Exploiting Control Systems" course, which focuses on safe methods of performing penetration testing on ICS environments. An alternative method of port scanning is performing passive network captures for asset owners trying to discover this information on internal ICS networks. This mostly eliminates the risk to ICS environments; however, it also prevents discovery of services with infrequent traffic.

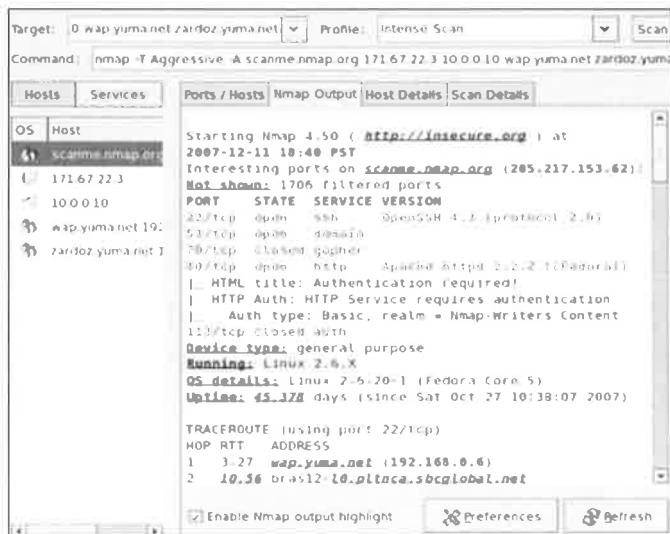
Also, it should be noted that port scanning in some countries is considered illegal without appropriate permission from the owners of the network and connected systems.

shenanigans started nmap for ipv6 in order to gather active ipv6

ZENMAP

Nmap available in a graphical version called Zenmap

Runs on Windows, Linux, and Mac OS X



This page intentionally left blank.

SHODAN.IO

Security geek's version of Google

- Shodan doesn't index contents of webpages
- Shodan indexes service banners and service headers

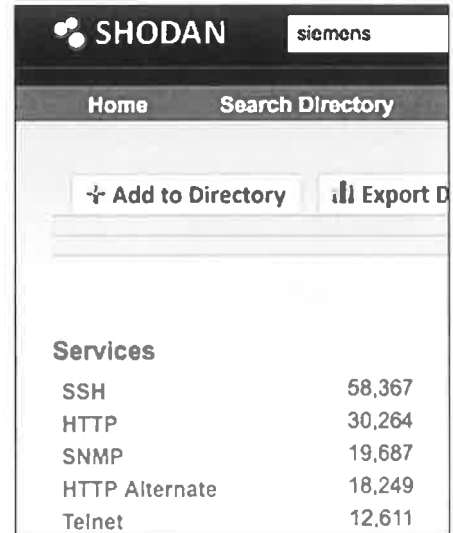
Protocols originally supported:

- HTTP and HTTPS
- Telnet and SSH
- FTP, SIP, and SNMP
- SMTP, POP, IMAP, and NetBIOS

Now scans over 300 ports including common ICS...

- <https://www.shodan.io/explore/category/industrial-control-systems>
- <https://ics-radar.shodan.io/>

ip:204.51.94.0/24 to search your networks



The screenshot shows the Shodan website interface. At the top, there is a search bar with the text 'siemens' entered. Below the search bar, there are two buttons: 'Add to Directory' and 'Export D'. The main content area displays a table of services and their counts:

Services	Count
SSH	58,367
HTTP	30,264
SNMP	19,687
HTTP Alternate	18,249
Telnet	12,611

SANS



ICS410 | ICS/SCADA Security Essentials 139

Google has all the knowledge about internet-exposed webpages, whereas Shodan is the source attackers go to for knowledge of other protocols. Shodan collects banner and response header information for HTTP, HTTPS, Telnet, SSH, FTP, SNMP, SIP, SMTP, POP, IMAP and NetBIOS. This allows attackers (and security defenders) to identify various services running in your environment without needing to do the port scanning themselves. Think of Shodan as an internet resource that provides portscan data for the 10 most common service ports. Shodan doesn't provide anything that Nmap wouldn't provide you, but it just decreases the time and risk associated with running those Nmap scans on the internet. Of particular interest is a new Shodan page that has been set up just for ICS systems. It is located at <https://www.shodan.io/explore/category/industrial-control-systems>

Security researcher Bob Radvanovsky of Project SHINE leveraged SHODAN to find internet-connected ICS devices starting in April 2012. By the end of 2013, Project SHINE had found over 1 million internet-connected ICS devices from over 64 different ICS device manufacturers. The project is finding between 2,000 and 8,000 new ICS devices each day. For more information on Project SHINE, see:

<http://www.tofinosecurity.com/blog/project-shine-1000000-Internet-connected-scada-and-ics-systems-and-counting>

The latest Project SHINE report findings can be found here:

<http://www.slideshare.net/BobRadvanovsky/project-shine-findings-report-dated-1oct2014>

Caution: Use of tools like Shodan may be prohibited by your organization.

SHODAN MAPS

Newly added map features

Geo positions the search results found

Example "Siemens" search



Shodan offers a number of additional features to members. Membership has a one-time-only fee of \$49.00, and it includes access to the API, mapping features, application integration to Metasploit, Maltego, Chrome, and search results with much higher limits (10,000 results).

ATTACK SURFACE MANAGEMENT

Catalog publicly available information

- Ongoing identification of information published
- Manage what information is published

Work with your communications department and procurement organization (identify new projects or modernizations)

Build strong remote access

- Identify people who have access to ICS
- Secure behavior training (job-specific awareness)
- Security "strong points" and process for moving files (e.g. jump servers, administrative controls)

ICS security and business unit management should consult their information security policy and map information classification requirements to ICS-related information. Processes should be established to evaluate information to be published in a public format or to specific audiences of forums. Cataloging what information is already available and developing capabilities to monitor information that becomes available in the future can help an organization to best understand its existing exposure. These efforts can be as simple as setting up Google Alerts with keywords to contracting a security intelligence service that is looking for entity-specific information on the internet.

All cybersecurity efforts begin with people, and that is where many cyber incidents also begin. The development of an overall cybersecurity strategy to protect both the corporate network and plant systems begins with system owners and cybersecurity specialists. The susceptibility of deployed technology can lie purely in the technology, but in many cases, it can be a combination of how the technology is configured or through actions taken by authorized system users and administrators. Many of the recent ICS-focused technical threats have included targeting and delivery techniques that have focused on human and work practice vulnerabilities. Attackers look for easy paths to gain footholds on trusted networks and place themselves within reach of important data and credentials that can help them achieve a freedom of movement and action.

Both corporate system and plant system users need security awareness and base knowledge. There are strong arguments for providing this training based on the type of work being performed and using the language that best aligns to the context of the technology and operation. Developing engineering-focused modules allows security program managers to relate common day-to-day activities to secure and less-secure behaviors.

Purdue Level 4: Plant's Local Business Network

Enforcement between ICS DMZ and Business Networks (Business pulls from or pushes to ICS DMZ)

ICS DMZ - Level 3 to 4
 ICS DMZ - Level 4 to 3
 ICS DMZ - Cloud Access
 ICS DMZ - Remote Access

Enforcement between Control Networks and ICS DMZ (Control pulls from or pushes to ICS DMZ)

Purdue Level 3: Plant-wide Supervisory	Master Servers, Historian, and HMIs	Workstations (per group/role)	Testing/Staging (per system)	Cybersecurity Operations	Jump Hosts (per vendor or group/role)
---	---	----------------------------------	---------------------------------	-----------------------------	---

Enforcement between Cell / Lines and Plant Supervisory (ACL on router / Layer-3 switch or Firewall)

Remote Access – Two Separate Steps

VPN from the internet (or Enterprise network) into the ICS DMZ

- Two-factor authenticates to Enterprise AD server or standalone authentication server
- Successful VPN gives IP from ICS DMZ to VPNed host

RDP/VNC/Citrix to jump host in Purdue Level 3

- Should authenticate to the Plant/Control AD server
- Different jump hosts per vendor or per employee role should be used
- All needed software and data should be pre-installed on jump host
- Solution properly hardened to prevent user from using software on their own host
- Separate scanning solution should be used when user needs to bring data with them

SANS
ICS410 | ICS/SCADA Security Essentials 142

This page intentionally left blank.

TAKEAWAYS AND RECOMMENDATIONS

Section takeaways

- Control systems exposed on the internet can be easily found
- Google, Shodan, and Nmap are the most common methods

Recommendations to owner/operators

- Take time to understand what the public knows about your systems
- Security assessments are a great time to do this

Recommendations to vendors

- Limit access to tools and documentation for your customers
- Don't limit access to your security bulletins, rather limit detail

This page intentionally left blank.

Course Roadmap

Day 1: ICS Overview

Day 2: Field Devices and Controllers

Day 3: Supervisory Systems

Day 4: Workstations and Servers

Day 5: ICS Security Governance

1. Introduction
2. Patching ICS Systems
 - Patch Decision Tree
 - Vendors, CERTS, and Security Bulletins
3. Defending Microsoft Windows
 - Windows Services
 - Windows Security Policies and GPOs
 - **Exercise 4.1: Baselining with PowerShell**
4. Defending Unix and Linux
 - Differences with Windows
 - Daemons, SystemV, and SystemD
 - Lynis and Bastille
5. Endpoint Security Software
 - Antivirus and Whitelisting
 - Application Sandboxing and Containers
 - **Exercise 4.2: Configuring Host-Based Firewalls**
6. Event Logging and Analysis
 - Windows Event Logs and Audit Policies
 - Syslog and Logrotate
 - **Exercise 4.3: Windows Event Logs**
7. Connectivity to the Internet
 - Honeypots
 - Attacks on the Perimeter
 - **Exercise 4.4: Finding Remote Access**

This page intentionally left blank.

EXERCISE 4.4: FINDING REMOTE ACCESS

DURATION TIME: 15 MINUTES

We are going to explore some of the methods attackers use to find our remote access and internet-exposed systems

OBJECTIVES

- Use Google to identify security bulletins
- Use Google to identify vulnerable systems
- Use Shodan-exposed ICS systems
- Use Shodan to find VPNs and other remote access systems

PREPARATION

Open a browser on your host to do this exercise; no virtual machine needed

Note: If the classroom is not equipped with internet, please do this exercise at home



Occasionally, SANS classrooms are not equipped with internet access due to venue limitations. If this classroom is not equipped with internet, please do this exercise at home or any other time that you have internet access.

Caution: Using reconnaissance tools against internet-accessible IP addresses may have some legal ramifications. In these labs, we are performing passive reconnaissance using search tools.

The goal of this lab is to practice using Google and Shodan to find vulnerable systems in your infrastructure. In this example, we use Siemens because of the popularity of its use and the maturity of its security efforts. Siemens' security team has been very proactive in addressing security vulnerabilities found in its products and Siemens releases vulnerability advisories to the public; this will make our lab a little easier than using more obscure vendors and products.

First, we will use Google to identify any existing vulnerabilities with the Siemens HMIs and Siemens S7-1200 PLCs. Next, we will use Google to identify Siemens HMI systems exposed to the internet using the search "Simatic HMI Miniweb on", which is a phrase that appears on the HMI home page. Then, we will use Shodan to identify any Siemens S7-1200 PLCs exposed to the internet.

As a final step of our lab, please use the remaining lab time to identify security issues with ICS products you have used in the past and attempt to find publicly exposed systems using Google and Shodan.

Caution: This lab shows the student a variety of capabilities found through online tools and requires an internet connection. While connected, please perform the labs as instructed and do not use any of the other tools found in the lab distribution environment.

EXERCISE 4.4: FINDING REMOTE ACCESS

USING GOOGLE TO FIND VULNERABILITIES

The image shows two side-by-side screenshots of Google search results. The left screenshot is for the search query "siemens hmi vulnerability" and shows results from ICS-CERT and DHS. The right screenshot is for "siemens s7-1200 vulnerability" and shows results from ICS-CERT and a security lab.

siemens hmi vulnerability

Web Images Maps Shopping More Search tools

About 34,900 results (0.23 seconds)

Siemens Simatic HMI Authentication Vulnerabilities | ICS-CERT
ics-cert.us-cert.gov/advisories/ICSA-11-056-01
Apr 22, 2013 - Overview, ICS-CERT is aware of a public report by independent researchers Billy Rios and Terry McCorkle concerning authentication ...

Siemens SIMATIC WinCC Multiple Vulnerabilities (UPDATE)
ics-cert.us-cert.gov/advisories/ICSA-12-030-01A
May 8, 2013 - This advisory is a follow up to a previous advisory titled "ICS-CERT - Siemens HMI Authentication Vulnerabilities" that was published ...

DHS, ICS-CERT Warn of Siemens HMI Vulnerabilities | Threatpost
threatpost.com/dhs-ics-cert-siemens-hmi-vulnerabilities
by Michael Milano - in 09 Google+ circles
Mar 21, 2013 - The Department of Homeland Security and the ICS-CERT issued an advisory yesterday warning of serious vulnerabilities in Siemens industrial ...

DHS Warns of HMI Vulnerabilities in Siemens WinCC TIA Portal

siemens s7-1200 vulnerability

Web Images Maps Shopping More Search tools

About 42,700 results (0.34 seconds)

Ad related to siemens s7-1200 vulnerability

New Siemens S7-1500 PLC - usa.siemens.com
www.usa.siemens.com/plc
Request a customer promo kit for our Programmable Logic Controller

Siemens S7-1200 Web Application Cross Site Scripting | ICS-CERT
ics-cert.us-cert.gov/advisories/ICSA-12-283-01
Oct 12, 2012 - This advisory provides mitigation details provided by Siemens regarding a vulnerability that impacts the Siemens S7-1200 Web Application Module.

Siemens S7-1200 Insecure Storage of HTTPS CA Certificate | ICS-CERT
ics-cert.us-cert.gov/advisories/ICSA-12-263-01
Sep 19, 2012 - Siemens has reported an insecure HTTPS certificate storage vulnerability in Siemens' S7-1200 v2.x. Siemens has provided guidance for mitigation.

PT-2012-50: Multiple vulnerabilities in Siemens SIMATIC S7-1200
en.securitylab.ru/lab/PT-2012-50

SIAS

ICS410 | ICS/SCADA Security Essentials 146

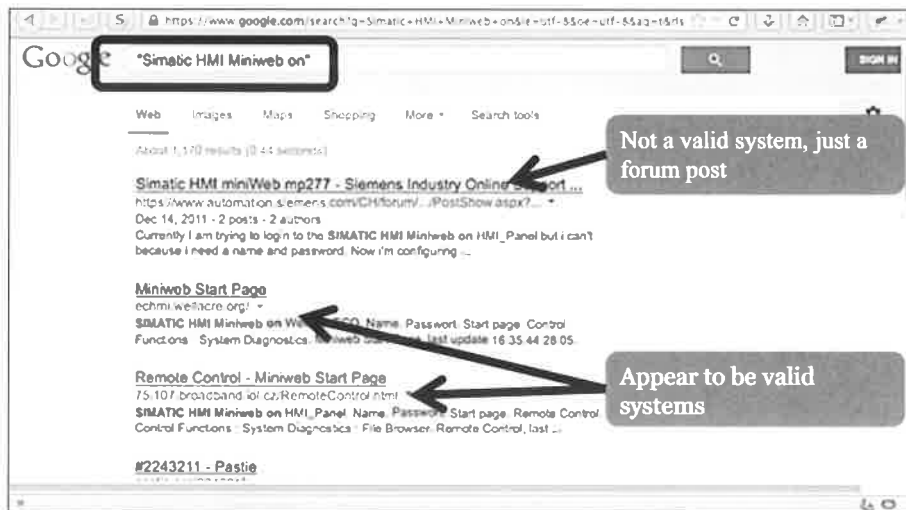
There are a few different paths you can take to identify Siemens' vulnerabilities. First, you could generically search for "Siemens vulnerability". This would provide search results like Siemens Vulnerability Handling page, which links to a page with all the security advisories. Or, you could Google for more specific terms like "Siemens hmi vulnerability" and "Siemens s7-1200 vulnerability", which will lead you to the ICS-CERT advisory pages for those specific products. Because Google's index is constantly changing, your exact search results may vary, but regardless, you should be able to find a lot of vulnerabilities in these two Siemens products.

A variety of vulnerabilities affect the HMI and the S7-1200 PLC. Here is a partial list of the most dangerous:

- **SQL Injection:** Allows attackers to read and write to the database and possibly gain a shell on the system
- **Cross-Site Scripting (XSS):** Allows attackers to run code in a user's browser
- **Authentication Bypass:** Allows attackers to log in without a username or password
- **Insecure Storage of SSL Certificates:** Allows attackers to steal the SSL certificate from the server and use it to decrypt traffic or clone the system
- **Denial of Service (DoS) on SNMP:** Allows attackers to stop the device from functioning

EXERCISE 4.4: FINDING REMOTE ACCESS

USING GOOGLE TO FIND HMI TARGETS



SANS

ICS410 | ICS/SCADA Security Essentials 147

If you know some exact keywords that appear on the login page of one of the systems, you can use Google to help identify systems running this software that are exposed to the internet. Now for our Siemens HMI, one of the strings that appear on the login page is:

Simatic HMI Miniweb on

If you use this search term with quotes surrounding it (as seen in the screenshot above), you should get a list of links to systems running that software. Note that not all the search results will be real systems running this software. For instance, the first search result above is a post by an individual about logging in to one of those systems where he mentions our search string in his post. Other systems may or may not be honeypots trying to lure in attackers to track their behavior. Regardless, if you visit any of these pages, do not attempt to log in or use a system, unless it is a system you have authorization to interact with.

If you suspect that you have Siemens HMI systems running on your company's network and are exposed to the internet, you can add **site:yourcompany.com** to the end of the search term to restrict the search to your company's domain. For instance, if you want to see if DOE is running any Siemens HMIs, you could do a search for:

"Simatic HMI Miniweb on" site:doe.gov

EXERCISE 4.4: FINDING REMOTE ACCESS

USING SHODAN TO FIND ONLINE HMIS

The screenshot shows the Shodan search results for the query "siemens hmi". The interface is divided into several sections:

- Home Search Directory Data Analytics/Exports**: Navigation menu.
- Services where the search term was found**: A callout box pointing to the "Services" table.
- Services**: A table listing discovered services:

Service	Count
SNMP	55
NetBIOS	1
SMB	1
- Top Countries**: A list of countries and their counts:

Country	Count
United States	11
Norway	7
Czech Republic	5
Sweden	5
France	5
- Top Cities**: A list of cities and their counts:

City	Count
7R 13& 3& 110	
- Host Profile: IP Redacted**: A detailed view of a host, including:
 - Summary**: IP: 213.115.12.43, Location: Sweden, Latitude/Longitude: 62.0, 15.0
 - HTTP**: HTTP/1.0 302 Object Moved, Content-Type: text/html, Content-Length: 0, Location: /Default.html
 - Other services running**: A callout box pointing to the "SNMP" service details.
 - SNMP**: Siemens SIMATIC HMI, XP277, 6ES7 643-0CB01-1AA0
 - HTTP**: HTTP/1.0 302 Object Moved, Location: /...

Caution: Use this tool with caution as your organization may prohibit Shodan access or restrict the use of company email when registering for Shodan. Also, depending on your region, there may be issues with connecting to resources discovered with Shodan. Please consult your IT security department with any questions prior to proceeding if you are unsure. Depending on your region, there may be issues with connecting to resources discovered with Shodan.

Now point your browser to www.shodan.io so we can see whether it can find any HMI targets. Shodan provides minimal functionality to guest users. If you would like to see the full details and service information, you will need a Shodan account. Shodan requires a username, an email address, and a password to set up an account. After supplying the basic information, you will receive a verification email. If you do not wish to create an account, some portions of this lab will not be available through Shodan.

If we do a simple search for "siemens hmi", Shodan returns a number of systems to us. If you look at the summary statistics on the left in the screenshot above, you can see that most of the results are SNMP services. Clicking on the details of the first result shows you a historic list of the services Shodan has found on this system, which the screenshot on the right shows. You can see that there are HTTP, SNMP, and NetBIOS services running on this system.

If you choose to visit any IPs you find in Shodan, be careful! Malicious attackers, security researchers, and government agencies often host honeypots that look like vulnerable machines to either attack or track the actions of visitors. Simply visiting a public IP address on the internet is legal in most jurisdictions around the world; performing any actions on that site could be considered illegal if you do not have permission to use that system.

EXERCISE 4.4: FINDING REMOTE ACCESS

USING SHODAN TO FIND ONLINE PLCS



Now let's use Shodan to search for Siemens S7 1200 PLCs. Doing such a search comes up with a handful of SNMP services claiming to be running that hardware. Sometimes it's hard to find the exact search term in Shodan to find the systems you are looking for; however, if you have access to any of those systems, you can do an Nmap scan of them with a `-sV` option and use keywords from the service banner that Nmap tells you. You can also usually use Google to help you find the appropriate strings as vendor documentation, blog posts, mailing lists, and other websites often contain this information.

If you look through the search PLC results, you can suspect that many of these systems are vulnerable to the SNMP DoS attack we found earlier. Once again, look at the details of the findings to see if any are running other services like HTTP.

EXERCISE 4.4: FINDING REMOTE ACCESS

FINDING VPN AND REMOTE ACCESS SERVICES

If you are willing to create a free account on Shodan, you will be able to access search filters

- city
- country
- geo
- hostname
- net
- os
- port
- before
- after

If you choose to register (or already have) an account, try the following search string

```
port:22,23,500,1194,1701,1723,3389,4500
```

Here are the remote access and VPN ports we were searching for above:

- 22 – Secure Shell (SSH)
- 23 – Telnet
- 500 – IPSec Internet Key Exchange (IKE)
- 1194 – OpenVPN
- 1701 – Layer 2 Tunneling Protocol (L2TP) VPN
- 1723 – Point-to-Point Tunneling Protocol (PPTP)
- 3389 – Microsoft Windows Remote Desktop Protocol (RDP)
- 4500 – IPSec NAT traversal – RFC 3947

EXERCISE 4.4: FINDING REMOTE ACCESS

LOOK FOR OTHER PRODUCTS AND VENDORS

Now pick some of the ICS products you have used in the past (or have heard of) and attempt to find publicly exposed systems using Google and Shodan

<https://www.shodan.io/explore/category/industrial-control-systems>

If you need some suggestions, try these manufacturers:

- ABB
- Emerson (DeltaV)
- Honeywell
- Invensys (with is now part of Schneider, but still appears in device outputs)
- Rockwell Automation
- Schneider Electric (Citect)

Now pick some of the ICS products you have used in the past (or have heard of) and attempt to find publicly exposed systems using Google and Shodan.

Or, you can try some of the links here:

<https://www.shodan.io/explore/category/industrial-control-systems>

If you need some suggestions, try these manufacturers:

- ABB
- Emerson (DeltaV)
- Honeywell
- Invensys
- Rockwell Automation
- Schneider Electric (Citect)

EXERCISE 4.4: FINDING REMOTE ACCESS

TAKEAWAYS AND RECOMMENDATIONS

Section takeaways

- Documentation and default settings are easy to discover and should always be changed
- Finding information on control systems connected to the internet is trivial

Recommendations to owner/operators

- Avoid connecting control systems directly to the internet
- Placing a web proxy in front of the web interface gains you very little security
- Using VPN connections is better if needed
- VPN followed by jump host with needed tools is best

Recommendations to vendors

- Ensure you leverage secure remote access models for customer support

This page intentionally left blank.

COURSE RESOURCES AND CONTACT INFORMATION



AUTHOR CONTACT

Justin Searle – justin@controlthings.io



SANS INSTITUTE

11200 Rockville Pike, Suite 200
N. Bethesda, MD 20852
301.654.SANS(7267)



ICS RESOURCES

ics.sans.org
Twitter: [@sansics](https://twitter.com/sansics)
SANS ICS Community
<https://ics-community.sans.org/signup>



SANS EMAIL

GENERAL INQUIRIES: info@sans.org
REGISTRATION: registration@sans.org
TUITION: tuition@sans.org
PRESS/PR: press@sans.org

This page intentionally left blank.

