AUD507 | AUDITING & MONITORING NETWORKS, PERIMETERS, AND SYSTEMS

507.1

Enterprise Audit Fundamentals; Discovery and Scanning Tools

SANS

THE MOST TRUSTED SOURCE FOR INFORMATION SECURITY TRAINING, CERTIFICATION, AND RESEARCH | sans.org

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 SANS INSTITUTE FOR THE COURSEWARE. YOU AGREE THAT THIS AGREEMENT IS ENFORCEABLE LIKE ANY WRITTEN NEGOTIATED AGREEMENT SIGNED BY YOU.

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

If User does not agree, User may return the Courseware to 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 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 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 trademarks 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.

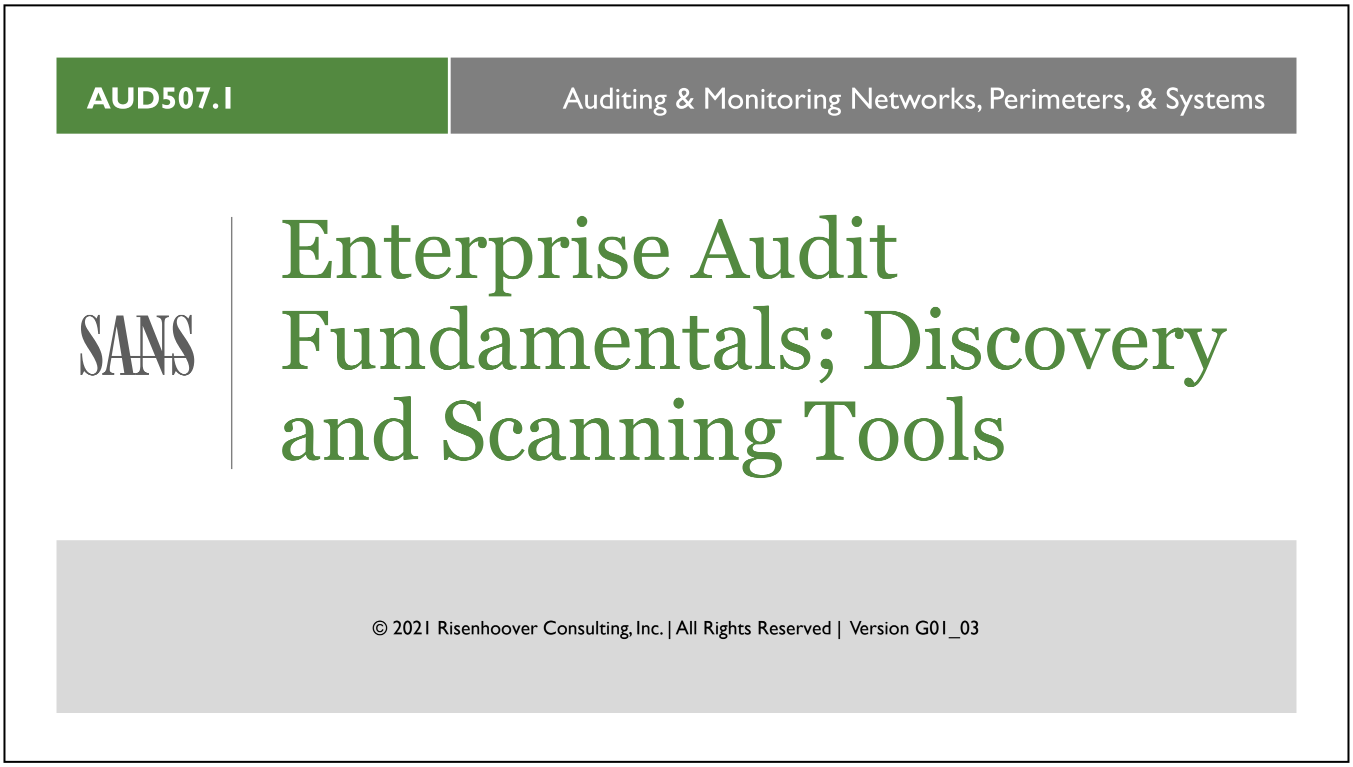

AUD507.1

Auditing & Monitoring Networks, Perimeters, & Systems

SANS

Enterprise Audit Fundamentals; Discovery and Scanning Tools

Welcome to the SANS AUD507 course! This course is written, maintained, and frequently taught by Clay Risenhoover. I am always looking for ways to improve this courseware. If you have questions or suggestions for how to improve the course, or if you need any additional materials referenced during the class, please let me know. If you find errors or inaccuracies in the course books, I encourage you to pass those on to me. You can email me at clay@risenhooverconsulting.com. Please put either “SANS” or “AUD507” in the subject line to ensure I see the email.

The entire content of this and every other volume in this course is © 2021 Risenhoover Consulting, Inc.

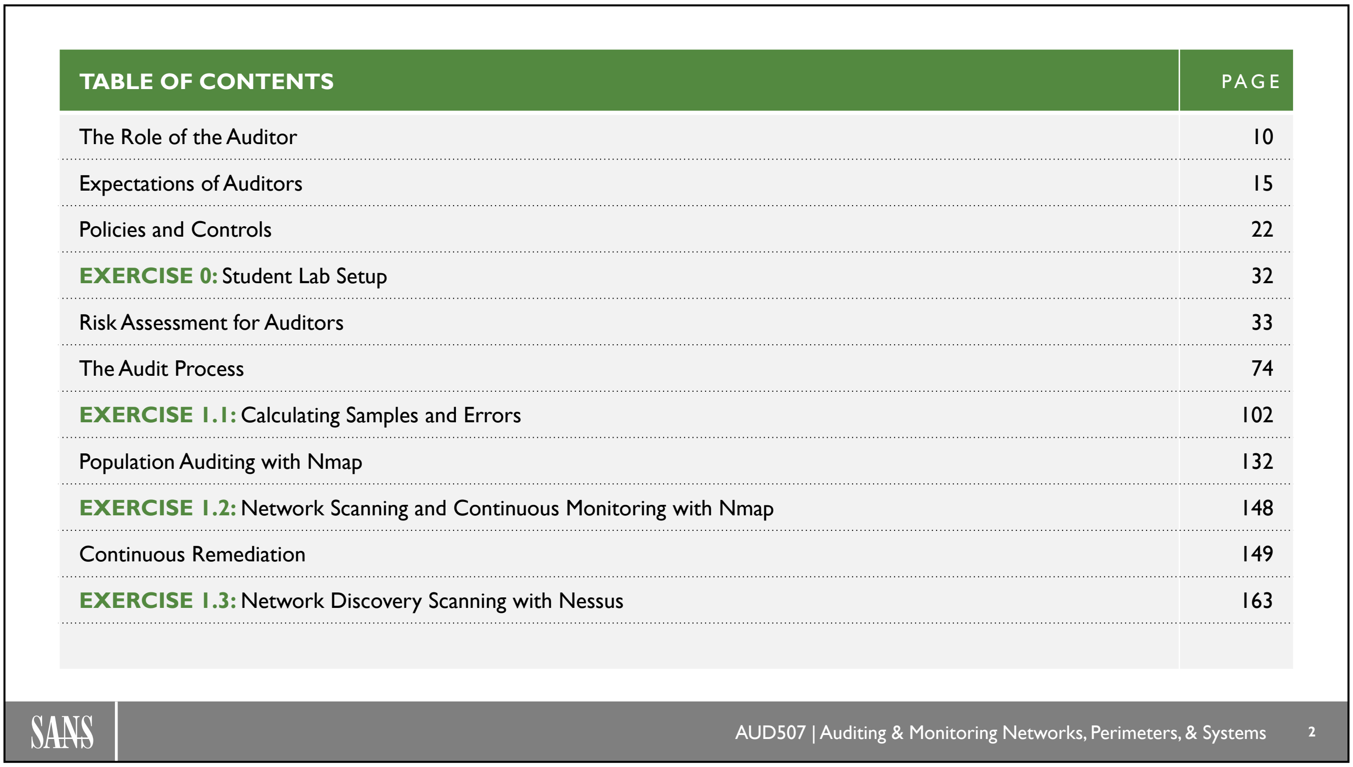

TABLE OF CONTENTS

TABLE OF CONTENTS	PAGE
The Role of the Auditor	10
Expectations of Auditors	15
Policies and Controls	22
EXERCISE 0: Student Lab Setup	32
Risk Assessment for Auditors	33
The Audit Process	74
EXERCISE 1.1: Calculating Samples and Errors	102
Population Auditing with Nmap	132
EXERCISE 1.2: Network Scanning and Continuous Monitoring with Nmap	148
Continuous Remediation	149
EXERCISE 1.3: Network Discovery Scanning with Nessus	163

This page intentionally left blank.

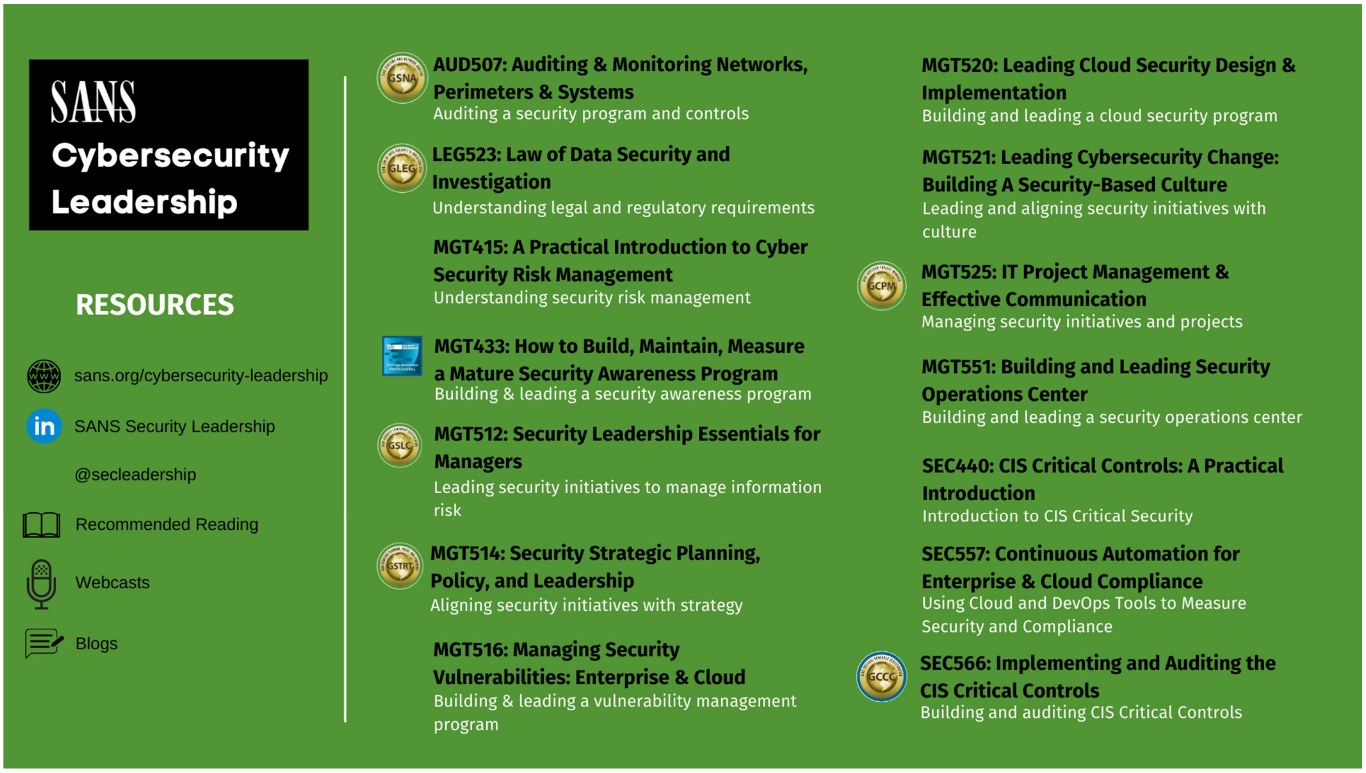

SANS
Cybersecurity Leadership

RESOURCES

- sans.org/cybersecurity-leadership
- SANS Security Leadership
- @secleadership
- Recommended Reading
- Webcasts
- Blogs

AUD507: Auditing & Monitoring Networks, Perimeters & Systems
Auditing a security program and controls

LEG523: Law of Data Security and Investigation
Understanding legal and regulatory requirements

MGT415: A Practical Introduction to Cyber Security Risk Management
Understanding security risk management

MGT433: How to Build, Maintain, Measure a Mature Security Awareness Program
Building & leading a security awareness program

MGT512: Security Leadership Essentials for Managers
Leading security initiatives to manage information risk

MGT514: Security Strategic Planning, Policy, and Leadership
Aligning security initiatives with strategy

MGT516: Managing Security Vulnerabilities: Enterprise & Cloud
Building & leading a vulnerability management program

MGT520: Leading Cloud Security Design & Implementation
Building and leading a cloud security program

MGT521: Leading Cybersecurity Change: Building A Security-Based Culture
Leading and aligning security initiatives with culture

MGT525: IT Project Management & Effective Communication
Managing security initiatives and projects

MGT551: Building and Leading Security Operations Center
Building and leading a security operations center

SEC440: CIS Critical Controls: A Practical Introduction
Introduction to CIS Critical Security

SEC557: Continuous Automation for Enterprise & Cloud Compliance
Using Cloud and DevOps Tools to Measure Security and Compliance

SEC566: Implementing and Auditing the CIS Critical Controls
Building and auditing CIS Critical Controls

The SANS Institute was established in 1989 as a cooperative research and education organization. Its programs now reach more than 165,000 security professionals around the world.

SANS is the most trusted and by far the largest source for information security training and security certification in the world. It also develops, maintains, and makes available at no cost, the largest collection of research documents about various aspects of information security, and it operates the Internet's early warning system - the Internet Storm Center.

SANS offers a number of courses that prepare cyber leaders and managers for building and leading world-class security teams. As security becomes more relevant to the business, we need to develop business, leadership, and technical skills to effectively interact with business leaders as well as lead and inspire our technical teams. The following courses give you the necessary skills to navigate in the new world of security:

AUD507: Auditing & Monitoring Networks, Perimeters, and Systems | GSNA | 6 Days

Learn how to audit and monitor your key controls.

LEG523: Law of Data Security and Investigations | GLEG | 5 Days

Understand key lessons on the law of data security and investigations that all managers, leaders, & executives should know.

MGT415: A Practical Introduction to Cyber Security Risk Management | SSAP | 2 Days

Understand how to assess and manage cyber security risk.

MGT433: Managing Human Risk: Mature Security Awareness Programs | 2 Days

Learn to build a security awareness program the right way.

MGT512: Security Leadership Essentials for Managers | GSLC | 5 Days

Get up to speed on information security issues and terminology. You don't just learn security; you learn how to manage security.

MGT514: IT Security Strategic Planning, Policy, & Leadership | GSTRT | 5 Days

Learn to build and execute strategic plans, develop and assess policy, and utilize management tools to lead, inspire, and motivate your teams.

MGT516: Managing Security Vulnerabilities: Enterprise and Cloud | 5 Days

Learn to build and manage a vulnerability management program for your enterprise and cloud systems.

MGT520: Leading Cloud Security Design & Implementation | 3 Days

Learn how to build and lead a cloud security program

MGT521: Leading Cybersecurity Change: Building A Security-Based Culture | 5 Days

Apply the concepts of change management to embed a strong security culture in specific security initiatives or organization wide.

MGT525: IT Project Management and Effective Communication | GCPM | 6 Days

Learn how to effectively drive and manage projects and key initiatives.

MGT551: Building and Leading Security Operations Centers | 5 Days

Learn how to build, operate, and continuously improve your Security Operations Center (SOC).

SEC440: CIS Critical Controls: A Practical Introduction | 2 Days

Introduction to proven techniques and tools needed to implement and audit CIS Critical Controls v8 as documented by the Center for Internet Security (CIS).

SEC557: Continuous Automation for Enterprise and Cloud Compliance | 3 Days

Teaching professionals tasked with ensuring security and compliance how to stop being a roadblock and work at the speed of the modern enterprise.

SEC566: Implementing and Auditing CIS Critical Controls | GCCC | 5 Days

Learn about the CIS Critical Controls v8 that can serve as the foundation for your security program, as documented by the Center for Internet Security (CIS).

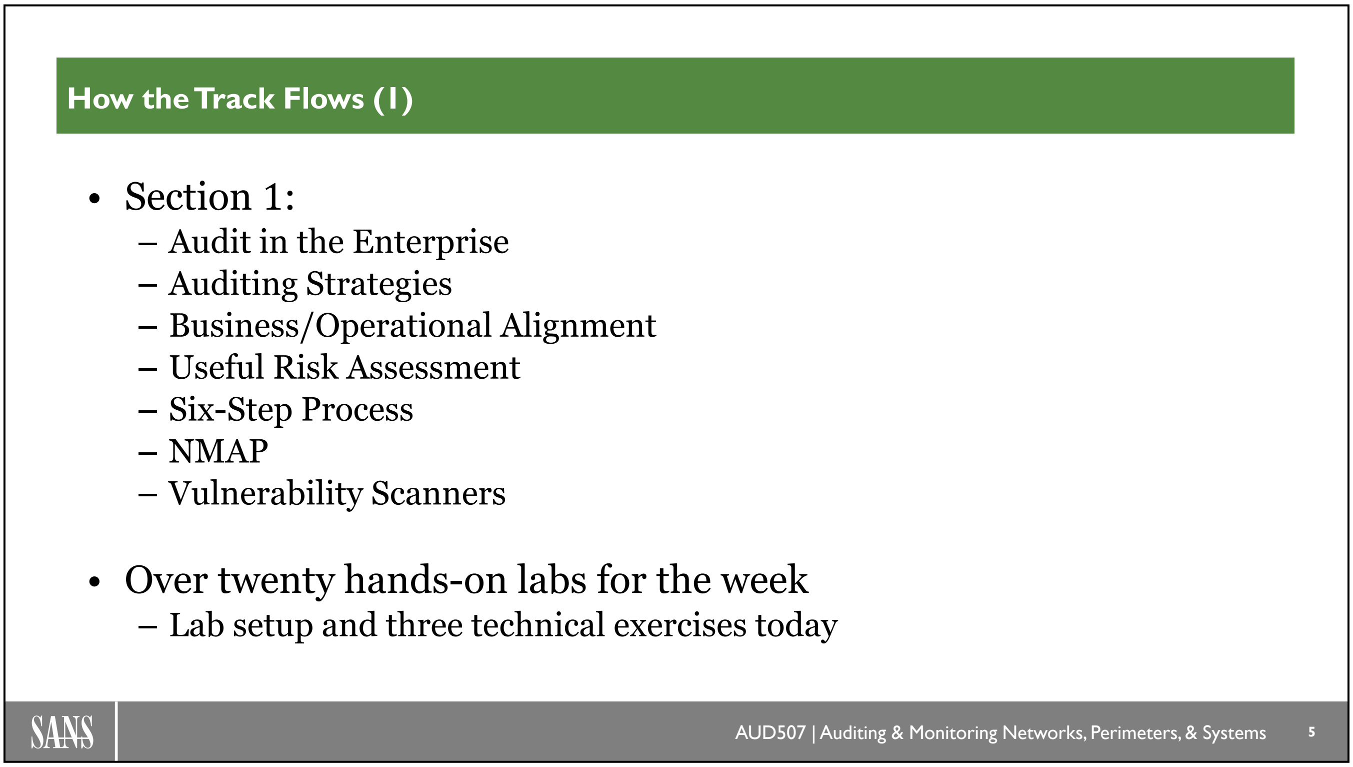

How the Track Flows (1)

- Section 1:
  - – Audit in the Enterprise
  - – Auditing Strategies
  - – Business/Operational Alignment
  - – Useful Risk Assessment
  - – Six-Step Process
  - – NMAP
  - – Vulnerability Scanners
- Over twenty hands-on labs for the week
  - – Lab setup and three technical exercises today

SANS | AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 5

Let's take a moment to see how this course will develop. Today, we define a few terms and try to develop a real and practical auditing mentality. We spend some time discussing recognized auditing standards from a 50,000-foot point of view and then spend some time with the specifics later in the week. We lay out a few real-world auditing strategies and discuss how to apply them to practical things we can all relate to. We also spend time working with two useful risk-assessment methodologies for auditors and anyone else involved in evaluating policies or specifying controls. In connection with this, one of the most important things we spend time with this week is how to ensure that what we report is aligned with the operational needs of the organization. If we fail to do this, then in many ways our work is irrelevant to the business. Without relevancy, funding becomes extremely difficult. We also go through the six-step auditing process. One of our goals is to describe a full six-step audit here in class today!

In addition to these less technical aspects, we spend a significant portion of the afternoon covering the technical tools that will make you "look smart" as an auditor. Nmap and your vulnerability scanner can be used to solve a wide range of audit problems; we will spend time with both of those tools today so you can leverage them throughout the week.

For the first part of the day, we want to make it clear that the material is *not* highly technical. In fact, we try hard to not be too technical! We focus on *what* you want to audit without being distracted by the *how*.

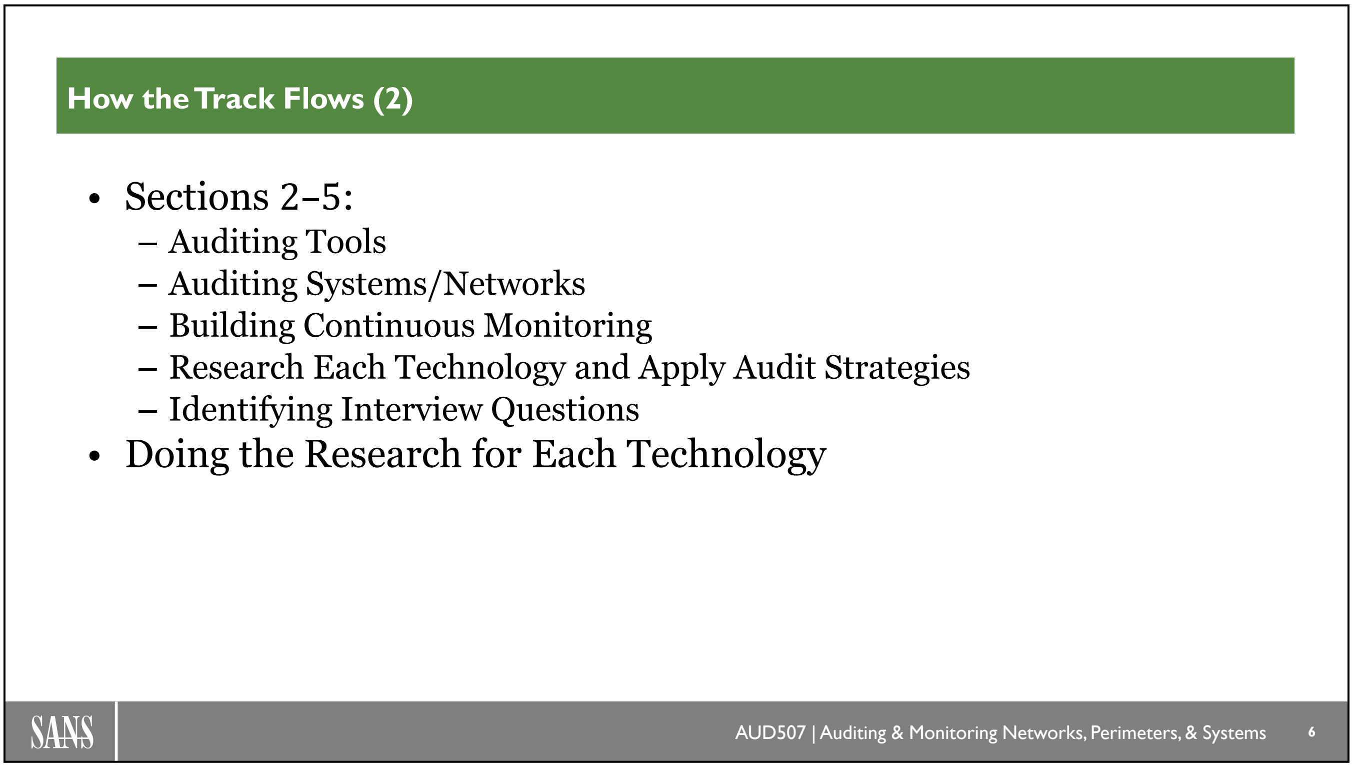

How the Track Flows (2)

- Sections 2–5:
  - Auditing Tools
  - Auditing Systems/Networks
  - Building Continuous Monitoring
  - Research Each Technology and Apply Audit Strategies
  - Identifying Interview Questions
- Doing the Research for Each Technology

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 6

The rest of the track is focused on applying the concepts from Section 1 to specific real-world situations we all run into. In other words, the *how*. For instance, we spend one section applying our strategies to Windows domains. We spend another section applying our strategies to UNIX systems. Another section deals with using Nmap to audit networks. We also focus on the application of audit strategies to web applications. This is what we call "real-world auditing." As we go through these sections, you'll see that there is a lot of material; however, remember that the goal is not to teach you about every possible tool and method for auditing a Windows domain or a specific Linux distribution, for instance, but to walk through how to apply the strategies to a specific system. Hopefully in the process, you learn as a side effect about some interesting tools and features in each system that we audit.

I'm a subscriber to the "Don't Repeat Yourself" school of development. We carry this same philosophy to Information Security and compliance auditing. What this means is that if there is a task or test that we need to do more than one time, we always take the time to automate it. We're going to take that one step further. If we're going to take the time to automate an audit activity, then we will always either take the time to create an automated continuous monitoring process or at least describe what that process would look like. Continuous monitoring and reporting are the real values for the organization. The main thing we do in this course is walk you through the type and depth of research necessary to create an effective audit program for a particular technology. It is our hope that you leave with a set of approaches you can use for any system you encounter in the future!

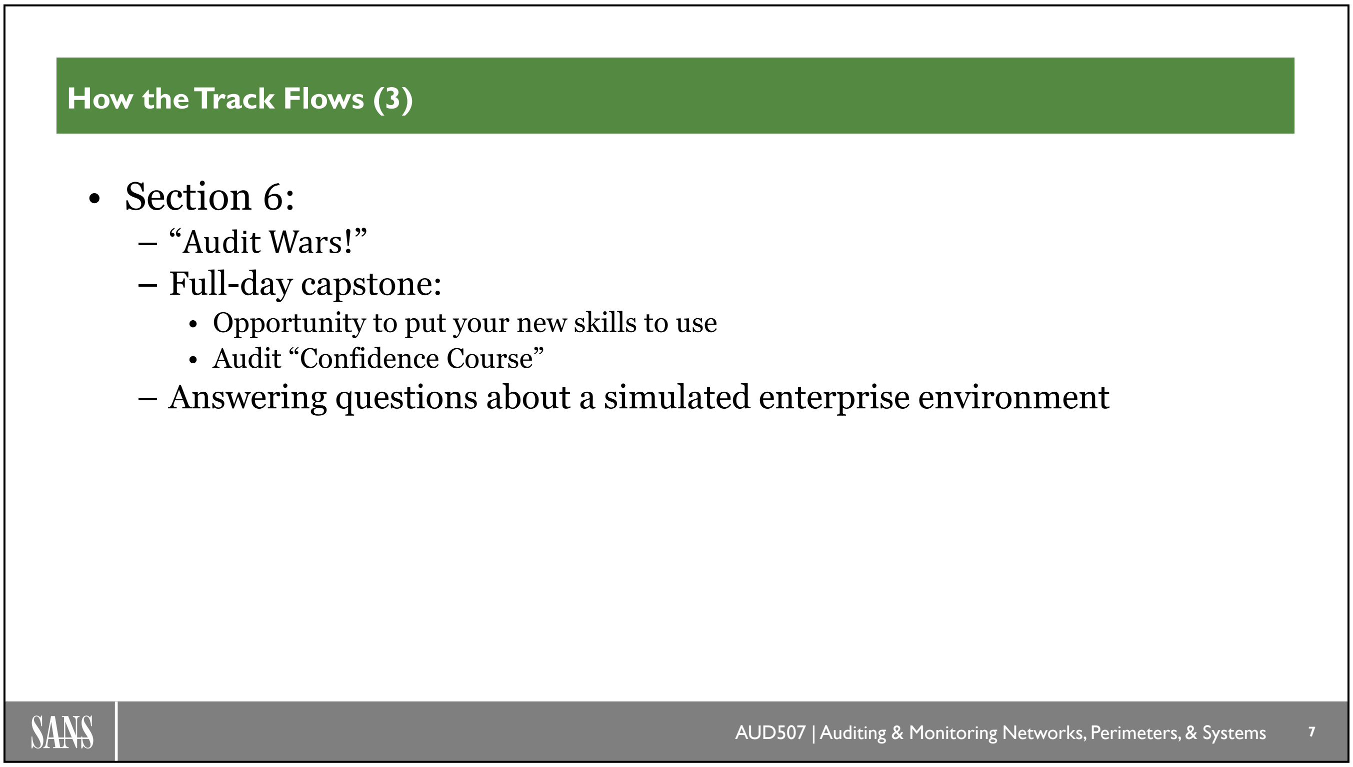

How the Track Flows (3)

- Section 6:
  - "Audit Wars!"
  - Full-day capstone:
    - Opportunity to put your new skills to use
    - Audit "Confidence Course"
  - Answering questions about a simulated enterprise environment

SANS | AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 7

Many students tell us that the capstone exercise is the most exciting part of this whole course. After we work through the various strategies and techniques, and after you learn plenty of audit ideas for technical systems, we unleash you on a virtualized enterprise environment. This capstone "audit the flag" challenge is an interactive environment in which you have the opportunity to work your way through a variety of audit exercises, applying the techniques you learned during the week. You can focus on just one or two aspects to drill in deeply, or you can approach it more broadly and try to cover as many of the systems as you can.

At the end of the experience, we talk about your findings and reflect on the possible root-cause issues. Together, we will answer questions that your findings might raise.

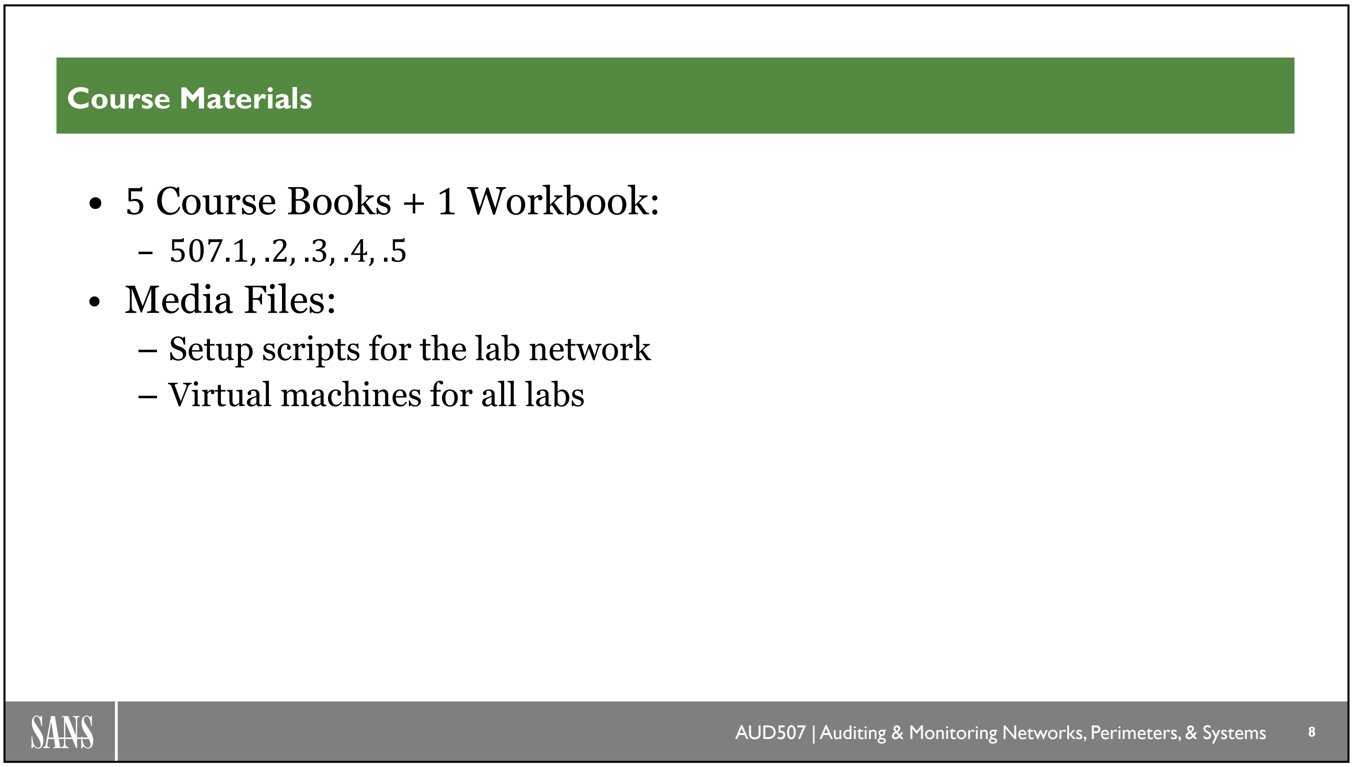

Course Materials

- 5 Course Books + 1 Workbook:
  - 507.1, .2, .3, .4, .5
- Media Files:
  - Setup scripts for the lab network
  - Virtual machines for all labs

Before the class, you should have received a course book for every section of the course *except* for the "Audit the Flag" capstone exercise. Audit the Flag is a full-day, hands-on lab, so there is no associated course book. You should also have received a book titled *Workbook*. It has all the instructions and walkthroughs for all the labs in the course, except for Audit Wars, of course.

You should have also received media files either on USB or by download. These files contain all the tools, virtual machines, and other materials you need to complete the hands-on exercises.

Course Web Sites

- Public Wiki - Open to public
  - https://www.aud507.com
  - References and cheat sheets
- Student Lab Wiki - Students only, please
  - https://lab-g.aud507.com
  - Username: **student-g**
  - Password: **pecantrackednoises**
  - Full content of lab workbook
  - Copy/paste into your lab VMs
  - Lab videos!
- Both are searchable

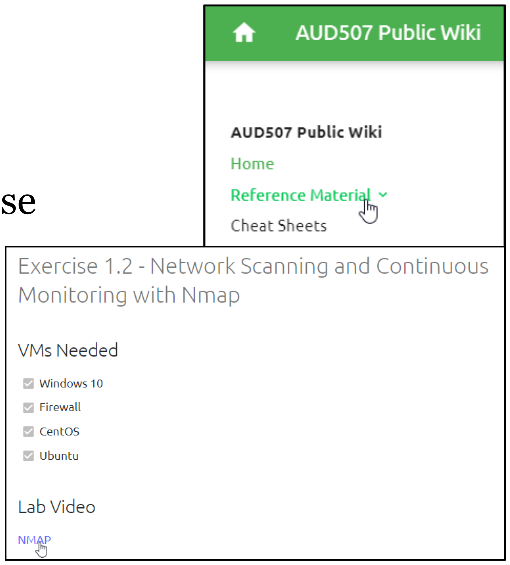

AUD507 Public Wiki

AUD507 Public Wiki

Home

Reference Material ⌄

Cheat Sheets

Exercise 1.2 - Network Scanning and Continuous Monitoring with Nmap

VMs Needed

☑ Windows 10

☑ Firewall

☑ CentOS

☑ Ubuntu

Lab Video

NMAP

Two of the more exciting resources for the class are the companion websites.

The public Wiki website is open for anyone (even non-students!) to view. Feel free to share this link with your friends and colleagues who might be interested in the content. The public Wiki contains over 100 links to web resources of interest to audit and security professionals. These references are where to go for more information about the topics we discuss together in class. The Wiki also contains digital copies of all of the cheat-sheets used in class.

The 507 lab Wiki is for use only by AUD507 students. This website contains a completely digital version of the student lab workbook and allows the student to copy commands for pasting directly into the VMs. No more typos making it difficult to complete a step! The lab Wiki also has links to videos of the course author working through and explaining every lab exercise. These can be very helpful when you get stuck working a lab or if you want to revisit a concept after class is over. The lab Wiki is custom-designed to match your version of the course and will be available for at least one year after you take your live or online course.

Both websites include the ability to search by keyword, making it faster and easier for you to find what you need.

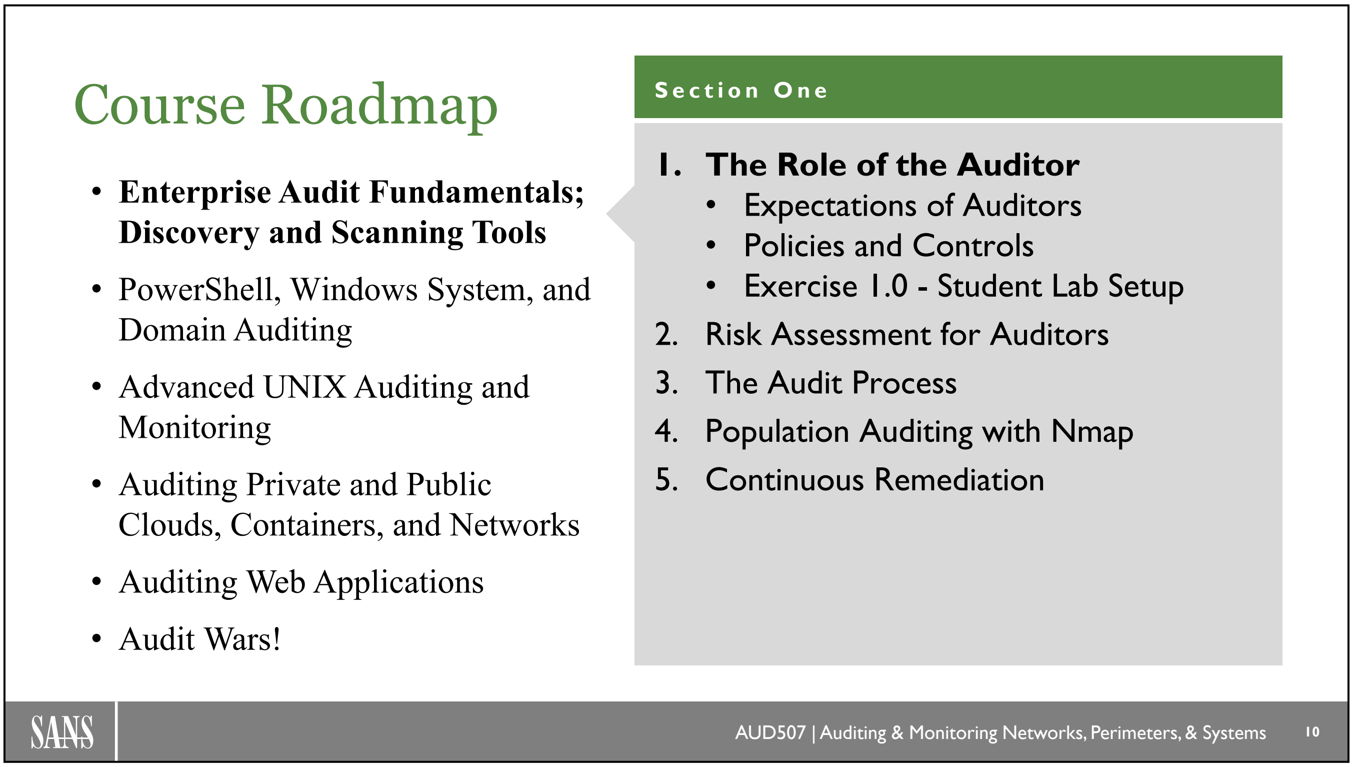

Course Roadmap

- **Enterprise Audit Fundamentals; Discovery and Scanning Tools**
- PowerShell, Windows System, and Domain Auditing
- Advanced UNIX Auditing and Monitoring
- Auditing Private and Public Clouds, Containers, and Networks
- Auditing Web Applications
- Audit Wars!

Section One

1. **The Role of the Auditor**
   - Expectations of Auditors
   - Policies and Controls
   - Exercise 1.0 - Student Lab Setup
2. Risk Assessment for Auditors
3. The Audit Process
4. Population Auditing with Nmap
5. Continuous Remediation

This page intentionally left blank.

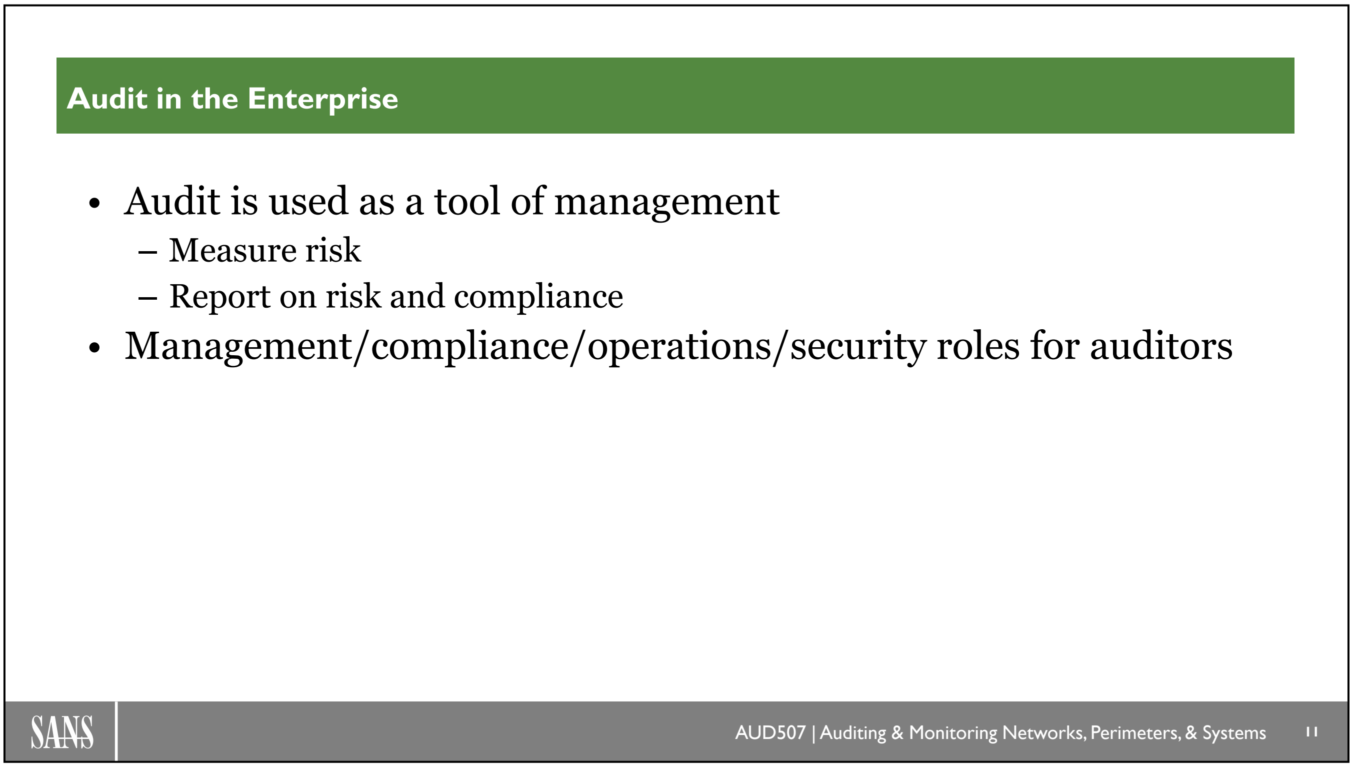

Audit in the Enterprise

- Audit is used as a tool of management
  - Measure risk
  - Report on risk and compliance
- Management/compliance/operations/security roles for auditors

The audit function performs many important services for the enterprise. Primarily, auditors help management to measure the risk which the organization faces and to understand how to best mitigate this risk. Auditors can help their organizations to better comply with laws and regulations and to move more effectively and efficiently toward meeting organizational objectives. Auditors often act in support of other compliance and operations roles within their organizations.

The techniques discussed in this class are useful for more than auditors, though. The skills you learn in AUD507 will be helpful in fulfilling roles related to organizational management, legal and regulatory compliance, IT management, IT operations, and information security.

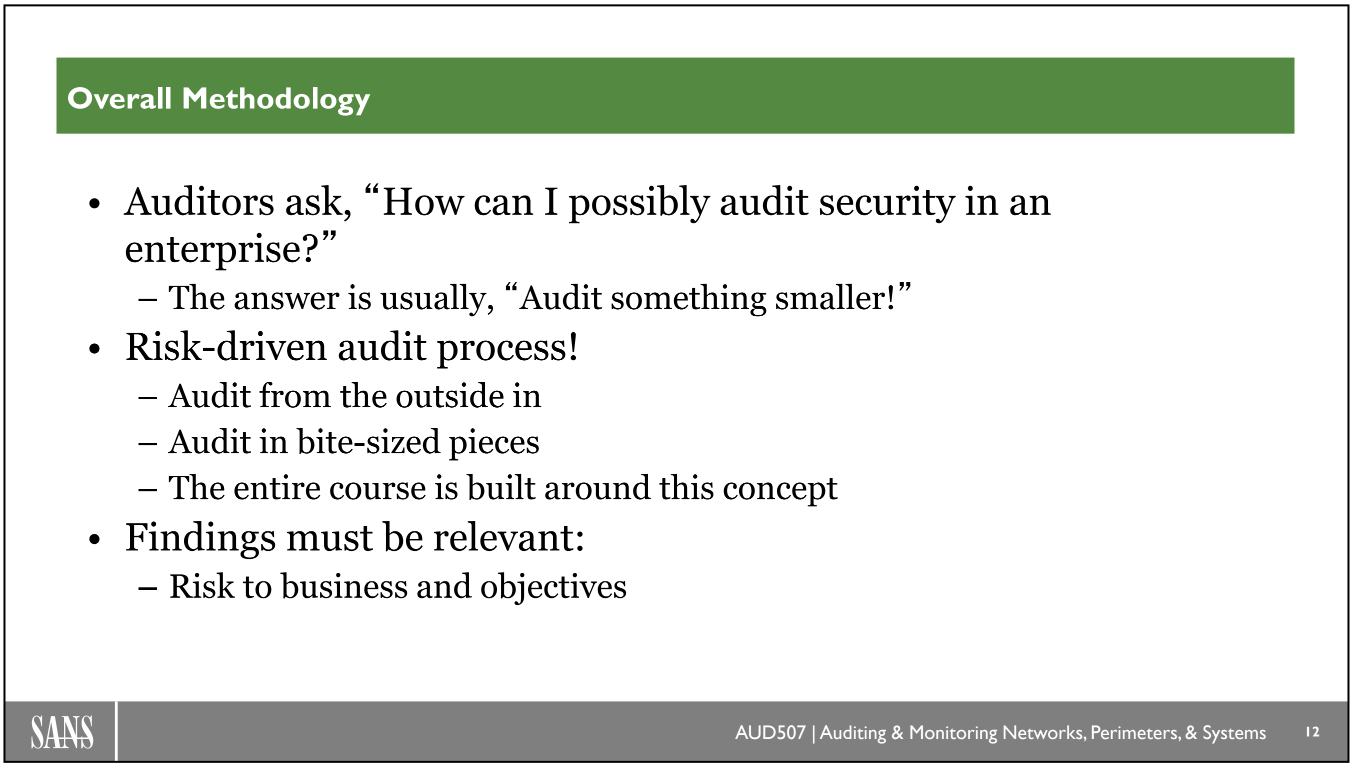

Overall Methodology

- Auditors ask, “How can I possibly audit security in an enterprise?”
  - The answer is usually, “Audit something smaller!”
- Risk-driven audit process!
  - Audit from the outside in
  - Audit in bite-sized pieces
  - The entire course is built around this concept
- Findings must be relevant:
  - Risk to business and objectives

SANS | AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 12

Quite frequently we have auditors in the class who ask, “How can I possibly audit <insert some really big scope here>?” The trouble is that for an auditor to do a good job, he needs to spend what can be a significant amount of time with each item that is in the scope of, or covered by, the audit. As a result, when an auditor says that he is going to audit “the security of the infrastructure,” he ends up with a huge number of systems and functions to analyze. How can this be turned into a manageable process?

The answer is to audit smaller pieces of the whole. To determine where to start, consider where the biggest risks are; for instance, although it would be important to make sure that internal servers are well secured, it would likely be riskier to leave external servers poorly secured. The simple result is that we can say that although the internal servers are more valuable, there would be a more immediate impact to the organization should it suffer a compromise of one of its external systems, so we audit from the outside in. This entire course works in this direction! We begin by covering audit strategies and legal concerns today, and then when we turn technical, we start from the outside and work our way into the infrastructure.

A CISO sitting through this class had a comment after the first lecture. He said, “I heard you say something that I’ve never heard from any other auditor, and if I had, I would have paid more attention to what they do.” What was it? As we cover the material in the books and have discussions, reflect on our focus on *effective* controls that *inform* the business process in connection with operational and business *objectives*. If we do not take these things into consideration during an audit, then our findings may be “interesting,” but they will likely also be irrelevant.

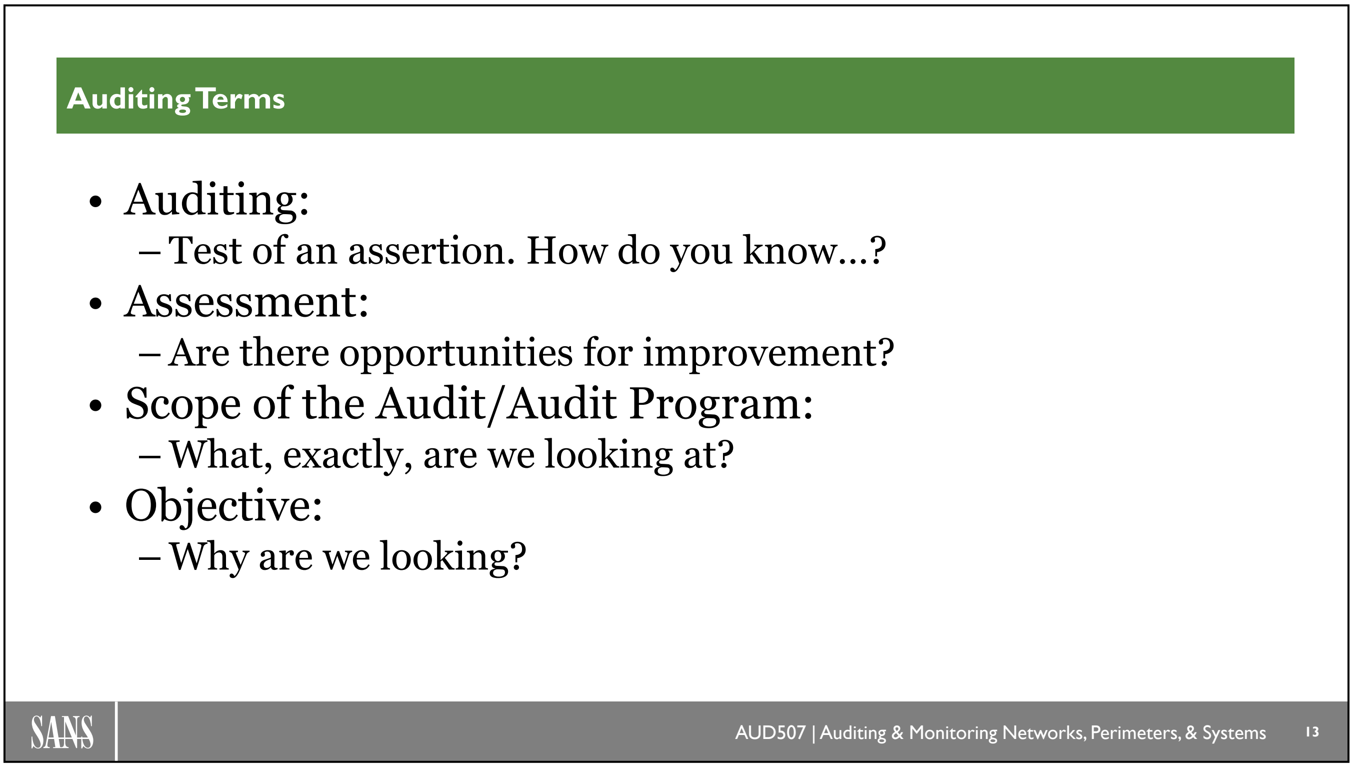

Auditing Terms

- Auditing:
  - Test of an assertion. How do you know…?
- Assessment:
  - Are there opportunities for improvement?
- Scope of the Audit/Audit Program:
  - What, exactly, are we looking at?
- Objective:
  - Why are we looking?

Let's start by defining auditing. *Auditing*, in general, is best described as the function of measuring something against a standard. Although system auditors tend to focus on using auditing to measure the security over time of a system, auditing can be applied to anything.

The three most effective places to apply auditing in Information Technology and Information Assurance are at the *policy* level, *procedure* level, and *system* level. You might also think of the system level as the *application* level. This doesn't mean software; rather, "application" here indicates the point at which we apply the policies or procedures.

Audits can come in all shapes and sizes as well. Examples of audits that you may experience (or conduct) include security audits, financial audits, compliance audits, and more. As Information Security (IS) professionals, we will most commonly be involved in conformance or compliance audits. A *conformance audit* is interested in measuring how well a system or process conforms to the policies and/or procedures that have been defined in the organization. A *security audit* is a more general audit that can be used to measure policy, procedure, or audits against industry best practice to determine if there is a need for improvement at a level higher than the application or system level.

A simple definition for auditing is that auditing answers the question, "How do you know you…?"

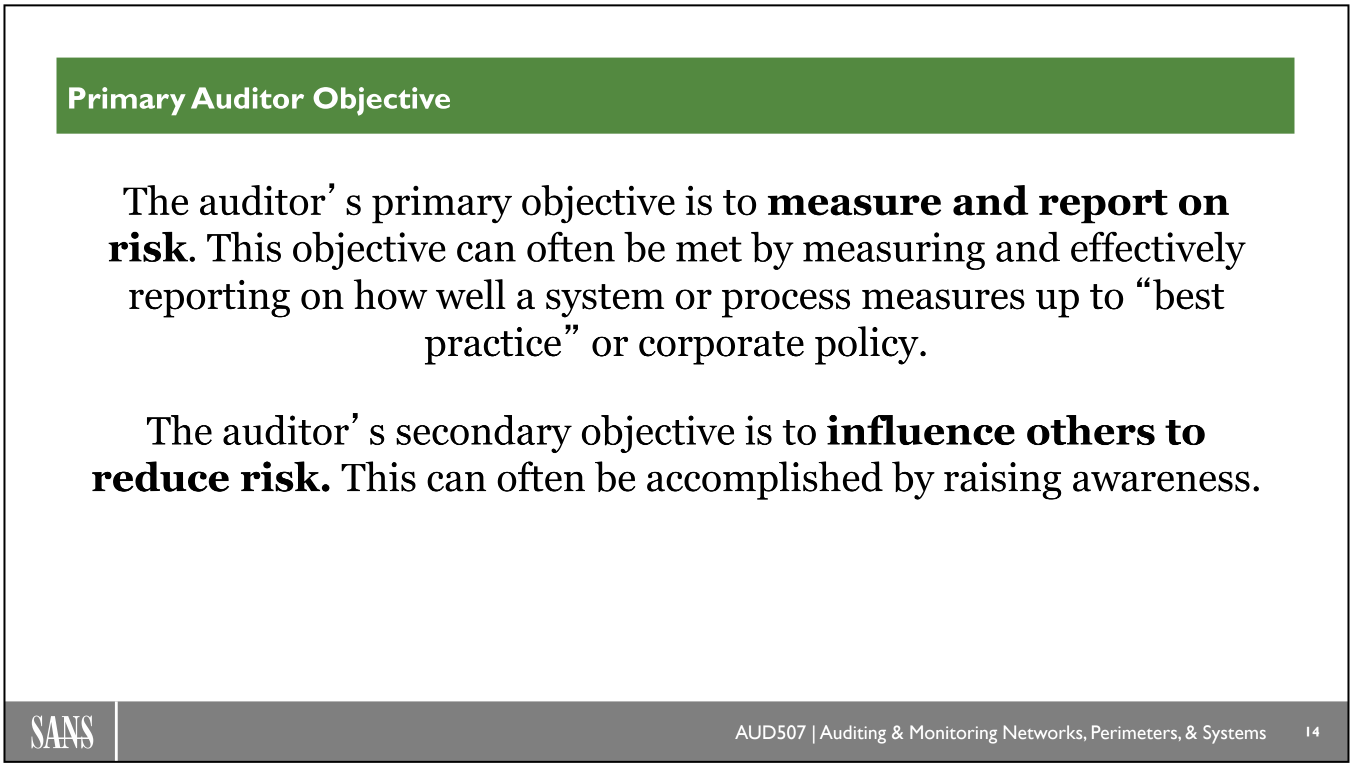

Primary Auditor Objective

The auditor' s primary objective is to **measure and report on risk**. This objective can often be met by measuring and effectively reporting on how well a system or process measures up to "best practice" or corporate policy.

The auditor' s secondary objective is to **influence others to reduce risk.** This can often be accomplished by raising awareness.

SANS | AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 14

Auditors in an organization are in a unique position. Traditionally, they are among the few who are invited to share thoughts and opinions as to how well an organization is doing at meeting its objectives. Auditors are also invested with a great deal of trust by upper management, who rely on the reports of auditors when making business-guiding decisions. Someone acting as an auditor must possess the highest level of integrity. An individual who performs as an auditor must also have a deep desire to gain the level of understanding and technical competence in a wide range of disciplines to perform functions effectively. Finally, this individual must be viewed as truly neutral.

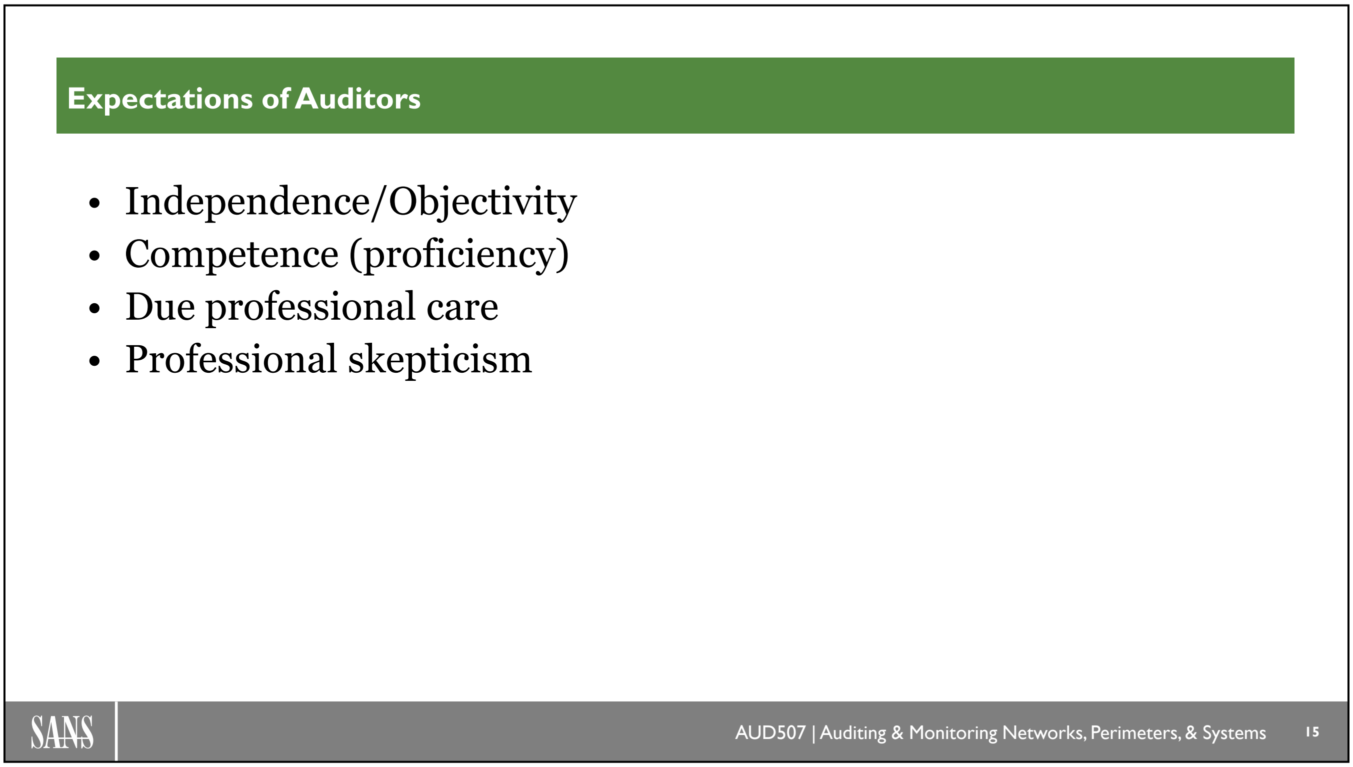

Expectations of Auditors

- Independence/Objectivity
- Competence (proficiency)
- Due professional care
- Professional skepticism

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 15

The successful IT auditor will exhibit certain characteristics. We cover some of the most important on the next few slides.

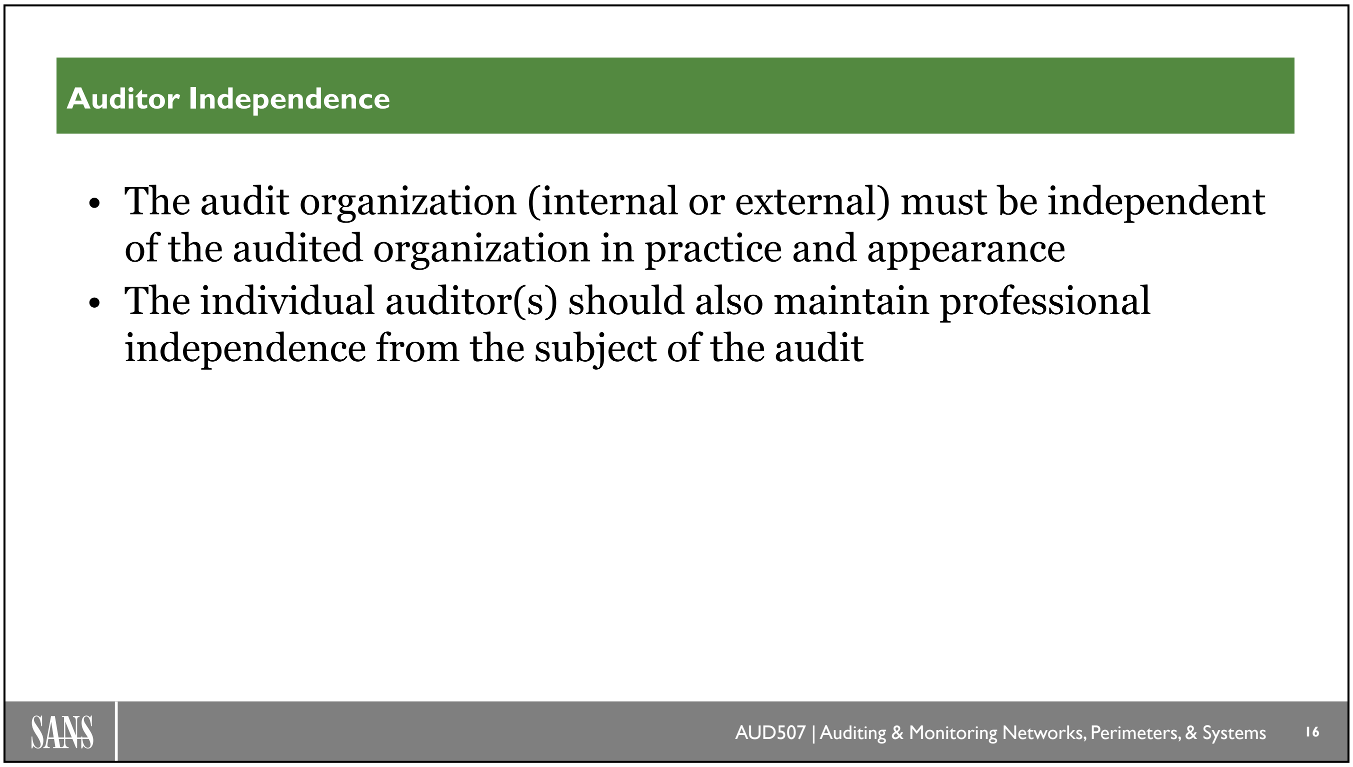

Auditor Independence

- The audit organization (internal or external) must be independent of the audited organization in practice and appearance
- The individual auditor(s) should also maintain professional independence from the subject of the audit

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 16

For an auditor to provide any value to an organization, they must provide an objective viewpoint on the systems and processes under review. This means that there must be no conflicts of interest which would impair the auditor's independence at either the organizational or individual level.

One of the results of the Enron / Arthur Anderson scandal was that the Sarbanes Oxley Act (SOX) required that accounting firms who provide consulting services to a client may not also provide audit services to the same client. This is a formalization of the concept of auditor independence, and it just makes sense. A consultant who designs a system cannot be expected to provide an objective and independent review of the effectiveness of the system. They already believe the system is effective because they think they designed it correctly in the first place.

This concept should extend to the relationship among individual auditors and the employees of the organization under review. While it is fine for the auditor to be friendly with the employees being audited, there are lines which should not be crossed. Auditors should avoid situations and actions which would cause their independence to be questioned.

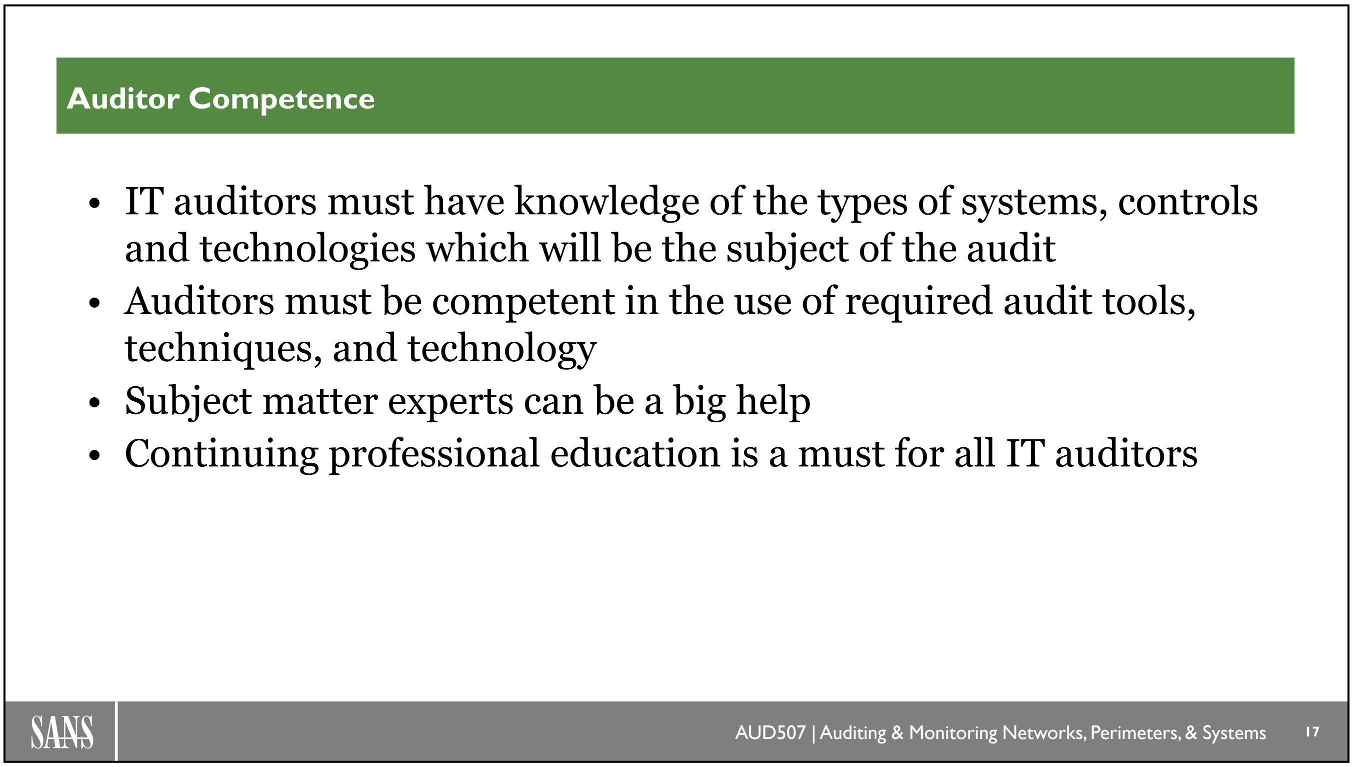

Auditor Competence

- IT auditors must have knowledge of the types of systems, controls and technologies which will be the subject of the audit
- Auditors must be competent in the use of required audit tools, techniques, and technology
- Subject matter experts can be a big help
- Continuing professional education is a must for all IT auditors

Auditors should have the skills and abilities necessary to complete the job at hand. This may include knowledge of certain technologies or business processes, or the ability to use the technical tools required to test controls. This also extends to the ability to handle other parts of the audit process, like performing planning activities, interviewing staff, running a meeting, or writing a report.

IT audit is a big field, and no one will be an expert in all areas. In the case that the auditor is tasked with something they are not technically competent to handle, they can do one of two things:

- Develop the competency by attending training (I hear SANS has some good courses) or doing self-study.
- Enlist the help of a qualified subject matter expert. This could even be someone from the organization under review, as long as independence can be maintained. For instance, the auditor may request the help of a qualified Unix administrator from another department to perform technical tests which the auditor is not qualified to perform.

Continuing professional education is a must for IT auditors. In fact, for many IT and internal audit certifications, it is a requirement of continued certification.

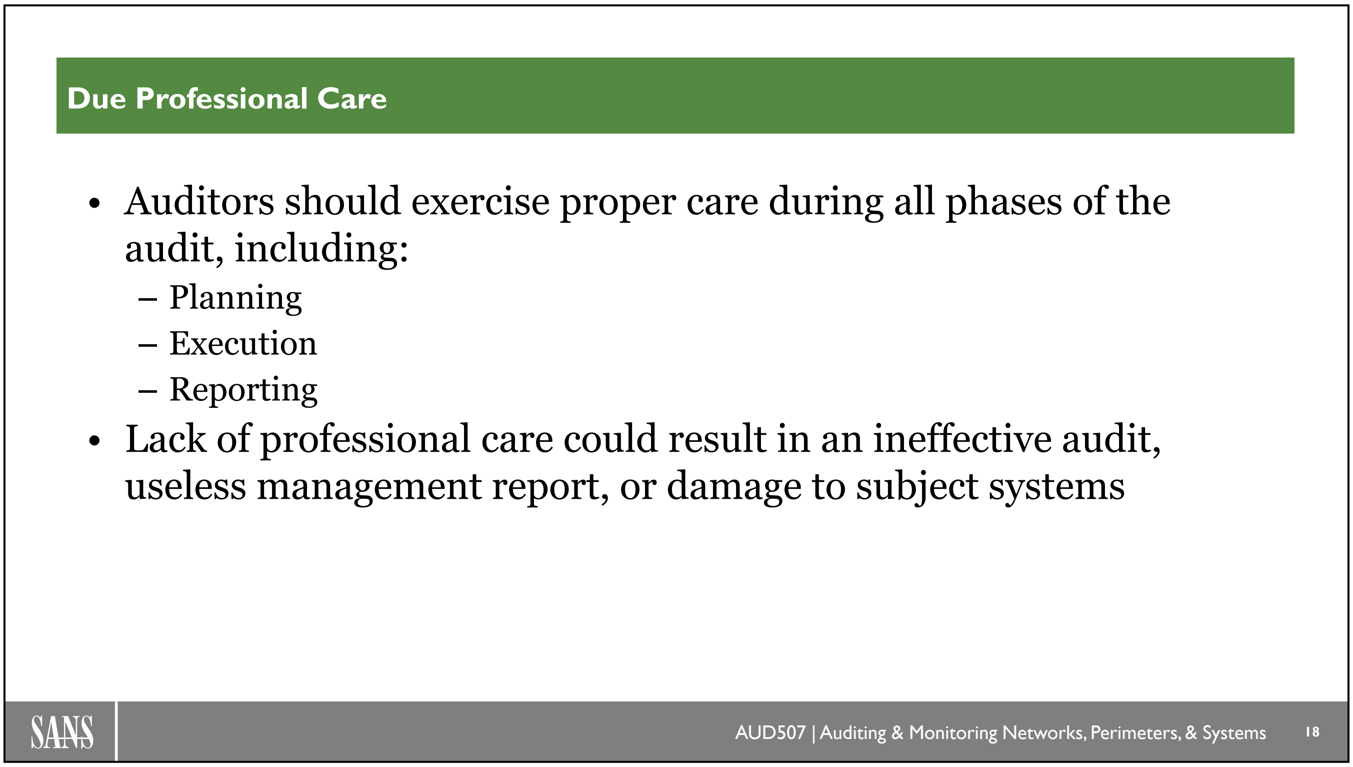

Due Professional Care

- Auditors should exercise proper care during all phases of the audit, including:
  - Planning
  - Execution
  - Reporting
- Lack of professional care could result in an ineffective audit, useless management report, or damage to subject systems

The IT auditor is a professional who is called upon to help an organization manage risk. If the auditor fails to exercise the care required of them, several bad things might happen:

- The audit might not meet the organization's needs. In this case they will have wasted time and money with only erroneous results to show for it.
- The audit might fail to uncover material risk, resulting in a catastrophic event in the future which may have been avoided.
- The auditor might uncover legitimate risks and communicate the findings so poorly that the organization fails to understand and mitigate them.
- An inept auditor performing tests they don't understand could even cause damage to the systems they've been tasked to help defend.

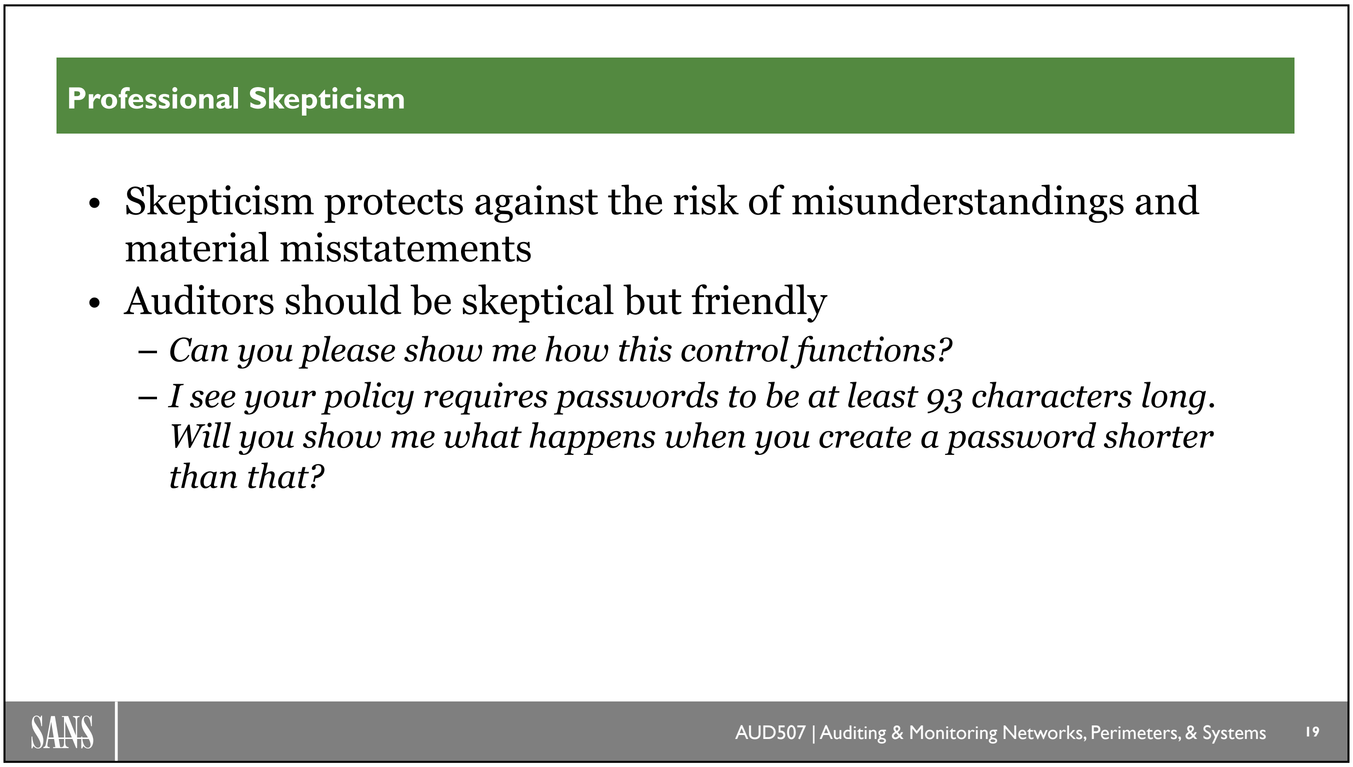

Professional Skepticism

- Skepticism protects against the risk of misunderstandings and material misstatements
- Auditors should be skeptical but friendly
  - *Can you please show me how this control functions?*
  - *I see your policy requires passwords to be at least 93 characters long. Will you show me what happens when you create a password shorter than that?*

Auditors should maintain a friendly but consistent “show me” mentality as they perform their duties. It is important to verify that systems, processes, and controls have been adequately explained and that these things actually exist and work correctly.

Both words in the term “professional skepticism” are important. Auditors should be professional, polite, and non-confrontational when assessing controls and interviewing staff. They should also show their skepticism by seeking to prove that controls work as designed. We have seen many instances where viewing the settings for a control indicated that everything was fine, only to do a substantive test and find that those settings were not actually applied or were applied incorrectly.

We frequently use the idea of “stimulus-response” testing to validate how a control actually *functions* in practice.

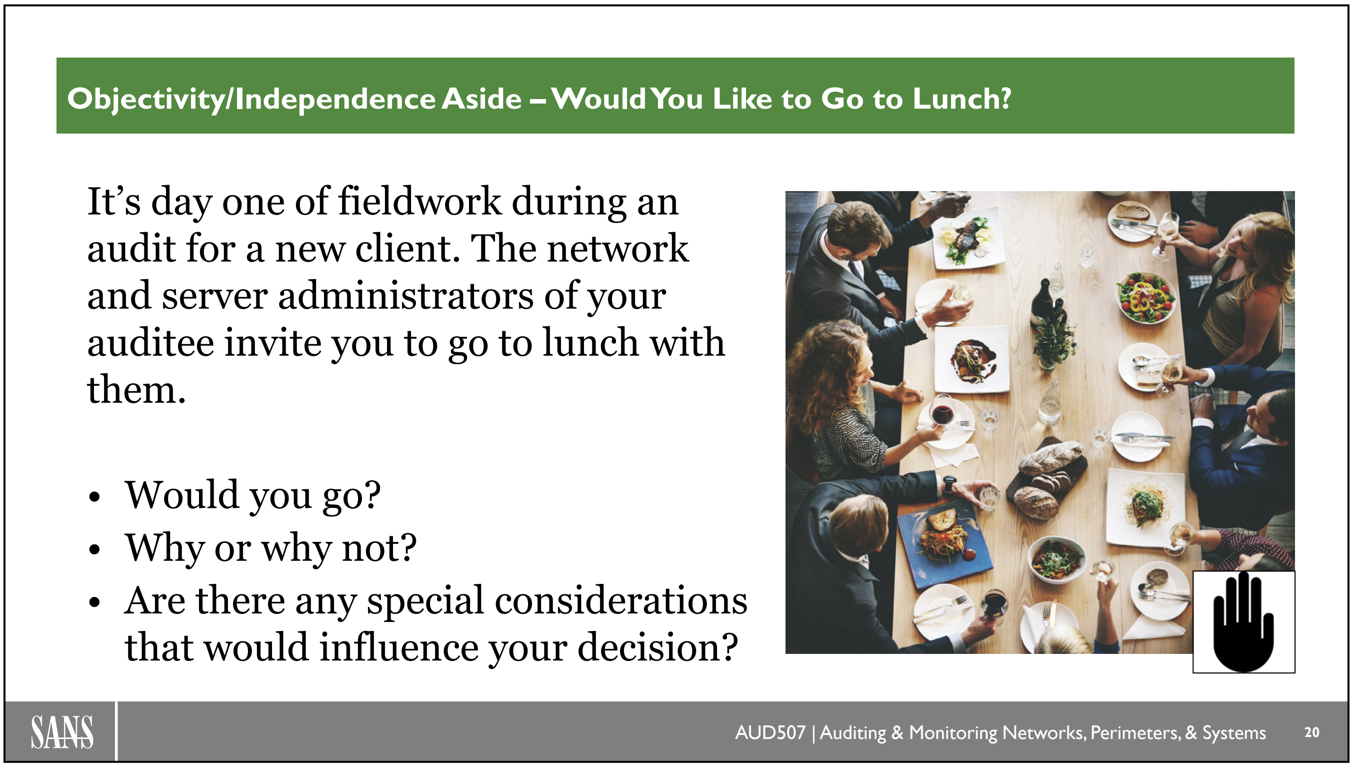

Objectivity/Independence Aside – Would You Like to Go to Lunch?

It’s day one of fieldwork during an audit for a new client. The network and server administrators of your auditee invite you to go to lunch with them.

- Would you go?
- Why or why not?
- Are there any special considerations that would influence your decision?

How do I maintain a cordial professional relationship with the staff of my auditee and stay independent and objective? What is the correct balance?

Take a little time to consider your answers to the questions on the slide. Here are some things to consider and balance as you make your decision:

- Would this help to build or maintain a good working relationship with the client?
- Would it jeopardize my independence or objectivity?

If you opt to accept the lunch invitation, think about these things:

- Who should pay for lunch?
- What topics of discussion should be avoided during the meal?

If you *would* accept the invitation, where is the line you would not cross? Would you accept an invitation to beer and karaoke after work?

Let’s wrap up the thought exercise with two observations. First, your answer to the questions will be colored by the working relationships and your organizational policies. Second, whether you go to lunch or not, you must be careful to avoid any appearance of impropriety. Auditors should be objective *in appearance* and objective *in fact.*

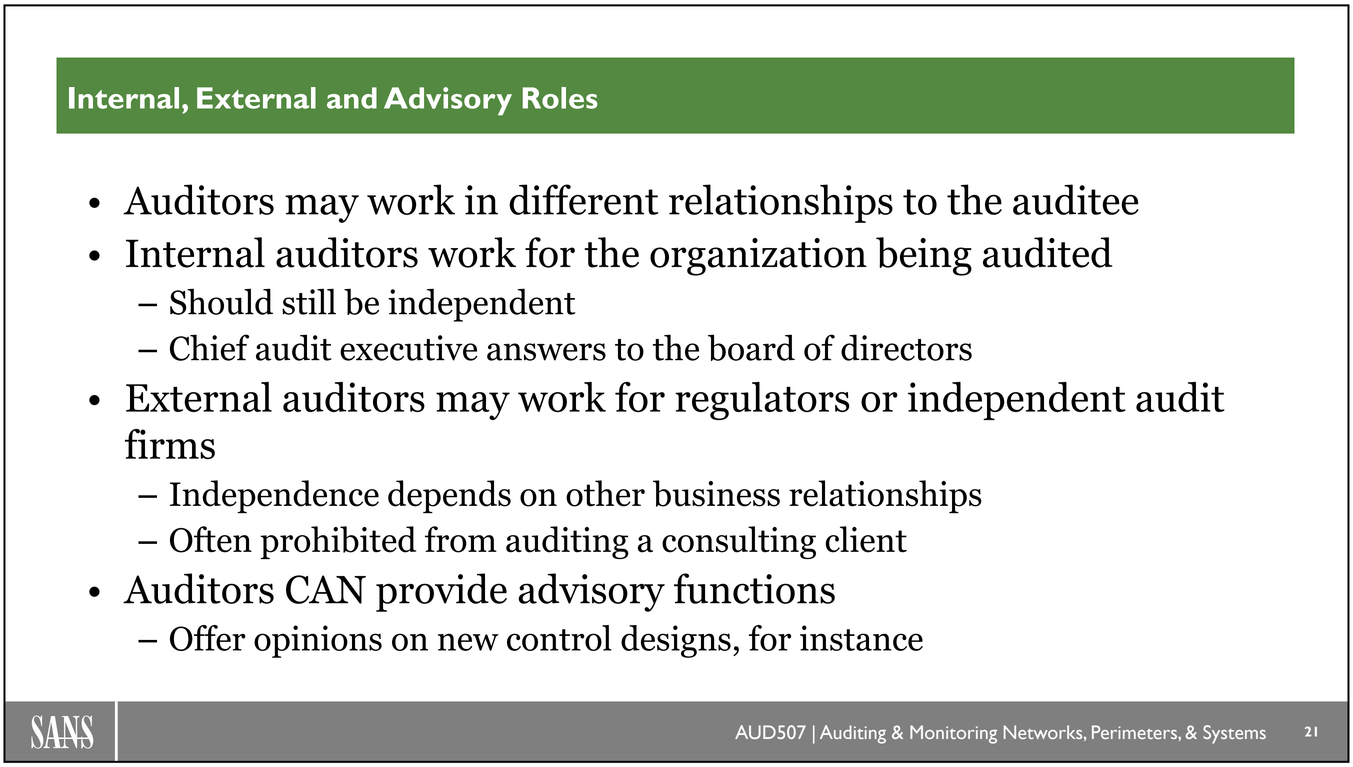

Internal, External and Advisory Roles

- Auditors may work in different relationships to the auditee
- Internal auditors work for the organization being audited
  - Should still be independent
  - Chief audit executive answers to the board of directors
- External auditors may work for regulators or independent audit firms
  - Independence depends on other business relationships
  - Often prohibited from auditing a consulting client
- Auditors CAN provide advisory functions
  - Offer opinions on new control designs, for instance

The relationship of the auditor to the auditee will help to determine how independence is maintained.

Internal auditors will often work for a department headed by a Chief Audit Executive (CAE). The CAE will often report directly to the audit committee of the board of directors, but sometimes will be placed under other executives, like the Chief Financial Officer. It is important for the organization to maintain independence of the audit function from the areas being audited. Internal auditors will often assist the organization in complying with requests and requirements of external auditors.

External auditors work for some organization other than the auditee. They may be employed by a regulator or an independent external audit or accounting firm. Independence is still important for external auditors, and external audit firms will frequently be prohibited from performing other work for the auditee.

Auditors are allowed to perform limited advisory functions for clients but cannot design or implement controls that the auditor would be expected to test later.

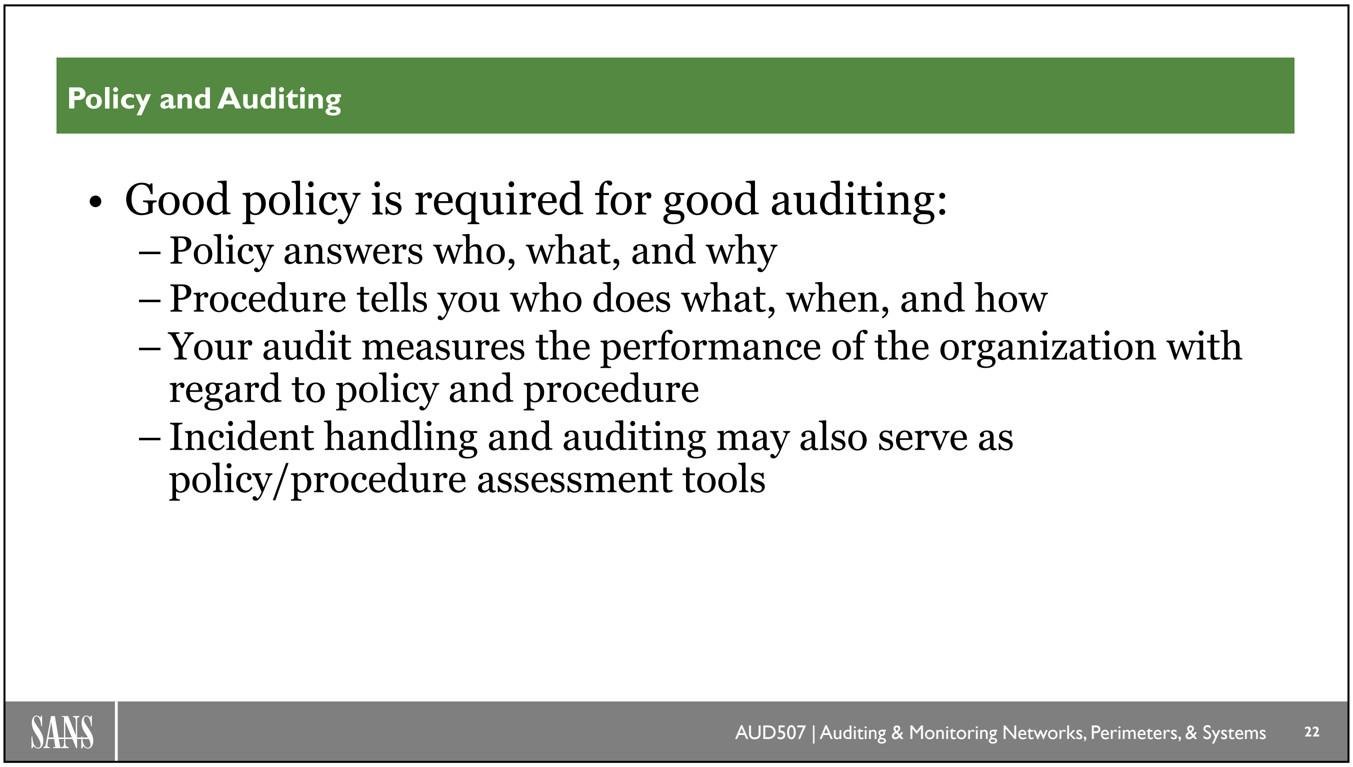

Policy and Auditing

- Good policy is required for good auditing:
  - Policy answers who, what, and why
  - Procedure tells you who does what, when, and how
  - Your audit measures the performance of the organization with regard to policy and procedure
  - Incident handling and auditing may also serve as policy/procedure assessment tools

Auditors are also often useful when it comes to the creation and definition of policy. We are not saying that auditors should write the policies! Rather, their input into the policy creation process is valuable.

If your organization does not involve auditors in this process, it is worth your time to convince it to include them! Your auditors are people who are accustomed to doing the necessary research and generally have a great deal of skill boiling down a policy into a practice. As a result, these folks can also quite often identify ambiguous or misleading policy statements and shoot them down before they ever become a part of an official document. Additionally, your auditors will likely have an eye on accountability to some degree, which is a major portion of policy definition. Let’s take a look at a few examples of auditing in a policy framework to see if we can get a feel for how they interrelate. In this case, we start with a policy statement and see if we can come up with a bullet item or two for an audit checklist or process.

Surprisingly, one of the most frequently asked questions (asked most frequently, but not necessarily by the majority of people taking this course) is, "What do I do if we don't have any documented policies and procedures?” The answer is deceptively simple: Get the organization to write some. Ultimately, auditing is measuring against a standard, and that standard is the organizational policies and procedures. If there are no documented policies and procedures, we cannot possibly perform an audit; we are limited to assessment. Further, it would be reasonable to say that if there are no documented policies and procedures, then management has abdicated its responsibility to control and manage risk, which it should articulate through the documentation of policies and procedures.

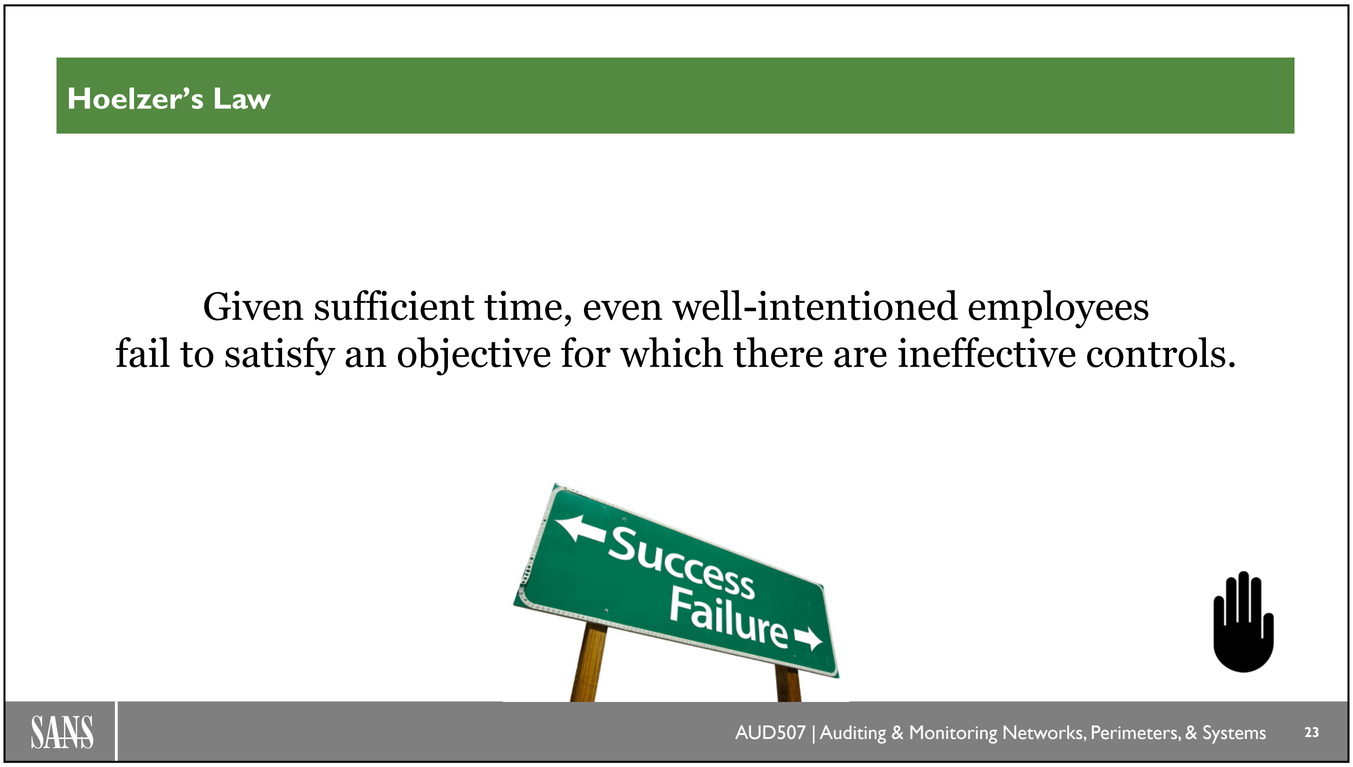

Hoelzer's Law

Given sufficient time, even well-intentioned employees fail to satisfy an objective for which there are ineffective controls.

Success
Failure

Let's take a quick step away from our procedure for a moment to have a look at Hoelzer's Law (named for the former author of this course). This law states that given sufficient time, even well-intentioned employees will fail to satisfy an objective for which there are ineffective controls. We must not just evaluate whether controls and objectives are specified, but we must make an objective assessment as to whether the controls in place are sufficient not just to enable but also to force the organization to meet an objective.

Why does this happen? It's quite simple. Your first answer might be that employees bypass controls because the controls are inconvenient, out of date, obstacles to efficient operations, or bureaucratic, or even because they are directed to bypass them by management. Ultimately, however, the real reason boils down to one simple fact: *Because they can.* If an employee cannot bypass a control, none of the previously mentioned reasons will come into play. Controls that can be bypassed or ignored are not effective and must be supplemented.

We can illustrate this using a simple test, which is found on the next slide.

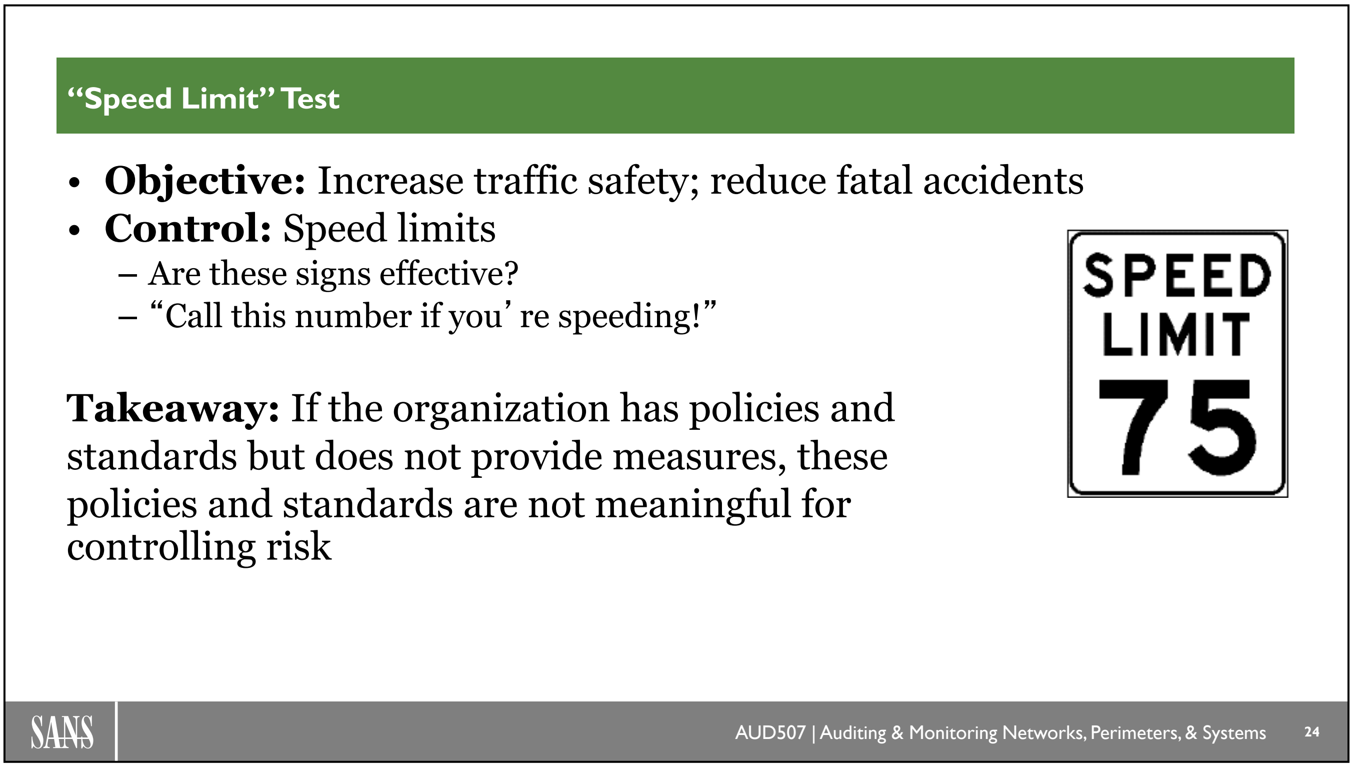

"Speed Limit" Test

- **Objective:** Increase traffic safety; reduce fatal accidents
- **Control:** Speed limits
  - Are these signs effective?
  - "Call this number if you' re speeding!"

Takeaway: If the organization has policies and standards but does not provide measures, these policies and standards are not meaningful for controlling risk

SPEED
LIMIT
75

The test that we perform is called the Speed Limit Test. We ask whether the controls in place have any reasonable chance of ensuring that the organization or those subject to the controls will be coerced into meeting the objectives set forth. If we consider speed limit signs found on highways and other roadways, the answer is clearly "no." These signs seem to be viewed as a suggestion for a minimum speed rather than a control for the maximum speed. For us to meet the objectives, we have to add additional controls, such as police with radar guns.

Imagine a policy that states that an employee who detects malware must call a certain number and speak to a certain person. This control has the objective of allowing the organization to respond better to malware incidents; however, it is not effective! Consider this in the context of the appropriate use policies of most organizations: If you encounter a piece of malicious code, most likely you were doing something you shouldn't have been doing. Now you're supposed to call a phone number and tell someone that you have violated those policies!

Whenever your controls fail the Speed Limit Test, you have found a strong indicator that there is a need for more controls. If your controls pass this test, then quite likely your controls are sufficient to meet the objectives, and there is no need for additional policy.

There's another takeaway from this: If your organization creates policies, procedures, and standards but fails to create measurement capabilities related to these, the controls become meaningless for risk prevention. Controls without measurement and enforcement mechanisms are effectively the same as speed limit signs with no speed cameras or police officers.

Communicating Risk: Auditing and Storytelling

- "Cooperative" audits enable remediation of thought:
  - They already know the facts
  - Contextualize the facts to change behavior

1983
What we say to dogs
Okay, Ginger! I've had it! You stay out of the garbage! Understand, Ginger? Stay out of the garbage, or else!
What they hear
blah blah GINGER blah blah blah blah blah blah blah blah blah GINGER blah blah blah blah blah blah...

Copyright Larson, 1983

You're probably familiar with the saying, "Give a man a fish, feed him for a day; teach a man to fish, feed him for a lifetime." That's the general idea here. Imagine two scenarios: In the first scenario, you conduct a traditional audit; you review the system with an admin account or interview the system administrator and make notes on your checklist. In the second scenario, you perform the type of cooperative audit that we recommend; you sit with the administrator and look at the checklist together. Whenever something needs to be done, the administrator performs the audit action, and you take notes on what occurs. You and the system administrator build a rapport, and the administrator feels free to comment on the checklist, even offering suggestions about it.

In the first scenario, the traditional audit, how does the administration respond to the auditor attempting to explain why something is risky? Those who have been administrators can tell you that they hear, "Blah blah blah blah **risky** blah blah blah."

In the second scenario, you and the administrator are communicating! Better still, because it's a cooperative audit, the administrator is free to comment on aspects of the audit that he doesn't understand or agree with. Also, if you take my advice, there's a references section on the checklist. This way, it's no longer the auditor telling the system administrator why or how he should do something; you can now suggest that you both look up the reference and see why the checklist (not the auditor) says you should do whatever is in question. You also have the opportunity to tell him stories. Remember that a primary objective is to reduce risk. A great way to reduce risk is by raising awareness. If you tell a story about someone he doesn't know (although it can be about you…) working for some other company, that will enable him to see the potential damage to the administrator; you don't have to tell him how risky it is; he can figure it out for himself!

Consider this: You discover that an administrator is browsing the internet while logged into an administrative-level account and you point out that this violates policy. Are you telling him something he doesn't already know? Or is he well aware of the policy and simply chooses to ignore it? Rather than telling him the facts ("That's against policy…"), you might use a story to contextualize the risk into something that may matter to him personally, allowing him to choose to modify his behavior.

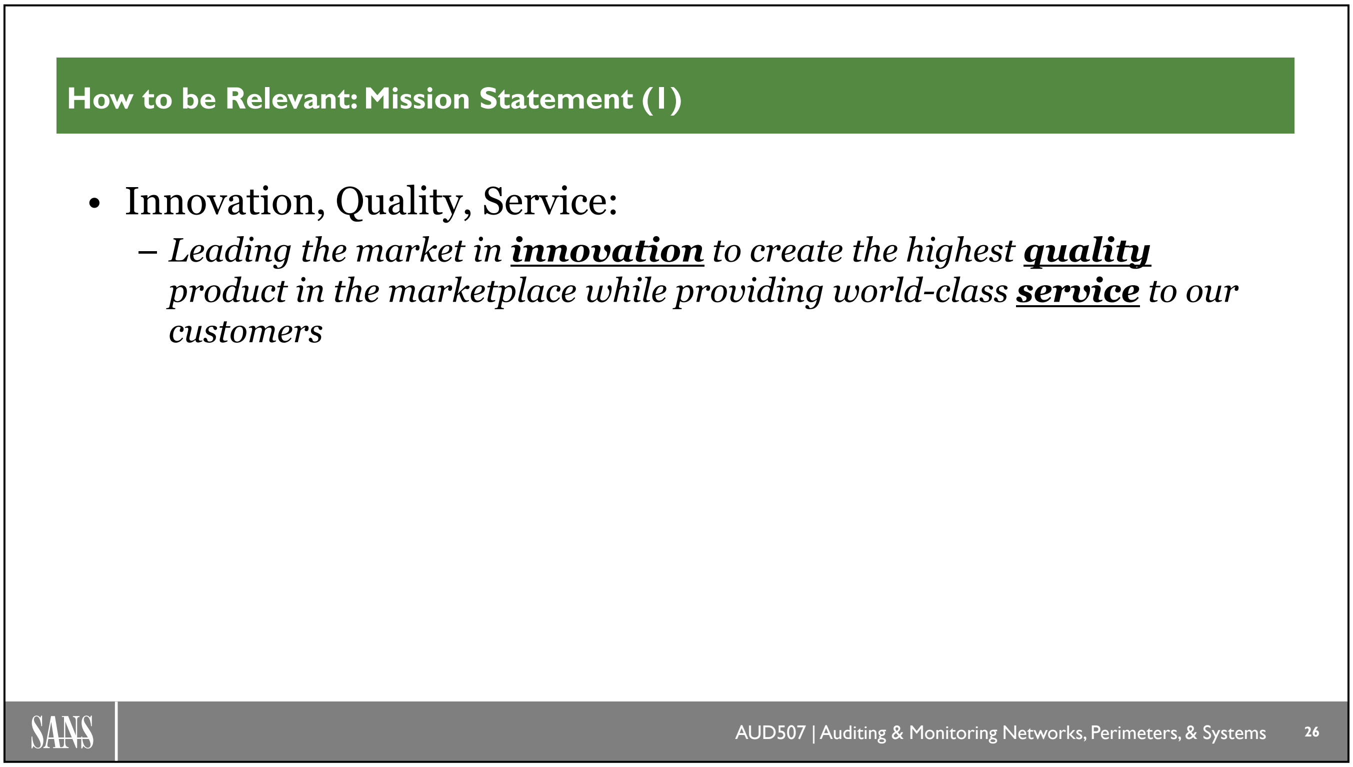

How to be Relevant: Mission Statement (1)

- Innovation, Quality, Service:
  - *Leading the market in **<u>innovation</u>** to create the highest **<u>quality</u>** product in the marketplace while providing world-class **<u>service</u>** to our customers*

Take this mission statement as an example. We hope that you agree that this is a nice-sounding mission statement. It is a forward-looking statement that appears to be largely customer-focused. Unfortunately, if you try to create controls and objectives around this, you will run into walls quickly. Part of the reason is that although this is a wonderful mission statement that should assist employees in making correct decisions in the face of a missing control, it's likely not honest.

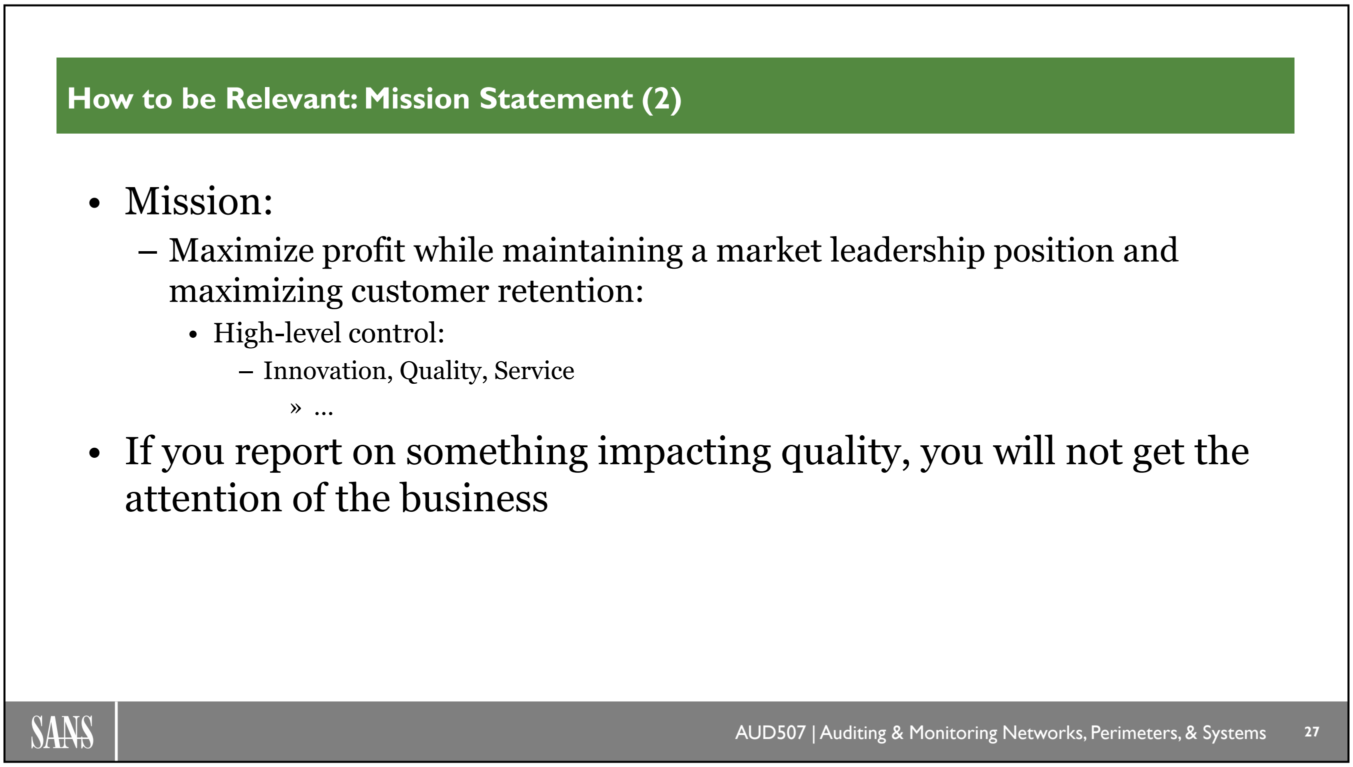

How to be Relevant: Mission Statement (2)

- Mission:
  - Maximize profit while maintaining a market leadership position and maximizing customer retention:
    - High-level control:
      - Innovation, Quality, Service
        - » …
- If you report on something impacting quality, you will not get the attention of the business

SANS AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 27

For a private enterprise, it is almost always safe to say that the mission of that enterprise is to make money. However, this is not always true. For example, if you consider a philanthropic organization such as the Carnegie Corporation of New York, its mission is, “To do real and permanent good” (https://u.aud507.com/1-29). If you look at the Bill and Melinda Gates Foundation, it has a mission that “…works to help all people lead healthy, productive lives” (https://u.aud507.com/1-30). These certainly aren’t “money-focused,” as is the mission statement on our slide, but they are powerful for us as auditors.

You may find that many profitable enterprises use as a mission statement something that is an expression of the *strategy* that the organization uses to achieve its actual mission. Understanding the true mission and the enterprise strategy is quite important, however, since it allows us to identify connections between the enterprise’s mission, sub-objectives, and controls. We should find that these can be organized into a pyramidal or other hierarchical structure, allowing us to identify the framework used by the enterprise to reach its goals.

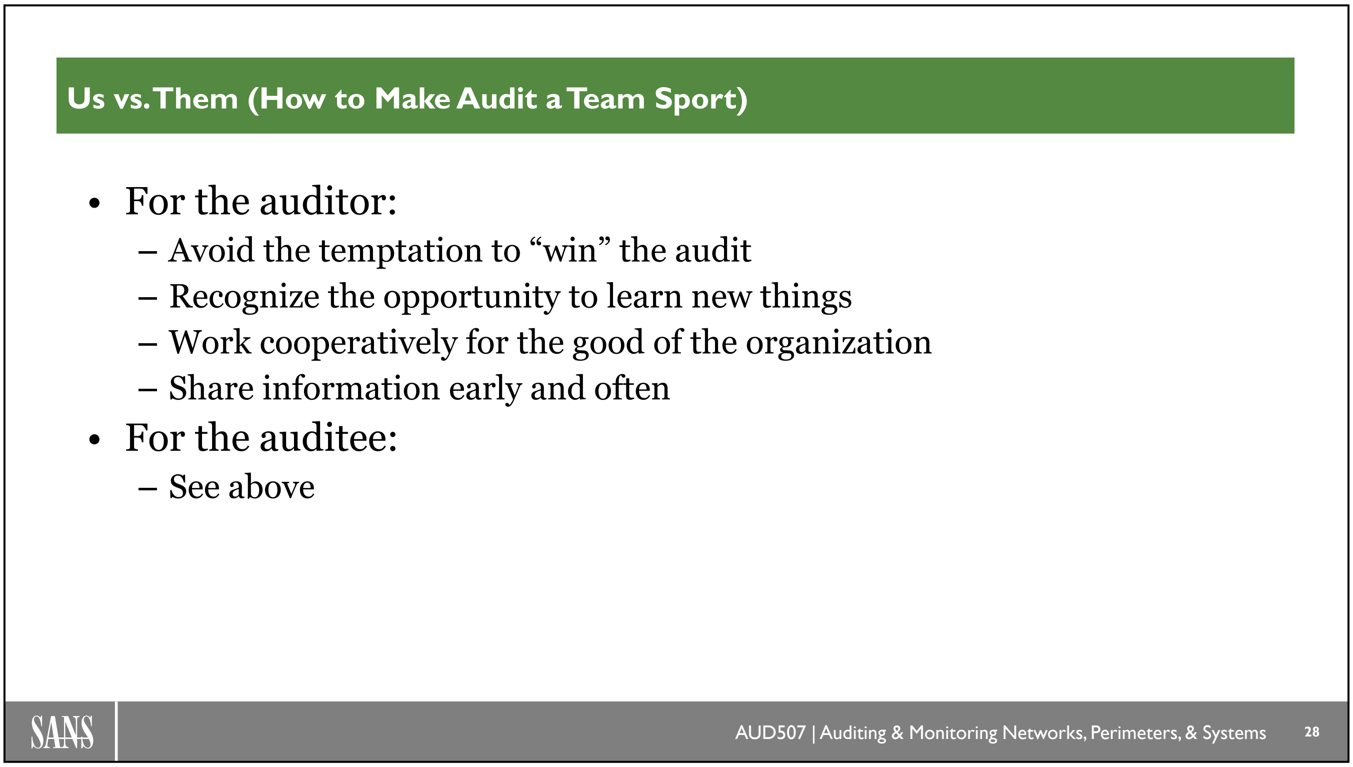

Us vs. Them (How to Make Audit a Team Sport)

- For the auditor:
  - Avoid the temptation to “win” the audit
  - Recognize the opportunity to learn new things
  - Work cooperatively for the good of the organization
  - Share information early and often
- For the auditee:
  - See above

A frequent problem that arises during audits is that the auditors (and the auditee, for that matter) try to “win” the audit. For the auditor, this might mean trying to come up with audit findings at all costs. For the auditee, it involves giving vague or scripted answers to questions rather than honestly facing the risk the organization might face.

To avoid this temptation as an auditor, I have a few recommendations. Avoid the desire to win and replace it with a desire for the organization to benefit from the audit process. Realize that the auditee’s staff are likely subject matter experts in their respective fields and embrace the chance to learn from them. Make every effort to perform a cooperative audit and to involve the auditee’s staff in the process wherever appropriate. Share information with the auditee on a frequent basis. On longer engagements, it is common to perform a daily status update meeting (keep it short, please) to keep everyone apprised of progress and issues to date.

If you are the auditee, all the same things apply. The audit should be a cooperative process, and you are key to making that happen.

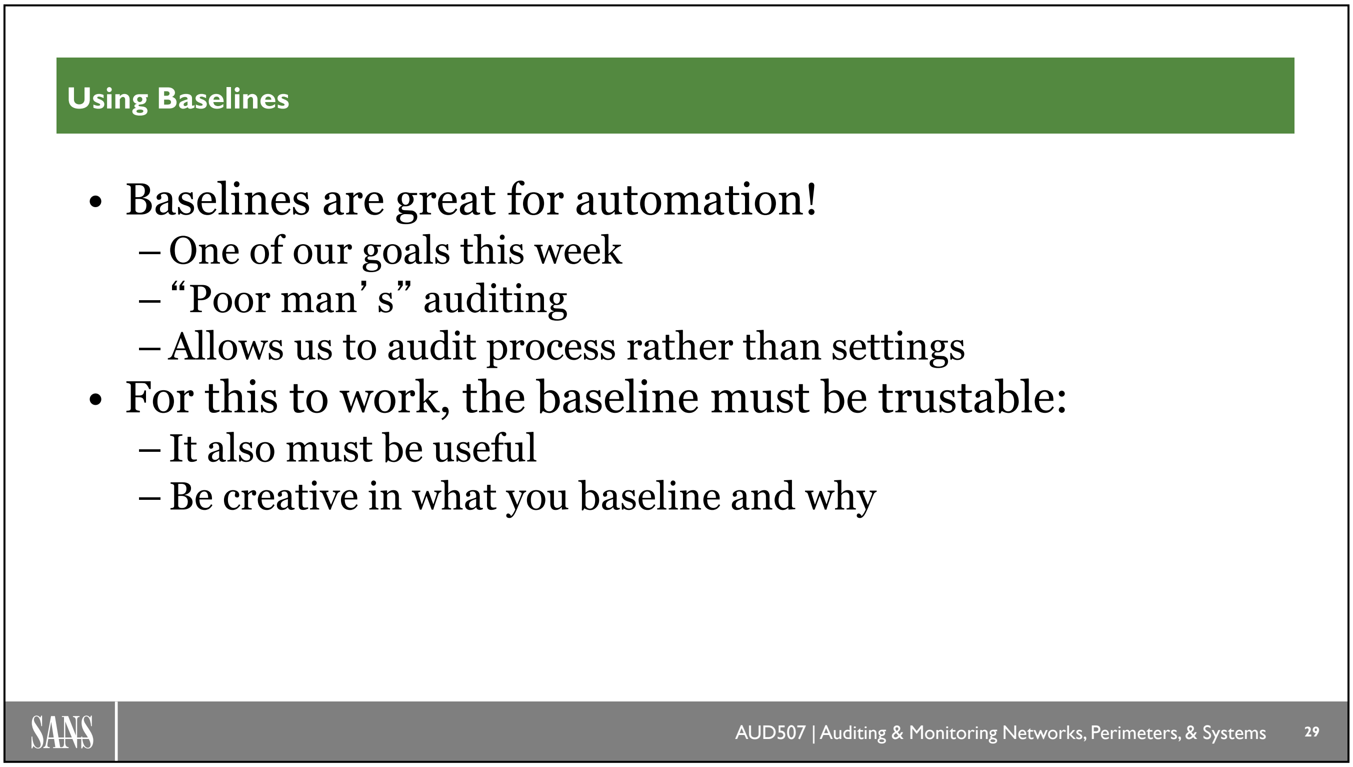

Using Baselines

- Baselines are great for automation!
  - One of our goals this week
  - “Poor man’ s” auditing
  - Allows us to audit process rather than settings
- For this to work, the baseline must be trustable:
  - It also must be useful
  - Be creative in what you baseline and why

Baselines will be discussed throughout the week. The reason for this is that one of our major goals is to try to present audit techniques that you can use to automate much of what an auditor would want to evaluate. There are two great reasons why any auditor should consider doing this.

First, if we automate the auditing of settings and things of that sort, it allows us to stand back from the intricate details and perform an audit of the process rather than the settings. This is a great thing! When our settings are in good order, process should be our focus as auditors. Process is where the real problems usually are.

The second reason for this is that we usually have vast numbers of computers in our environments, and manually testing every one of them is not only time-consuming but may not be possible. Automation usually reduces this problem nicely.

For this to work well, though, we have to trust the baseline. When the baseline is taken, we must know that the system being baselined is in a known good state. It also means that when we use that audit to check this or any other machine, we must know that the baseline data has not been modified in any way.

As you work with baselining, be creative! Don’t limit yourself only to settings. Instead, allow your baselines to be more flexible. We repeatedly look at how to do this during the week.

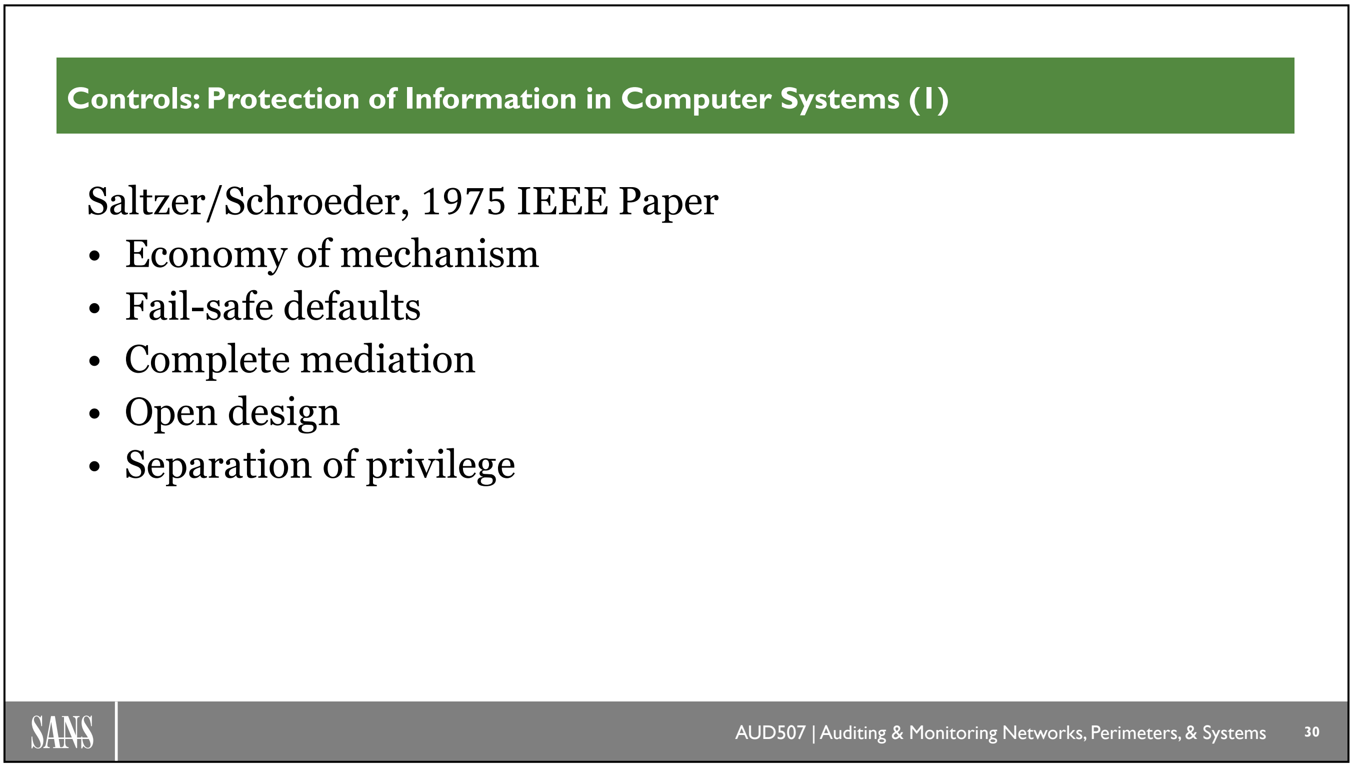

Controls: Protection of Information in Computer Systems (1)

Saltzer/Schroeder, 1975 IEEE Paper

- Economy of mechanism
- Fail-safe defaults
- Complete mediation
- Open design
- Separation of privilege

Throughout the week, we will be referencing the foundational Information Security referenced on this slide. These concepts were first presented in a 1975 paper in *Proceedings of the IEEE.* While we'll discuss these more this week, it's worth taking time to give brief definitions now.

- **Economy of mechanism:** Simple systems are easier to secure. Conversely, overly complex systems often have well-hidden flaws.
- **Fail-safe defaults:** Security systems should fail into a secure state. A firewall which is overloaded with traffic should drop new traffic, rather than forwarding traffic it shouldn't.
- **Complete mediation:** All authentication and access control decisions should go through a single decision-making process to ensure that all are handled in the same way.
- **Open design:** The opposite of "security by obscurity" – systems with widely understood designs are more likely to be secure than those with closed, proprietary designs.
- **Separation of privilege:** (Often called segregation of duties) Sensitive or risky tasks should require the action of more than one person to protect against fraud and abuse.

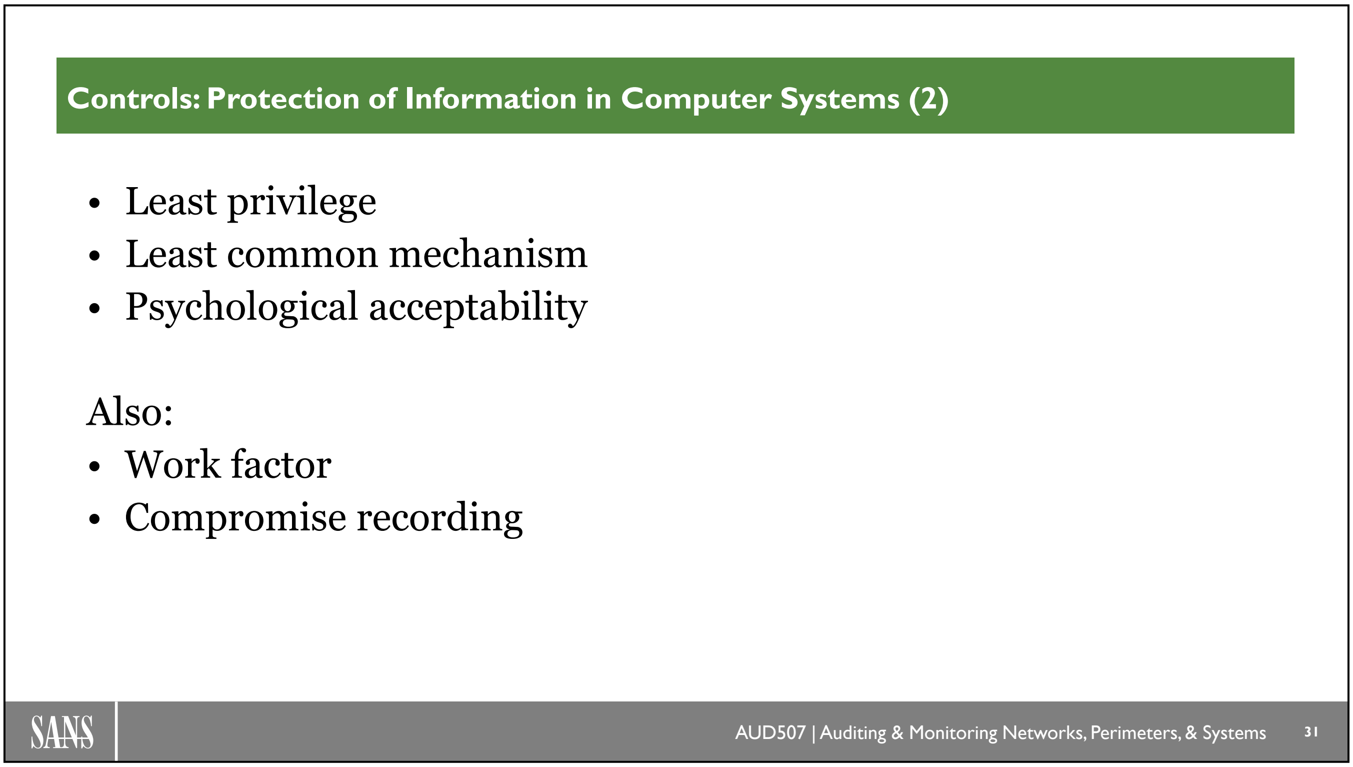

Controls: Protection of Information in Computer Systems (2)

- Least privilege
- Least common mechanism
- Psychological acceptability

Also:

- Work factor
- Compromise recording

- **Least privilege:** Users and systems should be given exactly the amount of access they need to data or other systems. Too little, and they can’t do their job; too much and they may abuse the system.
- **Least common mechanism:** Access to resources should not be shared among subjects. For example, each user has a unique username and password to ensure that only they can access the resources to which they have been given access.
- **Psychological acceptability:** Users must accept security controls as being easy to use, sensible, and not preventing legitimate access. When this principle is not followed, users are tempted to bypass the controls.

The paper also discussed work factor (the cost of circumventing a control should be beyond an attacker’s reach) and compromise recording (the system should continue to log activity even after a malfunction or compromise). This paper is available at:

https://u.aud507.com/1-1

Exercise 0 - Student Lab Setup

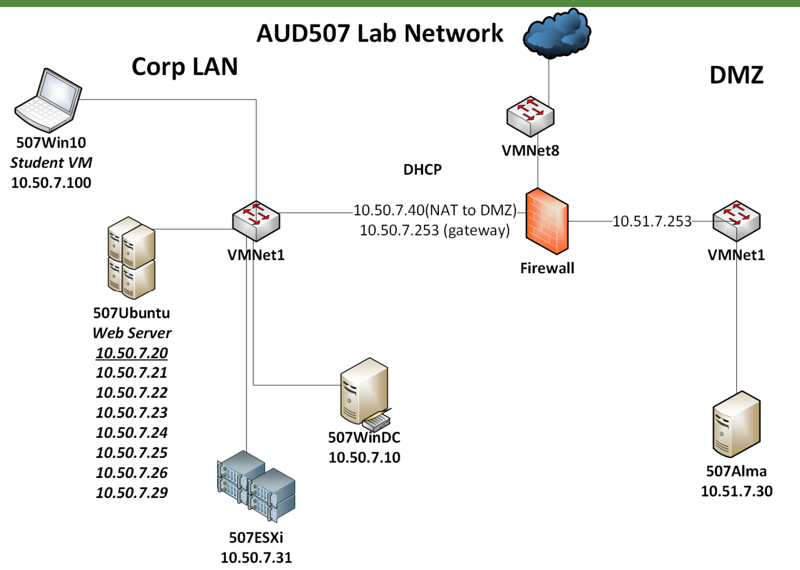

This page intentionally left blank.

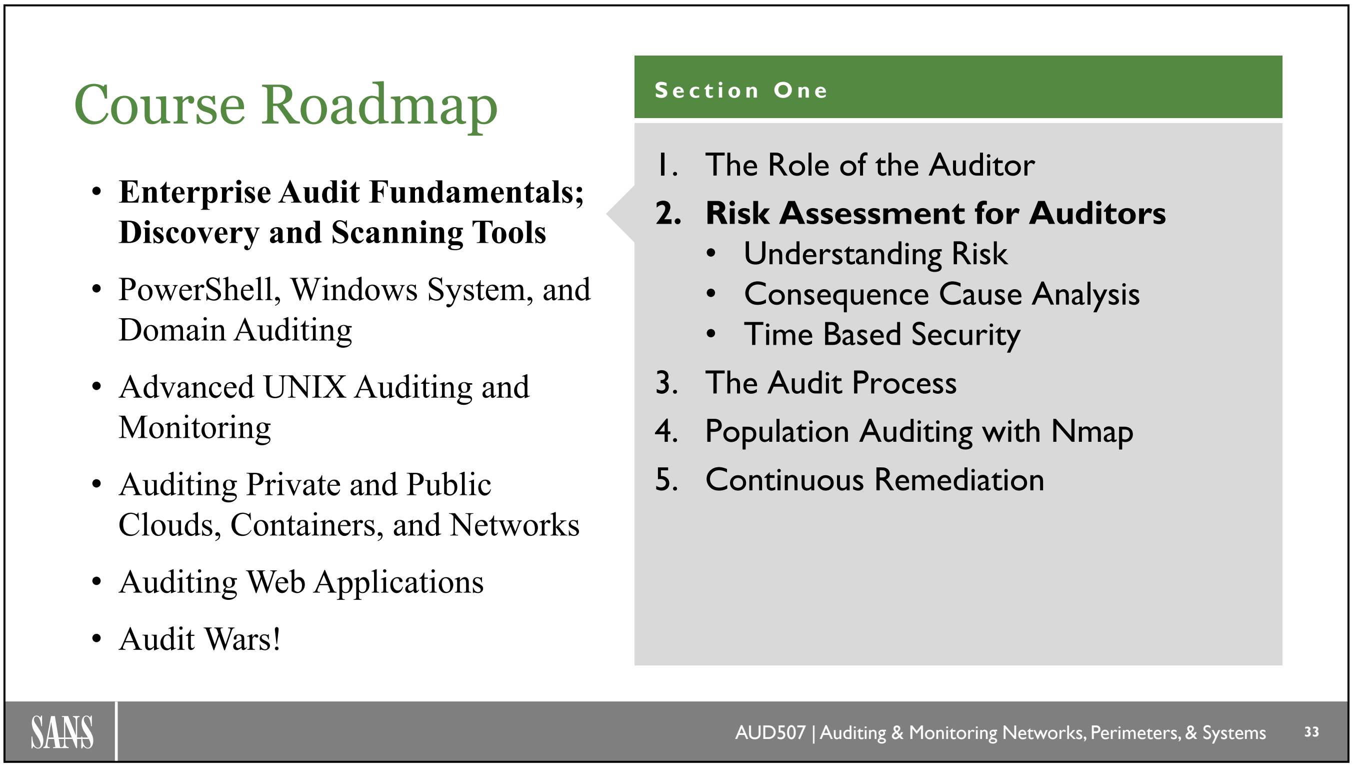

Course Roadmap

- **Enterprise Audit Fundamentals; Discovery and Scanning Tools**
- PowerShell, Windows System, and Domain Auditing
- Advanced UNIX Auditing and Monitoring
- Auditing Private and Public Clouds, Containers, and Networks
- Auditing Web Applications
- Audit Wars!

Section One

1. The Role of the Auditor
2. **Risk Assessment for Auditors**
   - Understanding Risk
   - Consequence Cause Analysis
   - Time Based Security
3. The Audit Process
4. Population Auditing with Nmap
5. Continuous Remediation

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 33

This page intentionally left blank.

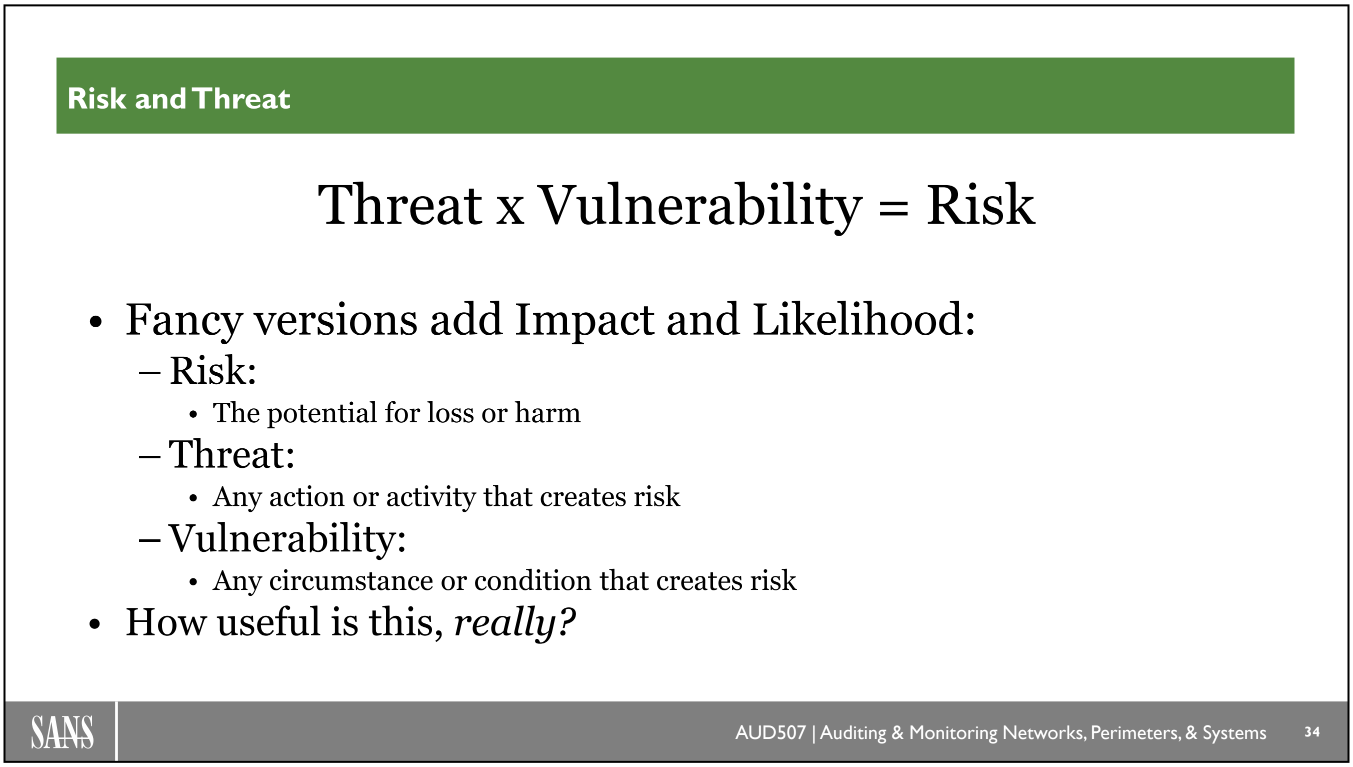

Risk and Threat

Threat x Vulnerability = Risk

- Fancy versions add Impact and Likelihood:
  - Risk:
    - The potential for loss or harm
  - Threat:
    - Any action or activity that creates risk
  - Vulnerability:
    - Any circumstance or condition that creates risk
- How useful is this, *really?*

We expect that just about everyone who takes this class has come across this generalized “formula” for calculating risk. It simply recognizes that risk, the potential for loss or harm, is the product of the amount of threat (actions or activities that create the potential for loss or harm), and the vulnerability (circumstances or conditions that create the potential for loss or harm).

Most risk assessments used in the Information Technology (IT) field can be traced back to this simple formula. It is certainly possible (and useful!) to add other terms, such as the likelihood of an event or the impact of an exposure, to arrive at a more useful risk value. In the end, though, all that this tells us is that something is or is not “risky.” Let’s explore this first and then examine two formal risk assessment approaches that allow us to produce something far more useful than a sense that something may be risky.

While this basic approach is good for a high-level understanding of risk, it doesn’t tell us what to do about the risk or why the risk exists. Since it’s so widely used, we’ll speak about it briefly, but the risk assessment strategies which follow are far more useful for making tactical decisions or recommendations for the remediation of risk.

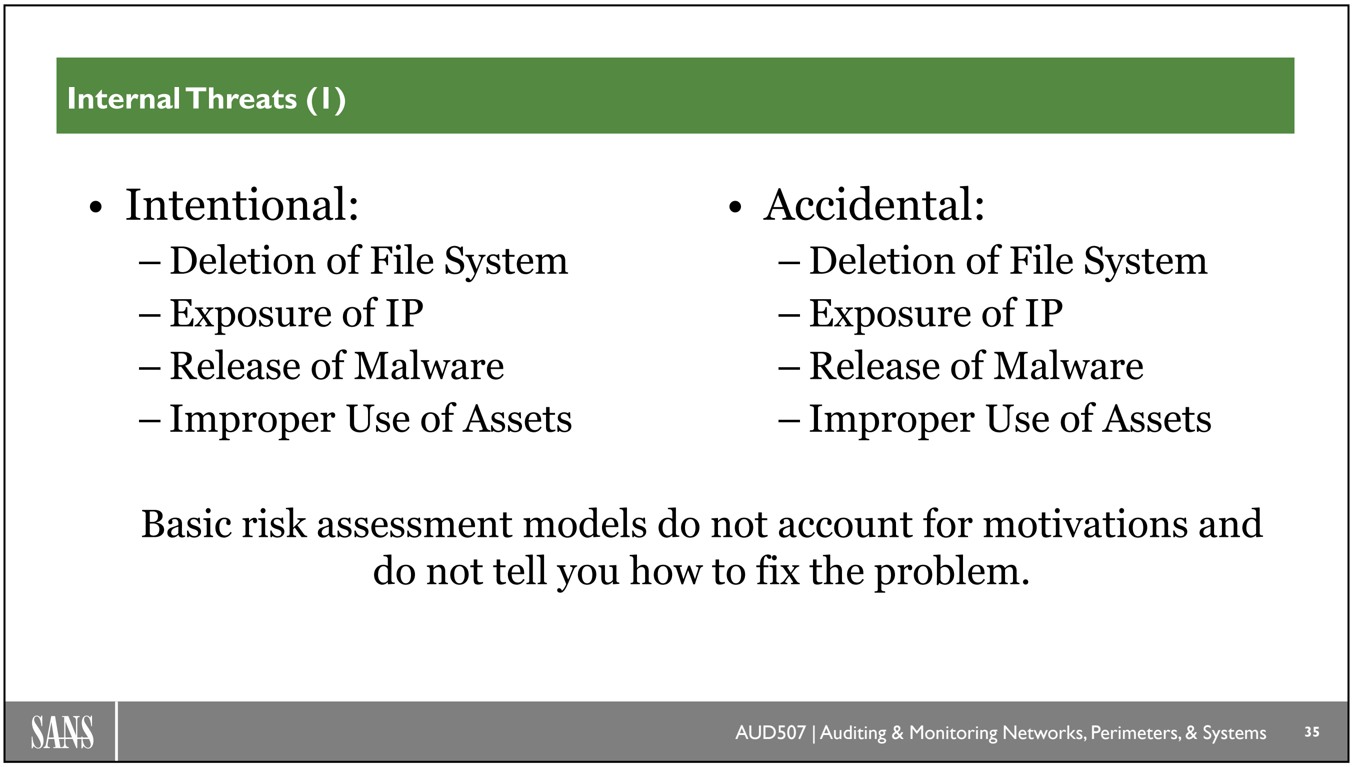

Internal Threats (1)

- Intentional:
  - Deletion of File System
  - Exposure of IP
  - Release of Malware
  - Improper Use of Assets
- Accidental:
  - Deletion of File System
  - Exposure of IP
  - Release of Malware
  - Improper Use of Assets

Basic risk assessment models do not account for motivations and do not tell you how to fix the problem.

Here are some examples of internal threat vectors. An interesting thing that you might notice about this slide is that the intentional and the accidental threats are exactly the same!

How do we classify a threat as being intentional or accidental? The motivation and orientation of the individual performing the threatening behavior are the answers. Think about the root cause! When we discover an internal individual bypassing the controls enforced by the firewall because he cannot accomplish a business-related function as a result of those controls, what is the root cause? Although the individual is still in violation of the policies, clearly there is a need for remediation elsewhere in the information control infrastructure. Although this individual was still "intentional" in his behavior, his orientation (accomplishing a business requirement) indicates that he likely did not realize that this behavior was a threat and could create a vulnerability.

However, a firewall administrator who punches holes in the firewall to allow himself to violate the appropriate use policy in the organization is clearly intending to violate policies in such a way that creates the potential for loss or harm for the organization.

Ultimately, however, the major failing of these really basic risk models is that they only serve to tell us the magnitude of the risk. Rarely, if ever, do they tell us how to remediate the risk.

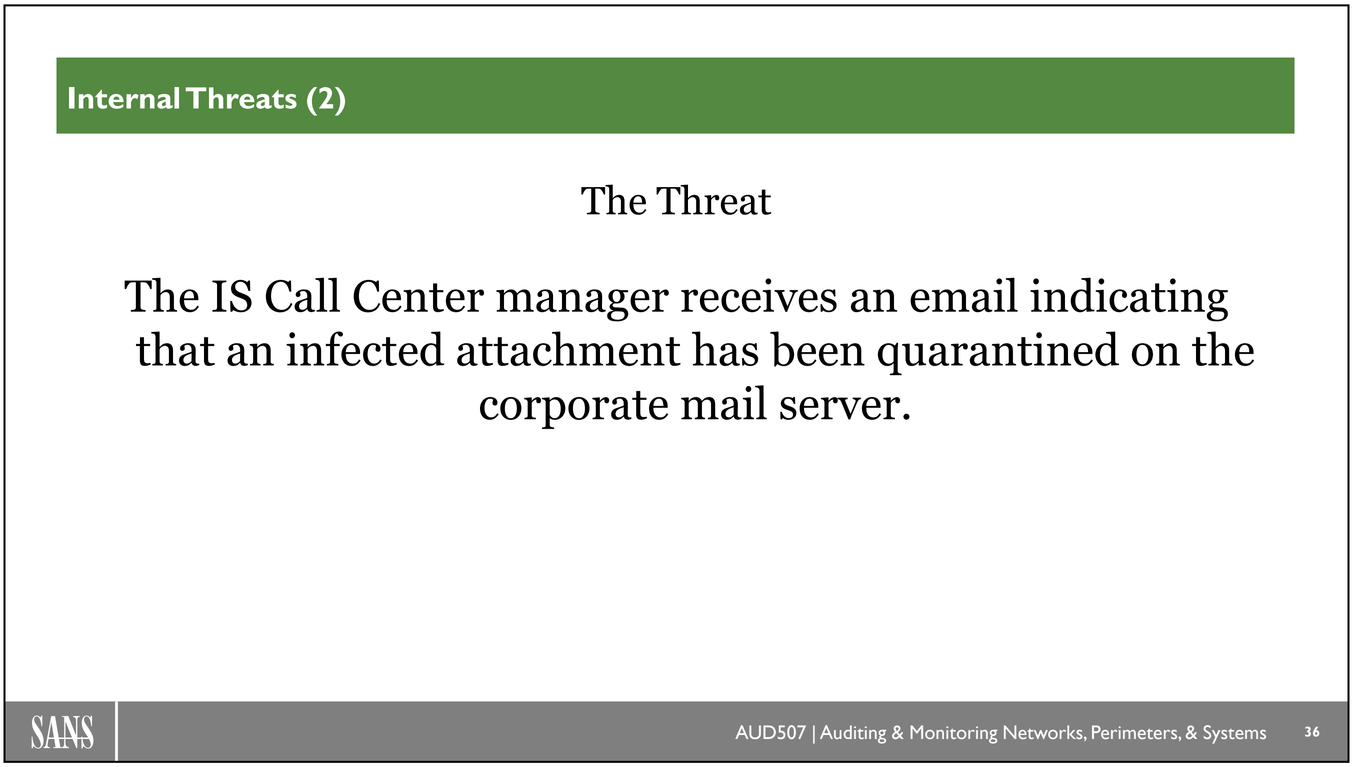

Internal Threats (2)

The Threat

The IS Call Center manager receives an email indicating that an infected attachment has been quarantined on the corporate mail server.

Let’s walk through an example of an internal security event that really occurred at an organization. The IS Call Center manager received an email in his mailbox from the Antigen virus scanner running on the corporate mail server. Antigen informed him that it had detected an email message infected with a virus that it couldn’t clean, so the email had been quarantined on the mail server. Included in the mail message were links to the quarantine area in the event the manager wanted to examine what had been captured, especially if he wanted to forward it to the antivirus manufacturer so that a new signature could be created to identify the virus.

What’s the threat? The threat is that this internet worm is attempting to enter the corporate infrastructure via email.

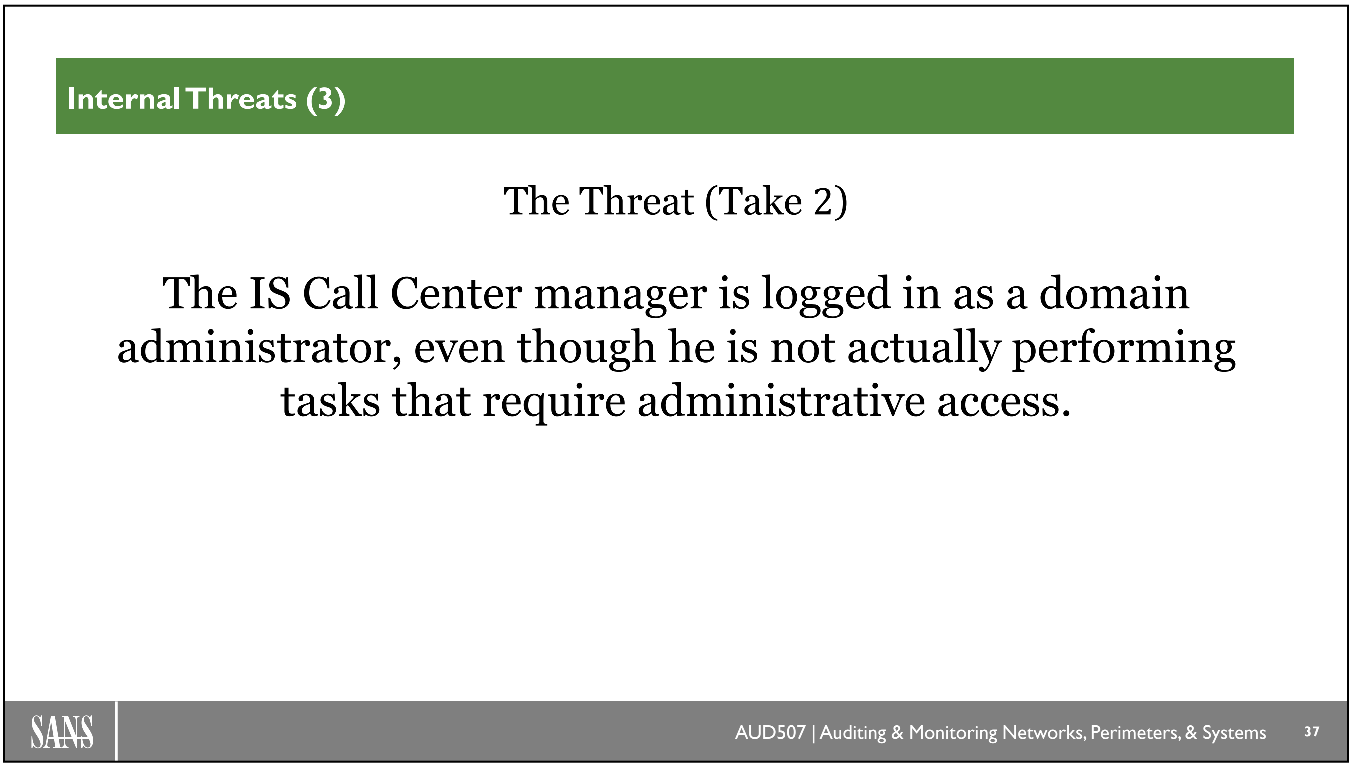

Internal Threats (3)

The Threat (Take 2)

The IS Call Center manager is logged in as a domain administrator, even though he is not actually performing tasks that require administrative access.

Enter threat #2. The IS Call Center manager is logged in with his administrator account to the Windows Domain. In this instance, it is important to note that this is threatening behavior because it is a violation of the Principle of Least Privilege, which requires that users have only the minimum privilege level required to perform the current task. We should also note that this organization did not have a policy requiring that this best practice be adhered to.

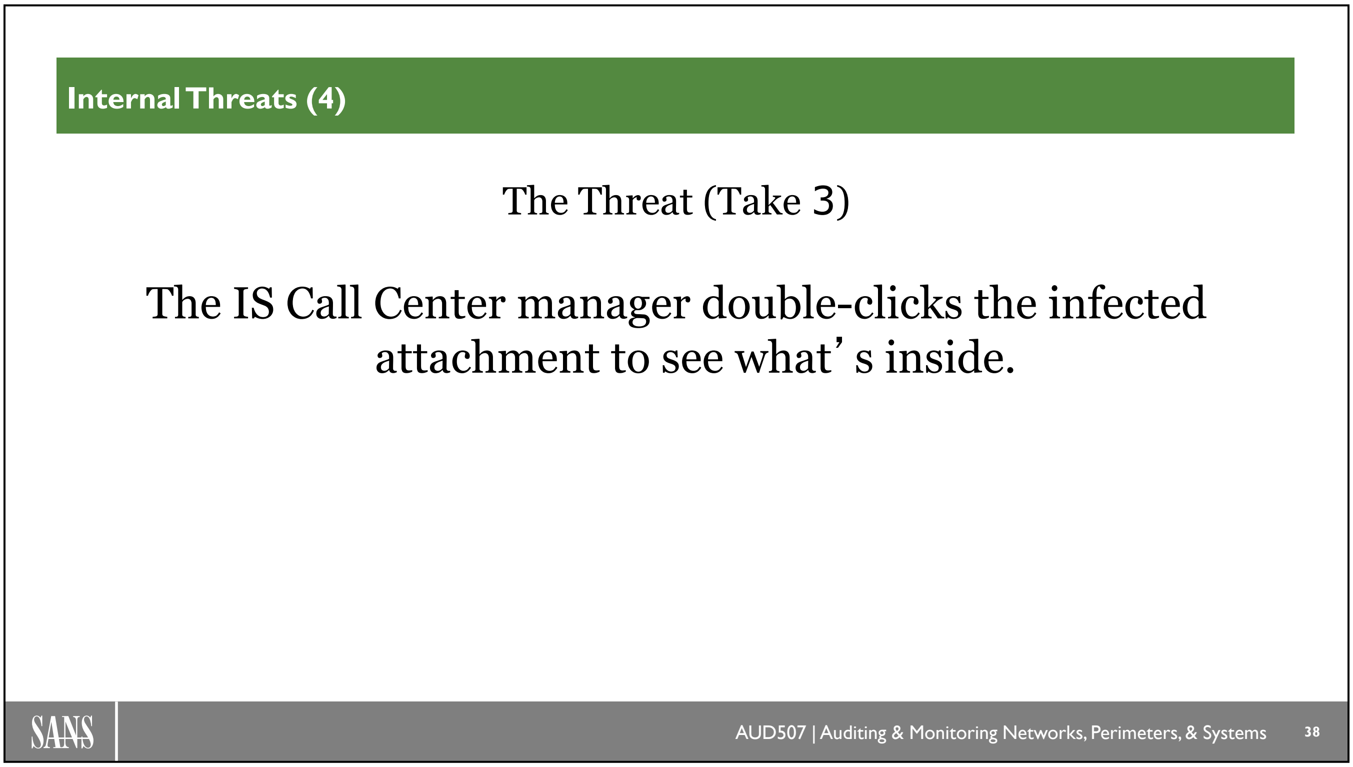

Internal Threats (4)

The Threat (Take 3)

The IS Call Center manager double-clicks the infected attachment to see what’ s inside.

Now, the IS Call Center manager that we’ re talking about is a fairly competent guy. He’ s been working in Information Technology for approximately six years and started out in a technical position before being promoted to manage the call center. As is true of all of us at one time or another, however, he experienced something of a brain freeze. Most of us would look at the word “quarantine” and know that we don’ t want anything to do with what’ s inside, but for some reason his mind connected “quarantine” with “cured,” so he decided to open the attachment up and see what was inside. Would you also classify this as a threat?

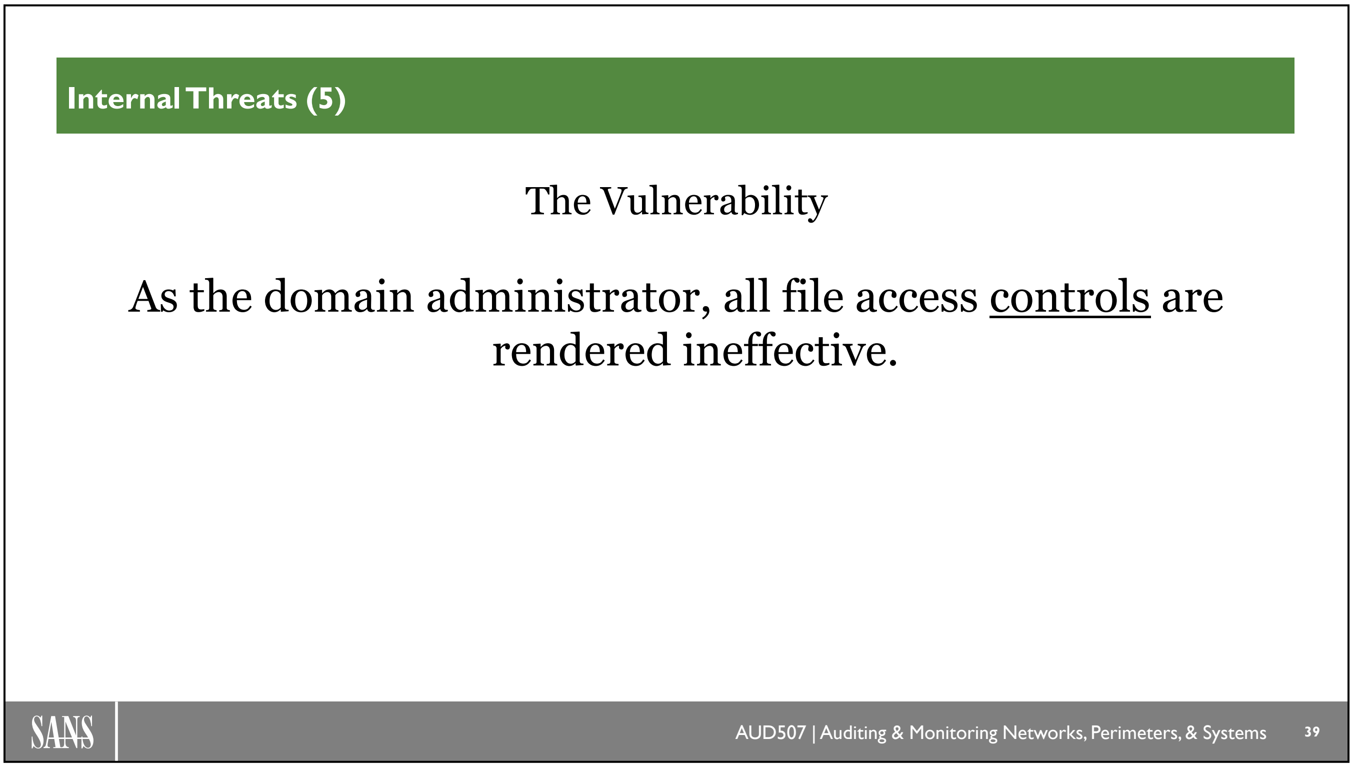

Internal Threats (5)

The Vulnerability

As the domain administrator, all file access controls are rendered ineffective.

Now let’s connect the rest of the dots. We’ve got a lot of threat here. How about some vulnerability and some loss or exposure?

For vulnerability, the window that the manager opened was logged in as administrator. The administrator user of the Active Directory has full rights and permissions to all files in the system. In Information Assurance (IA) terminology, the admin accounts defeat all controls, preventing the systems from meeting their security objectives.

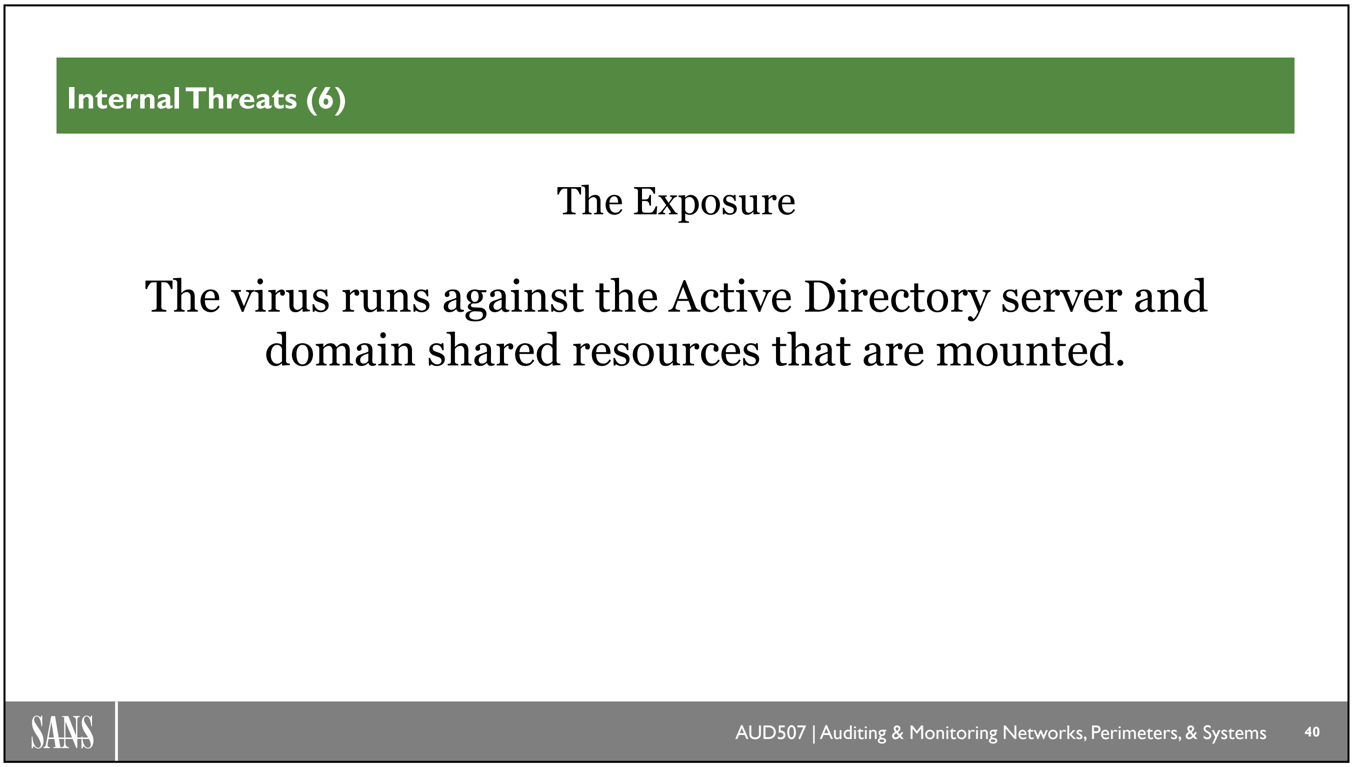

Internal Threats (6)

The Exposure

The virus runs against the Active Directory server and domain shared resources that are mounted.

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 40

In terms of loss or exposure, this organization fared quite well. It got off almost free and clear. Although the manager did have quite a few drives mapped on his system, the organization experienced little actual loss.

The malware in question can search local drives as well as network drives to identify certain files that it overwrites or moves. During this particular incident, the malware deleted many gigabytes of data from the corporate servers. How is it there was no actual loss, you ask? They were all MP3 files that were residing on the server in violation of other policy directives.

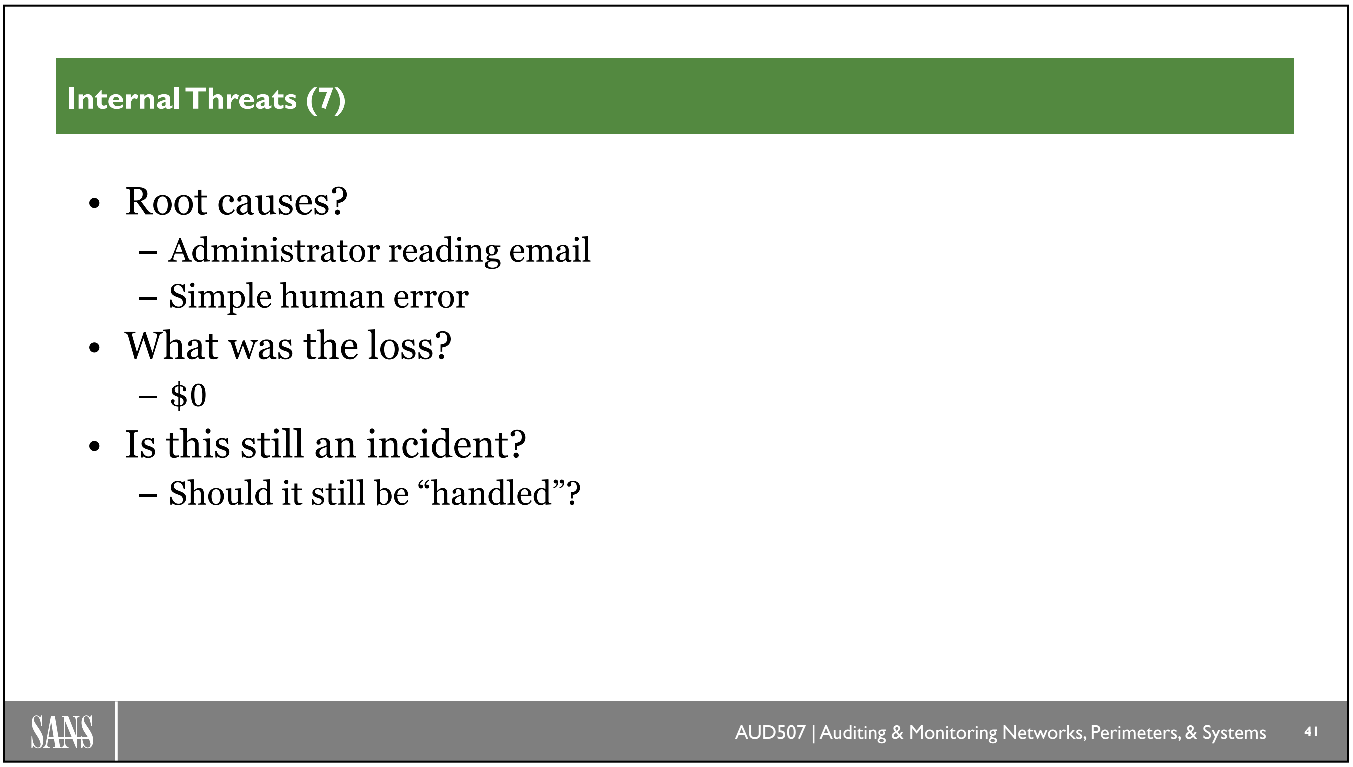

Internal Threats (7)

- Root causes?
  - Administrator reading email
  - Simple human error
- What was the loss?
  - $0
- Is this still an incident?
  - Should it still be “handled”?

Thinking about this example that has been laid out, we see that the root causes are clear. You may be thinking that you’ve seen similar examples. Because the administrator can do most tasks and it’s “easy” to use the administrative account, a lot of bad things happen as a result. Not only does one generally lack accountability when administrative accounts are used, but as in this case, there’s huge potential for something to go terribly wrong.

In this particular case, though, there was no actual or measurable loss. The question, then, is whether this still counts as an incident. The answer is absolutely “yes.” What about the final question on the slide, though? Should the incident-handling team still get spun up, run through their tasks, and generate a report? Especially because there was no actual loss, it can be tempting to answer “no.” Think about this for a moment, though. The incident has revealed that we are either lacking controls or have controls that failed to meet objectives. Incident handling sits in the layers of our Information Security program in a position that allows us to learn about precisely these issues. What this means is that we *must* handle this as an incident, even if there did not seem to be any real impact. If we fail to do so, it is difficult to learn properly from this experience!

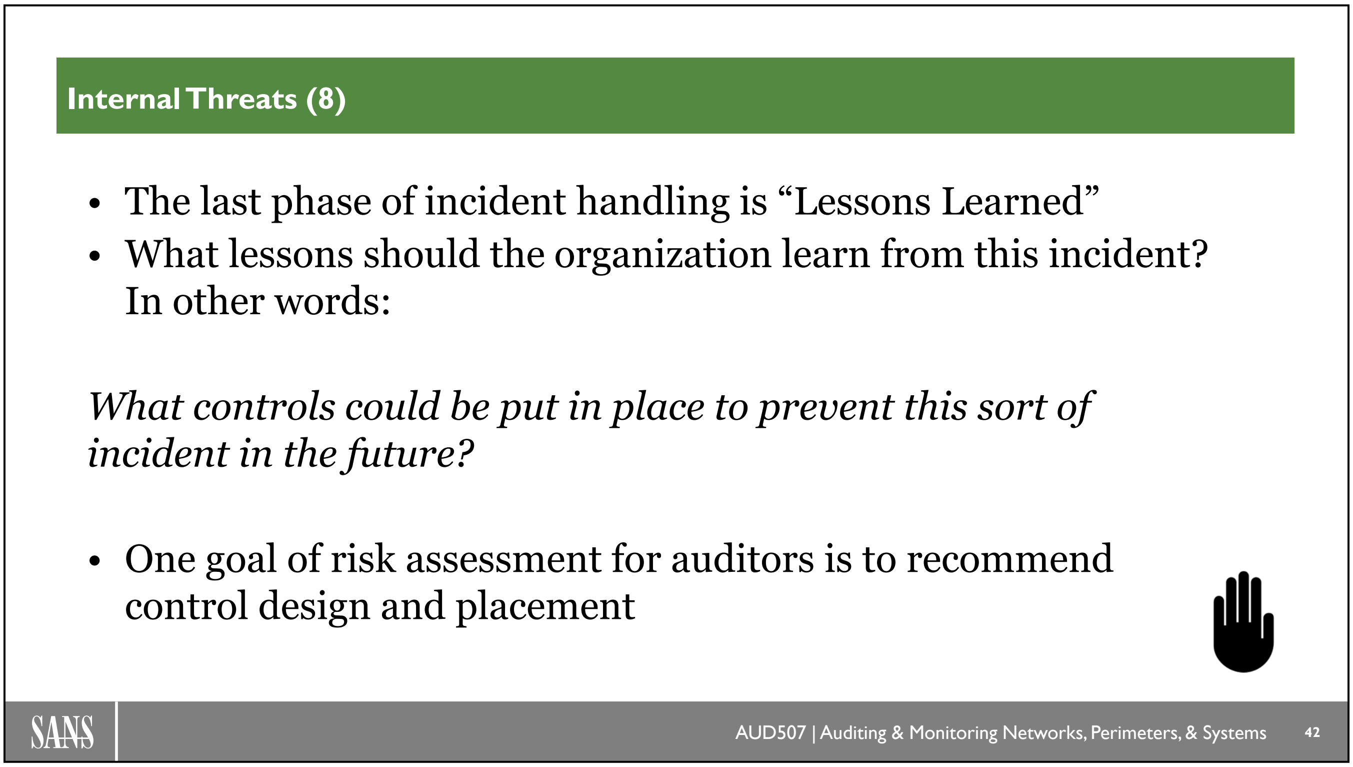

Internal Threats (8)

- The last phase of incident handling is “Lessons Learned”
- What lessons should the organization learn from this incident? In other words:

What controls could be put in place to prevent this sort of incident in the future?

- One goal of risk assessment for auditors is to recommend control design and placement

When SANS teaches classes on incident handling, we use a six-step process which ends with a phase called “Lessons Learned.” In this phase, the organization attempts to discern the root causes of the incident and design new or enhanced controls to try to prevent a similar incident in the future.

When auditors do risk assessments, it is usually for one of two reasons:

1. To enhance the possibility that the audit will be useful for the organization.

2. To determine where recommendations should be made for implementing new or redesigned controls to better mitigate the risk to the organization.

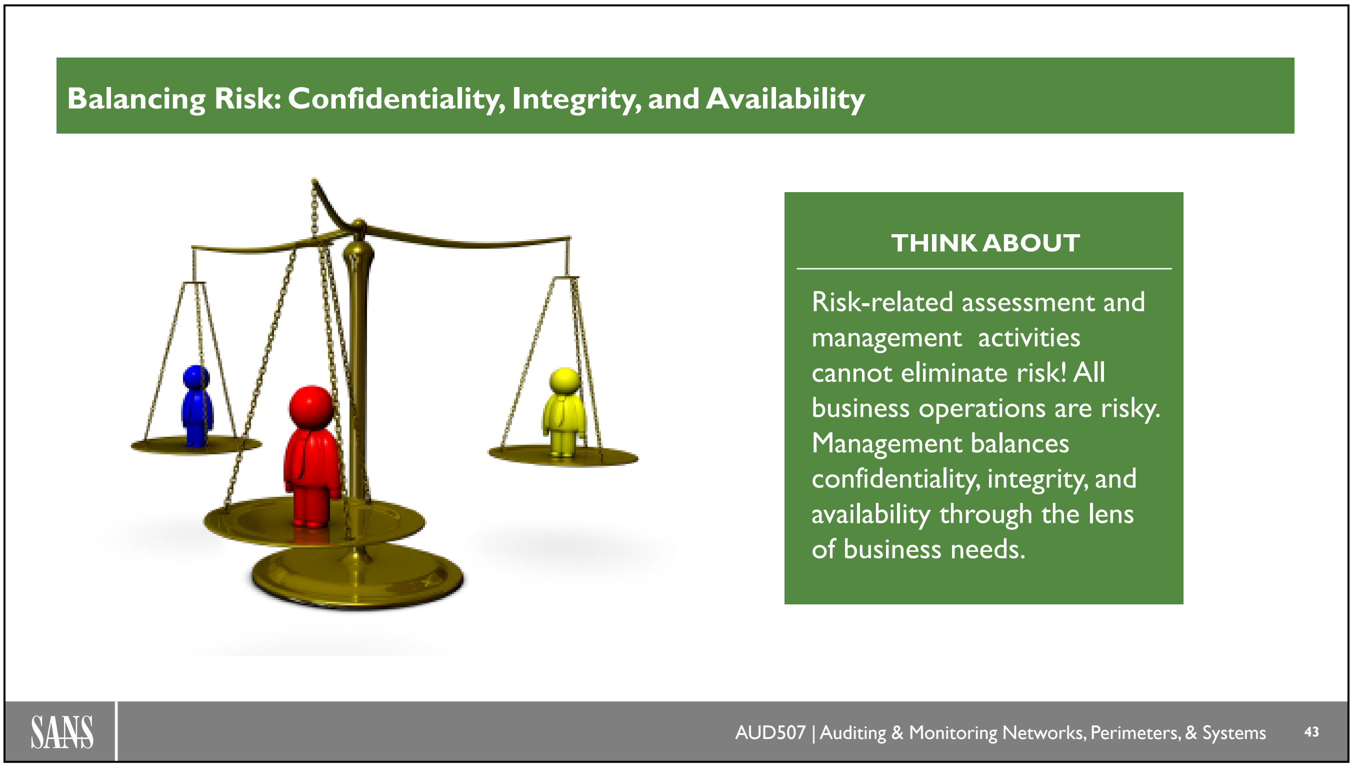

Balancing Risk: Confidentiality, Integrity, and Availability

THINK ABOUT

Risk-related assessment and management activities cannot eliminate risk! All business operations are risky. Management balances confidentiality, integrity, and availability through the lens of business needs.

SANS
AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems
43

Keep in mind that when you're performing risk assessments and managing risk, it is always a balancing act. The three primary components of Information Security are confidentiality, integrity, and availability. While it might be our natural inclination to focus on the confidentiality of the data in a system, business and mission requirements may dictate that availability of the data is far more important.

Consider this example: It might seem to you and me that the confidentiality of banking information is the most important consideration. After all, how much money a person has seems to be extremely private. Imagine, however, that your bank experienced an issue of some kind, and your money became unavailable. How critical is it that your money is available when you need it? Would you rather that your balance remains confidential, or would you prefer that you can access your funds when you need to? How important is it that the bank maintains the integrity of your account by properly applying all the deposits you make to the correct account?

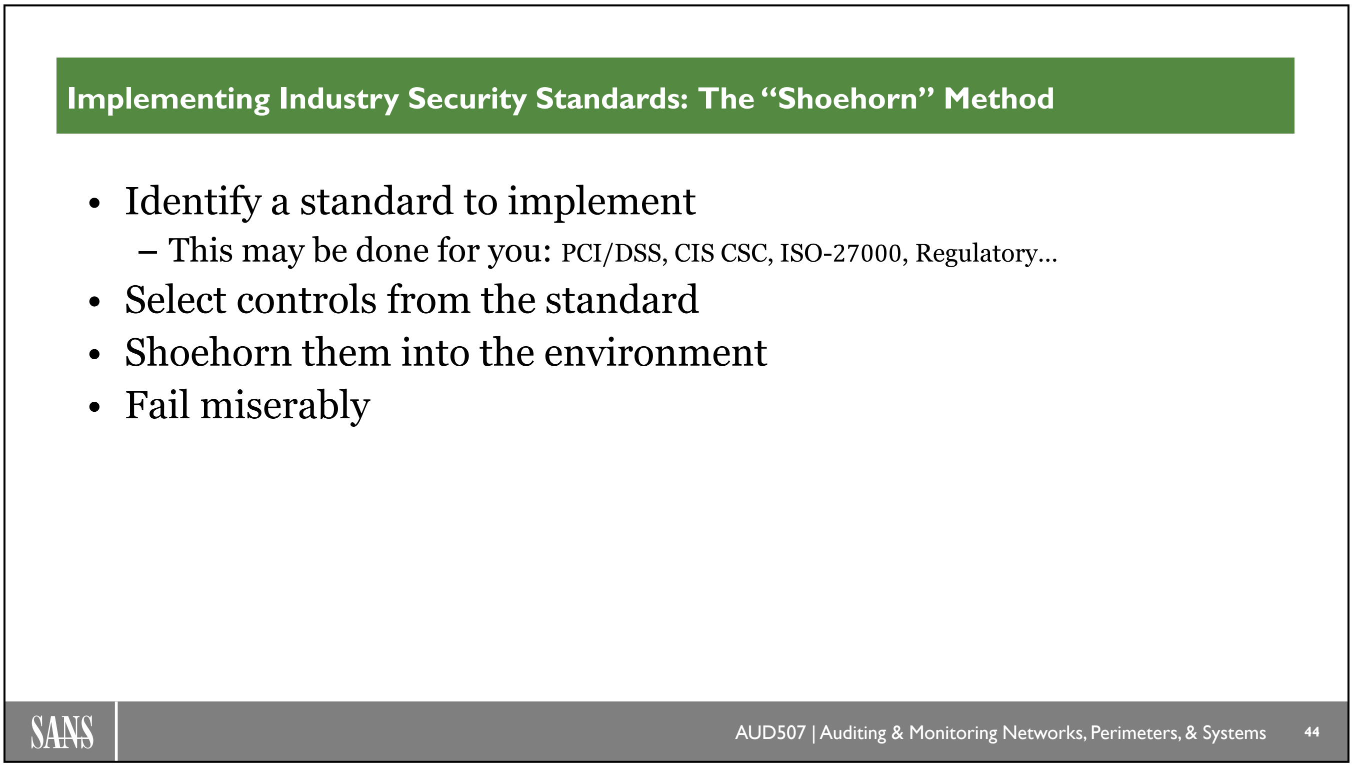

Implementing Industry Security Standards: The "Shoehorn" Method

- Identify a standard to implement
  - This may be done for you: PCI/DSS, CIS CSC, ISO-27000, Regulatory…
- Select controls from the standard
- Shoehorn them into the environment
- Fail miserably

We frequently see organizations that have failed in their implementation of a new industry or regulatory standard. The controls from the standard are implemented poorly, placed incorrectly, or otherwise ineffective in mitigating the very risks they were designed to address.

Here's a post-mortem of the sort of process we have seen fail in the past:

1. Identify the standard with which you are required to or want to be compliant. It is possible that this has already been done for you. For instance, if you work for an organization that accepts and processes credit cards, you are required by the payment card industry to adhere to the PCI Data Security Standard. If you work at a federal organization using information systems, you are required by law to adhere to specific federal requirements. Perhaps, in an effort to reduce risk, your management team has chosen to require adherence to the Center for Internet Security Critical Security Controls. Whatever the case, your standard is selected.
2. This where things often start to go wrong. The organization appoints a committee to read the standard and begin applying the controls from the standard into the organization. If the committee fails to plan and design for the proper placement and implementation of the controls, things may not go well.
3. Often, it is difficult to decide exactly how to implement the controls, so the committee begins to haphazardly apply them anywhere just so they can say they are implemented. Without a good understanding of the environment, the controls are applied in the wrong places or using the wrong techniques.
4. Eventually, something bad happens that should have been prevented or corrected by one of the controls. This could result in a damaging incident, or maybe even a regulatory violation that harms the organization.

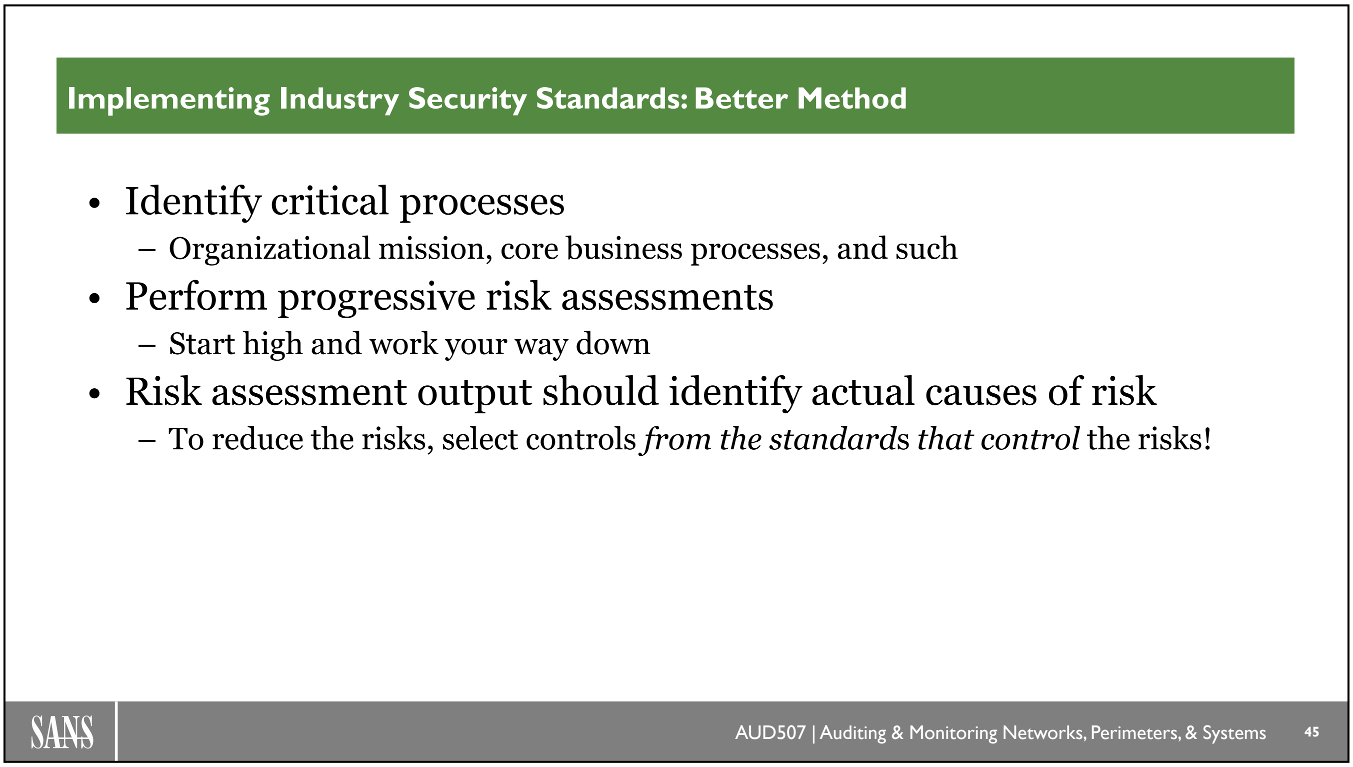

Implementing Industry Security Standards: Better Method

- Identify critical processes
  - Organizational mission, core business processes, and such
- Perform progressive risk assessments
  - Start high and work your way down
- Risk assessment output should identify actual causes of risk
  - To reduce the risks, select controls *from the standards that control* the risks!

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 45

As before, the first step is to identify the standard we wish to implement. Then, perform these steps to ensure better results:

1. Form a high-level steering committee or team, appointed and empowered by senior management, to analyze your enterprise to identify the critical processes of your business. These are the processes that help you to achieve your organizational mission and are commonly your core business processes.

2. Perform progressive risk assessments for confidentiality, integrity, and availability issues related to each of these processes. We say "progressive" because you typically start with an assessment of the entire process at a high level and progressively drill into the process as time goes on.

3. The risk assessments that you perform should identify risks that your business processes face. Knowing what these risks are, you then use the standard that has been selected to identify controls that control the risks!

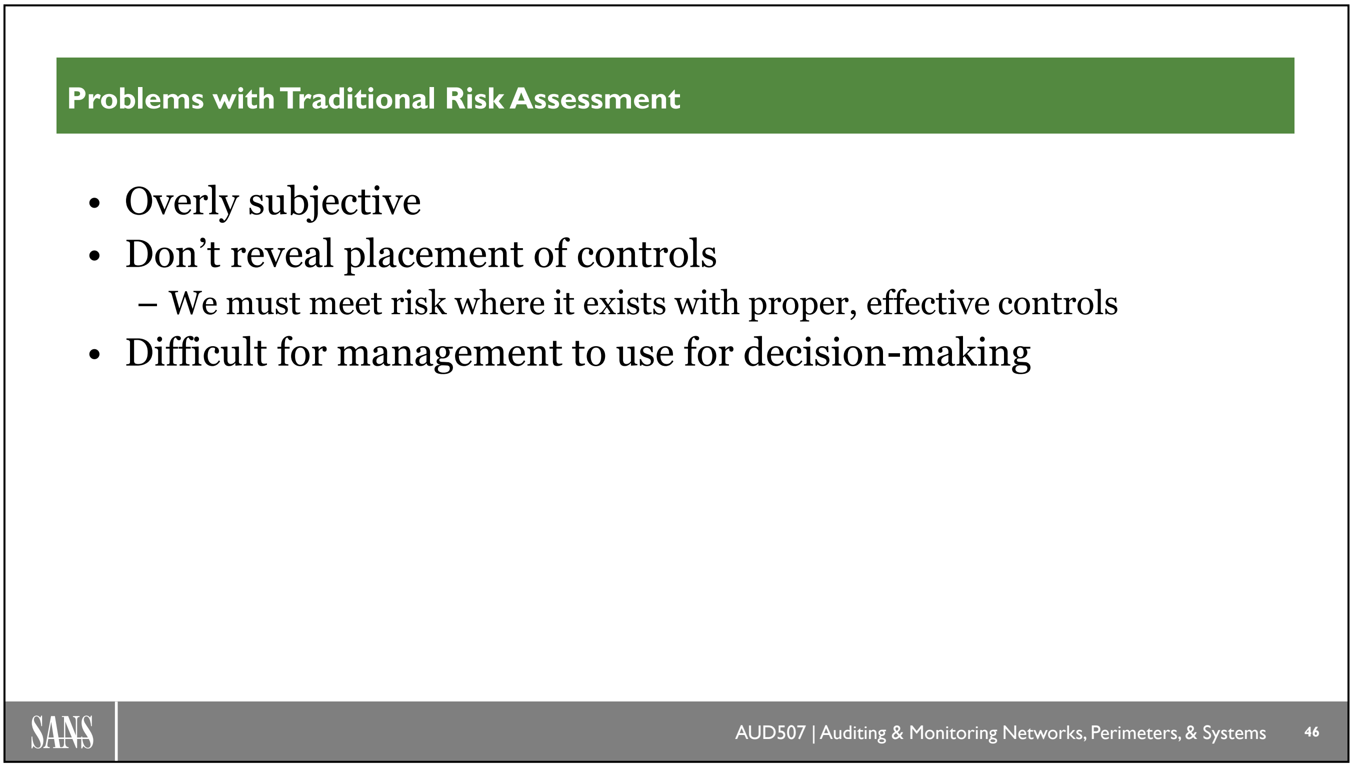

Problems with Traditional Risk Assessment

- Overly subjective
- Don't reveal placement of controls
  - We must meet risk where it exists with proper, effective controls
- Difficult for management to use for decision-making

Many of the risk assessment methods used in information technology share some common flaws. The first, and most important, is that the results are often overly subjective. Sure, we use numerical values and formulas to report on risk, but those numbers are often made up to match our qualitative and subjective assessment of the risk to the organization.

The second big problem is that knowing about a risk does not inform me of how to mitigate that risk with controls. We need assessment methods that allow us to determine how controls should operate and where they should be placed.

Finally, many of our risk assessment methods give very little actionable information to management. How can management be expected to make good cost/benefit decisions when they know neither the relative cost nor the relative benefit of the controls recommended to them?

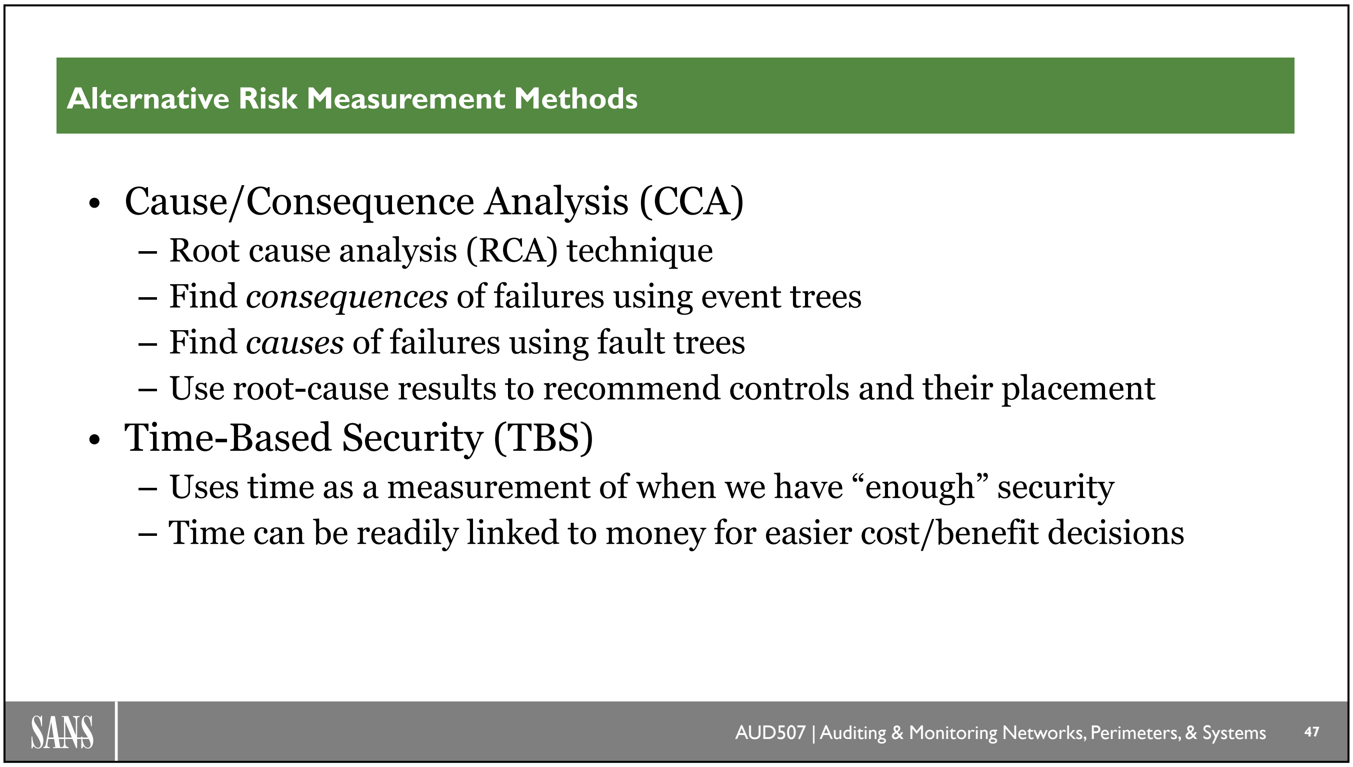

Alternative Risk Measurement Methods

- Cause/Consequence Analysis (CCA)
  - Root cause analysis (RCA) technique
  - Find *consequences* of failures using event trees
  - Find *causes* of failures using fault trees
  - Use root-cause results to recommend controls and their placement
- Time-Based Security (TBS)
  - Uses time as a measurement of when we have "enough" security
  - Time can be readily linked to money for easier cost/benefit decisions

To address the problems with traditional risk assessment techniques, we will discuss two alternative methods for understanding risk and recommending control placement to mitigate it.

The first is called Cause/Consequence Analysis (CCA). This is a form of root cause analysis which can be very helpful in designing controls for the organization.

The second is a concept for understanding risk known as Time Based Security (TBS). TBS gives the auditor a framework to use to understand and frame discussion of risk and control.

We'll begin with a conversation about CCA.

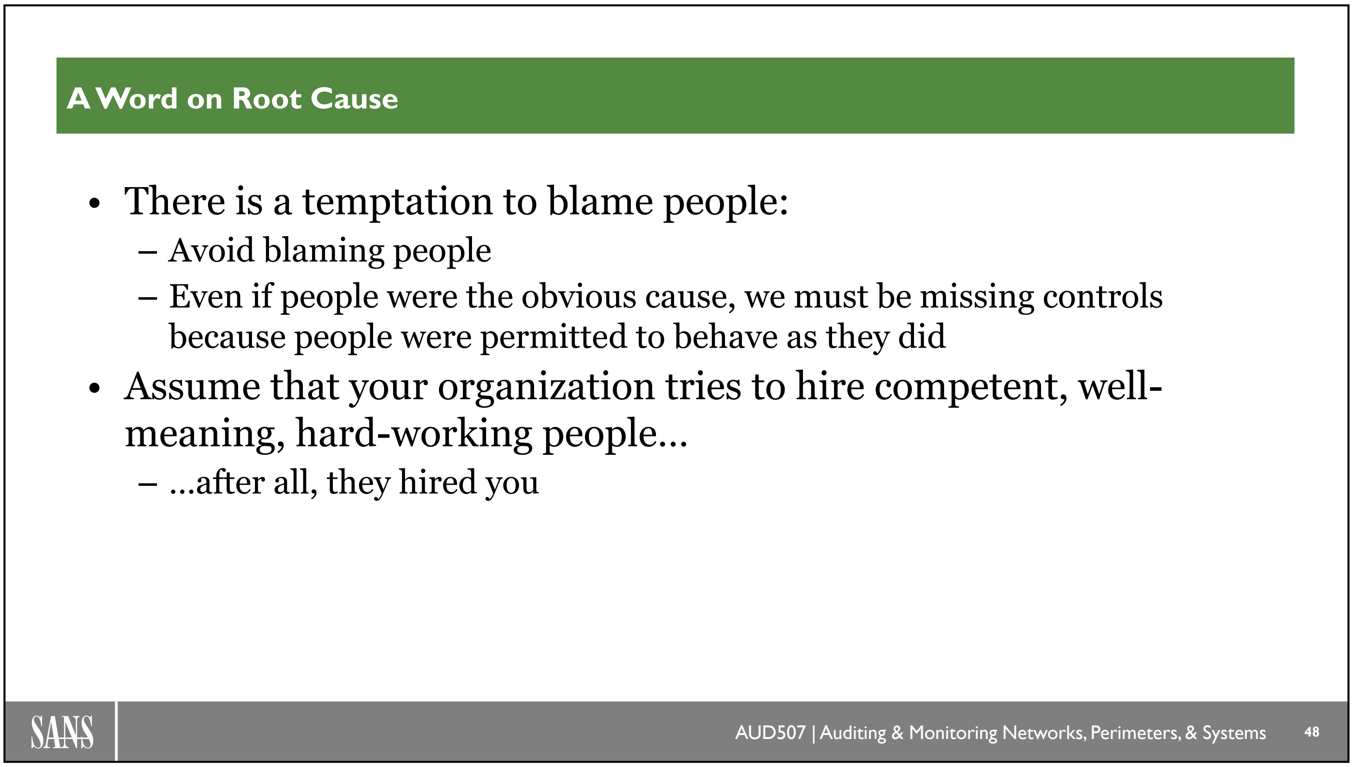

A Word on Root Cause

- There is a temptation to blame people:
  - Avoid blaming people
  - Even if people were the obvious cause, we must be missing controls because people were permitted to behave as they did
- Assume that your organization tries to hire competent, well-meaning, hard-working people…
  - …after all, they hired you

Before we continue with our discussion, let me state something up front: People are almost never the root cause of your problems.

There is a real temptation to blame individuals. Although this does not usually happen during the risk assessment phase while creating controls, when performing the risk assessment as a result of an incident that has occurred, the temptation is to blame people. Some of this is cultural. Western cultures tend to feel better about problems if they can blame those problems on someone. The reality is, however, if a person were the problem, then the problem would never recur after replacing the person. What we find in practice is that the problems do recur, which, anecdotally proves that the person was not the problem.

The real answer is that even if a person were the problem in some way, something is lacking in our controls that allows that person to behave in the way that he did. If there is a failure, it means that we are missing preventative controls. If no control detected the failure, we are also missing detective controls. Finally, depending on how long it took to discover the problem, the failure indicates that we are missing reactive controls. In none of these cases is the person the problem.

When dealing with the question of fault, here’s a great way to explain and think of it. Simply assume that your organization tries its best to hire competent, well-meaning, hard-working people. When you say this out loud, don’t be surprised to hear some snorts and giggles. You might be tempted to snort and giggle yourself until we consider the second part of this: It hired you, didn’t it?

So, if you are an example of the kind of people your company hires, then what is the reason that risk persists? Is it truly the people, or is it more likely a lack of controls (or simply misapplied controls)?

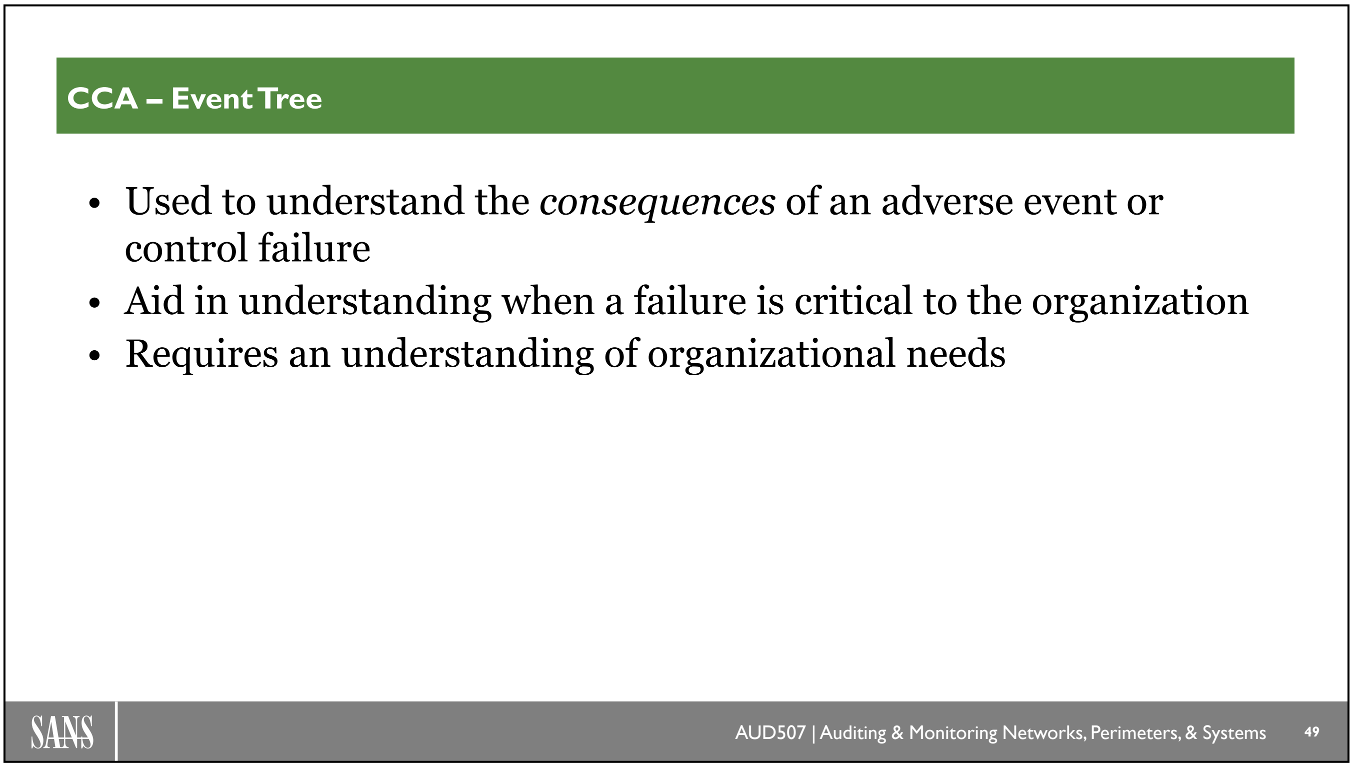

CCA – Event Tree

- Used to understand the *consequences* of an adverse event or control failure
- Aid in understanding when a failure is critical to the organization
- Requires an understanding of organizational needs

Event trees are used to identify the consequences of a potential adverse event or the failure of some control. The objective of the event tree is to identify which controls, if they should fail, would create a critical risk to the organization and its objectives. We will often require an understanding of organizational needs and operating environment to do a good job with developing the event tree.

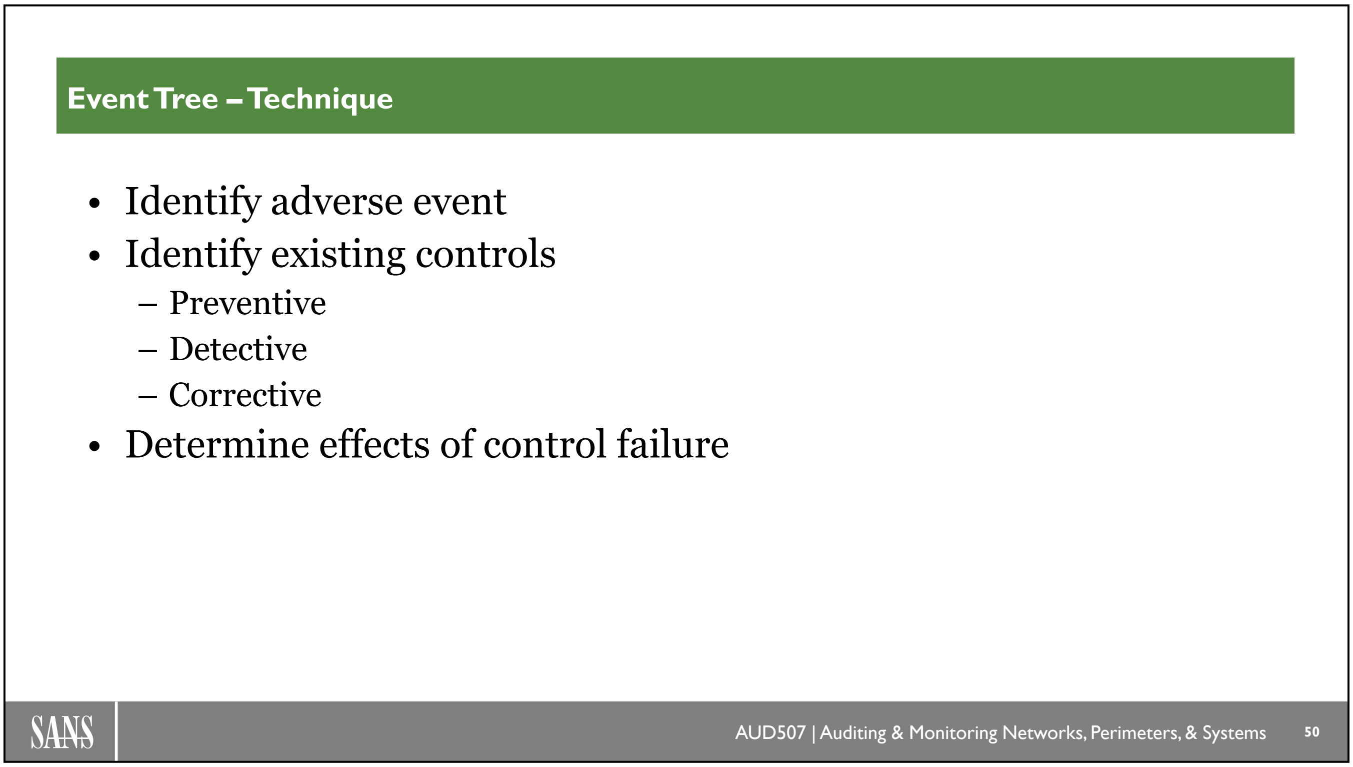

Event Tree – Technique

- Identify adverse event
- Identify existing controls
  - Preventive
  - Detective
  - Corrective
- Determine effects of control failure

An *event tree* is a method of analyzing a system's detective and reactive controls. In addition to forcing us to inspect these two control types, the other output of the assessment method is the identification of critical failures.

Remember that if we are to apply controls correctly and efficiently, we first need to identify where the critical failures to the system or process can occur. For this reason, event trees are ideal. Let's see how we create one.

Generally, controls are separated into three categories: *Protective* (or *preventative*), *detective* (sometimes considered to be *audit* controls), and *reactive* (or *corrective*). For the event tree analysis, we will often not examine protective controls. Where protective controls exist, they frequently come in a form that includes a detective or reactive component. For example, in your data center, you likely have an uninterruptible power supply. The Uninterruptible Power Supply (UPS) acts as a protective control, but the way that it achieves this protection is by detecting and reacting.

Since we are interested in examining what happens when controls fail, purely protective controls are not especially useful for this exercise. Our tendency can be to assume that they cannot fail.

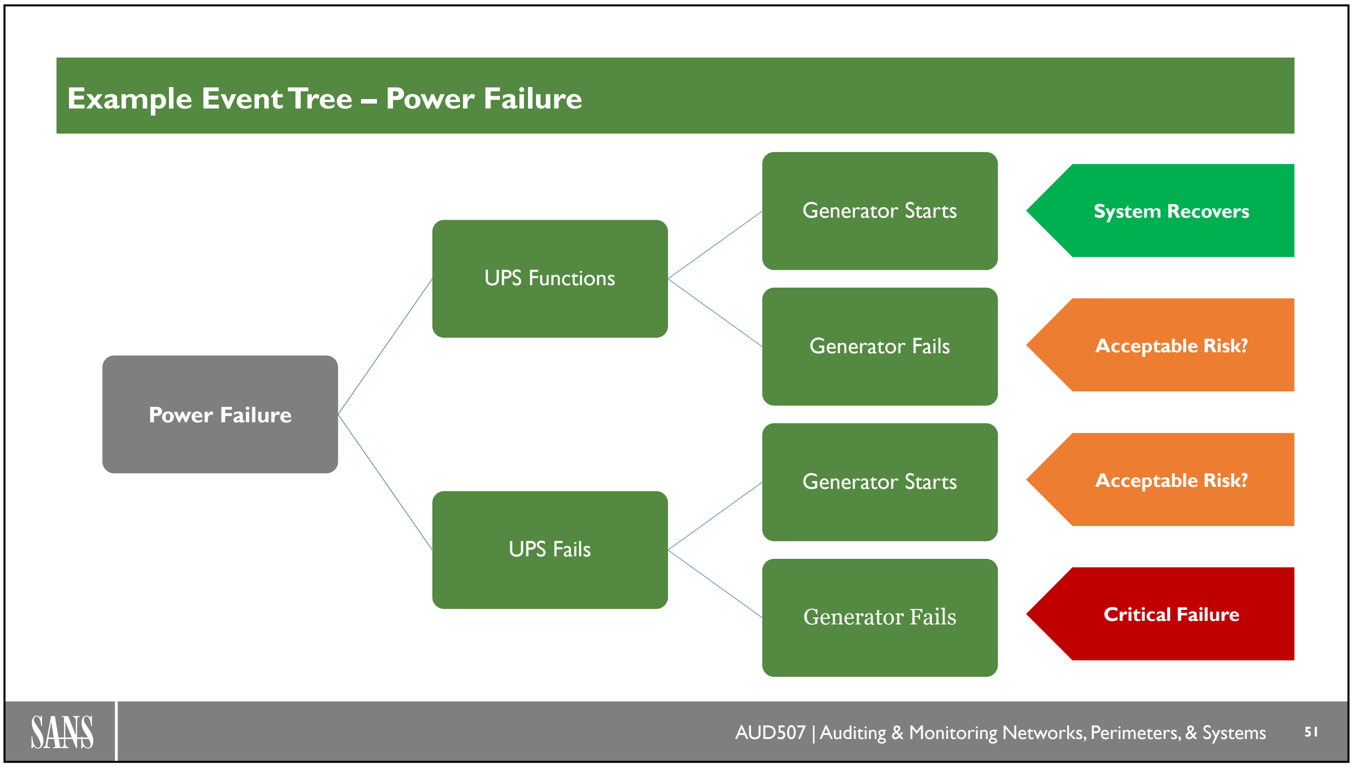

Example Event Tree – Power Failure

Power Failure

UPS Functions

UPS Fails

Generator Starts

Generator Fails

Generator Starts

Generator Fails

System Recovers

Acceptable Risk?

Acceptable Risk?

Critical Failure

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 51

Here is an example of an event tree. In this case, we examine the recovery controls surrounding the potential event of a power failure. Rather than getting into whether this is a failure on the part of the power company or a spilled cup of coffee, the risk remains the same: Availability is affected. For the sake of simplicity, we have touched only on two detection and recovery controls in the diagram. For two of the cases, the outcomes are obvious: If the Uninterruptible Power Supply (UPS) functions correctly and the generator starts, we determine that the system has recovered at an acceptable level, whereas if the UPS fails to function and the generator fails to start, we have an obvious critical failure. Most likely, the decisions on what is critical and what is not have already been worked out in a business impact analysis during DR/BCP planning.

What is not entirely clear (and would depend on organizational requirements) is whether the other two events are critical failures or system recoveries. For instance, it may be that although power failures are not desirable, it is okay if the generator does not start as long as the UPS works long enough for the systems to be brought down successfully. In this case, availability is likely not as critical as data integrity.

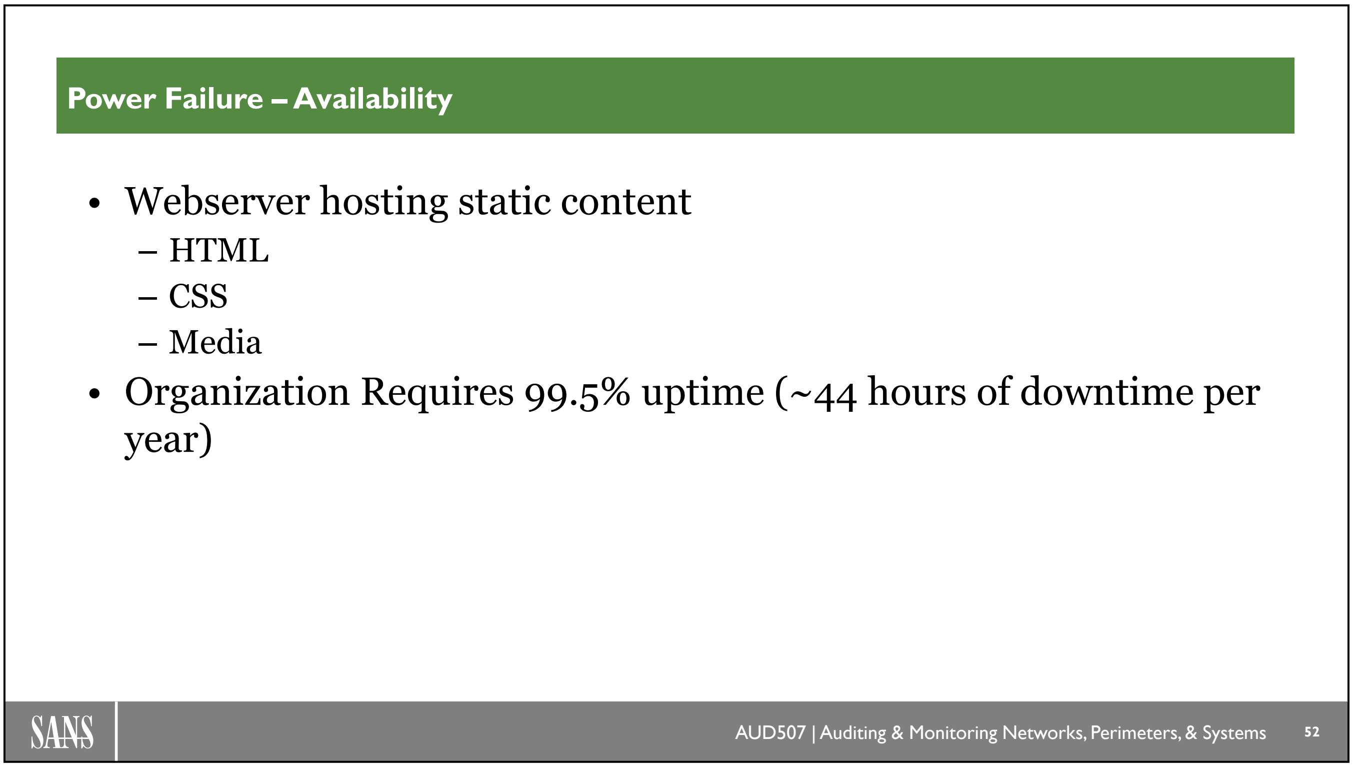

Power Failure – Availability

- Webserver hosting static content
  - HTML
  - CSS
  - Media
- Organization Requires 99.5% uptime (~44 hours of downtime per year)

In this example, we have chosen to examine a specific system: A static web server.

For our sample organization, we will imagine that we have had a conversation with management where we explore which aspects of confidentiality, integrity, and availability matter most. This decision is always a balancing act; most of the time, all three will matter, but at least one will matter *more* than the others.

For instance, in this case, our management interviews informed us that the server being considered is used to host static web content, like images, videos, and cascading style sheets (CSS). For this server, management's biggest concern is that the web server remains reachable a large percentage of the time.

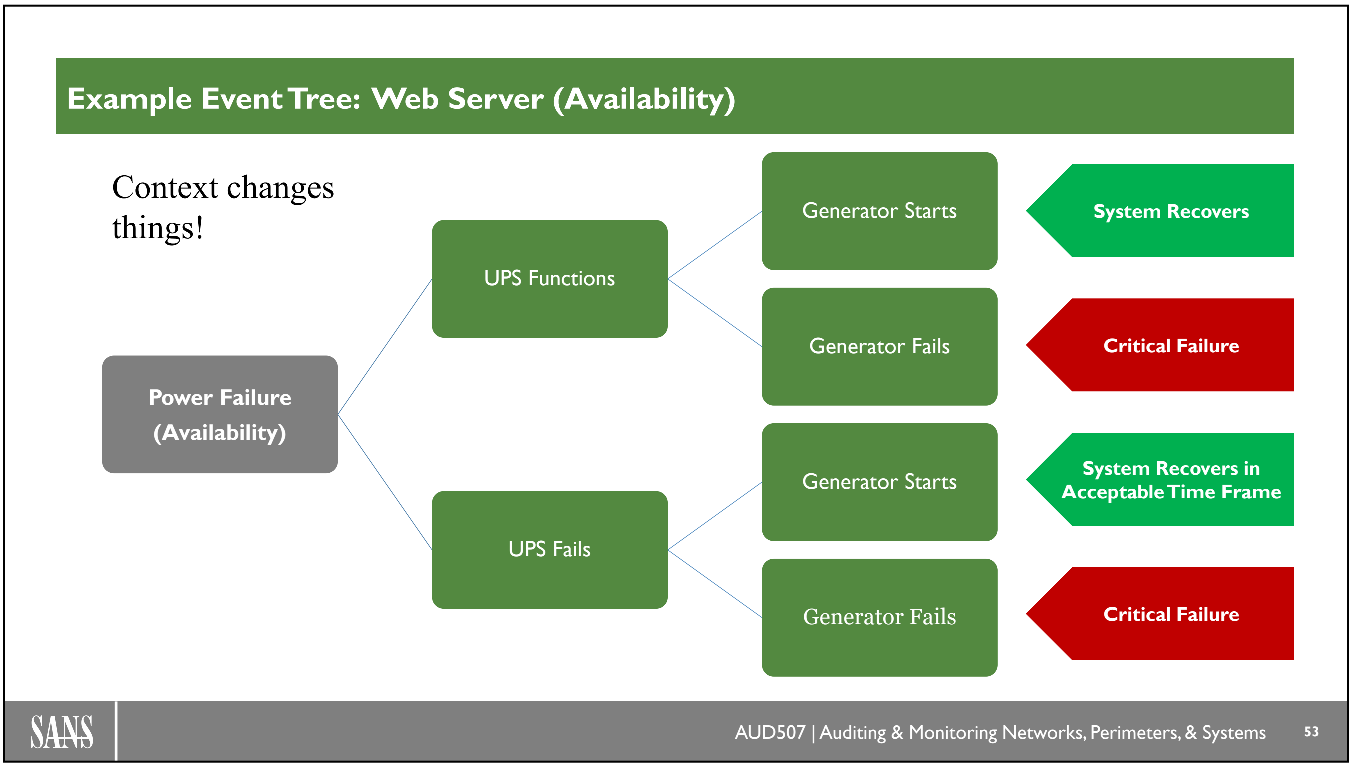

Example Event Tree: Web Server (Availability)

Context changes things!

Power Failure (Availability)

UPS Functions

Generator Starts

System Recovers

Generator Fails

Critical Failure

UPS Fails

Generator Starts

System Recovers in Acceptable Time Frame

Generator Fails

Critical Failure

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 53

With that in mind, we examine the failure cases for the web server. We know that as long as the UPS works and the generator starts, everything is good. It's also pretty clear that if the UPS and generator both fail during an extended power outage, a critical failure has occurred. What is less clear is what happens if the UPS fails, but the generator works.

After discussing this with management and the technical teams, we conclude that if the UPS fails but the generator functions, the web server will usually be able to recover within about three minutes, returning to service. Based on this information, management decides that this is an acceptable risk for this system.

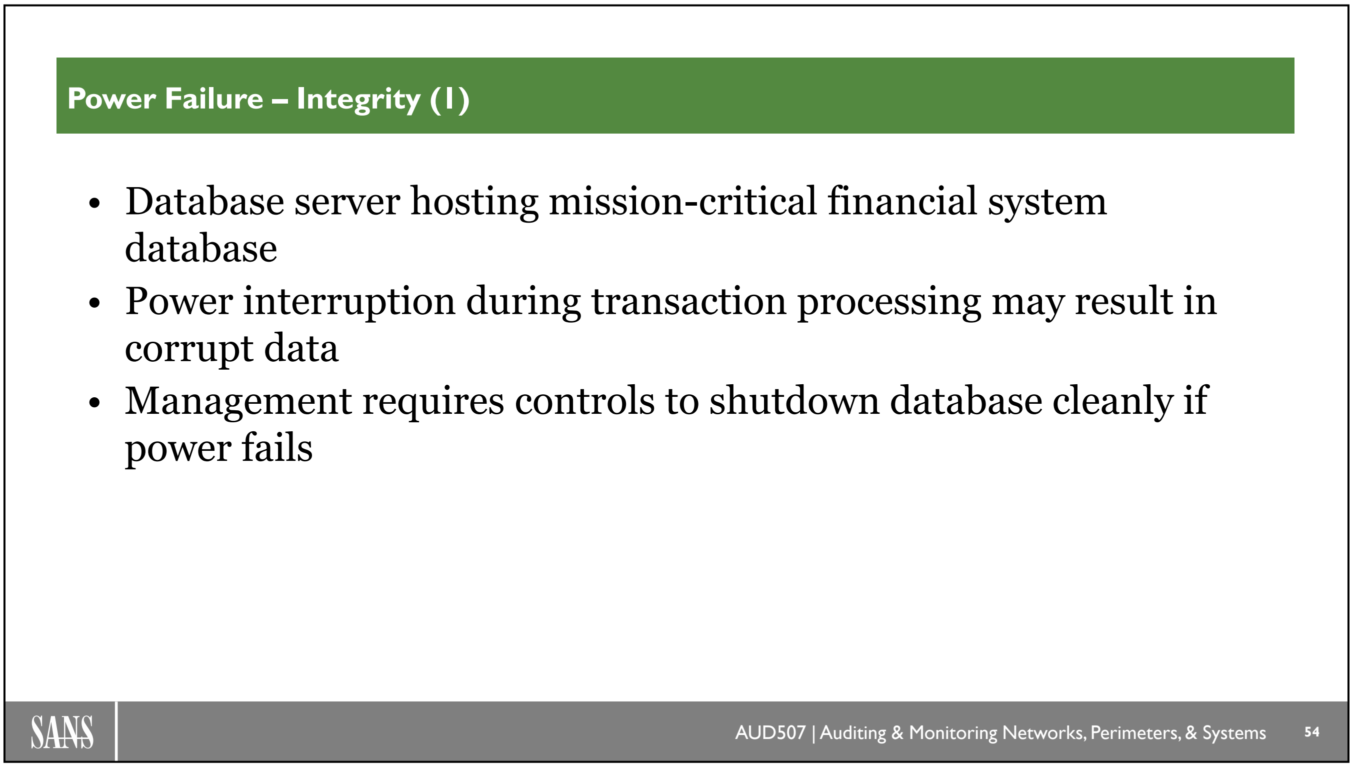

Power Failure – Integrity (1)

- Database server hosting mission-critical financial system database
- Power interruption during transaction processing may result in corrupt data
- Management requires controls to shutdown database cleanly if power fails

Event trees can serve to show us something else, though. In this example, we are looking at a database server. Unlike with web servers, data integrity is usually the greatest concern for a database server. If a database system is abruptly powered off (rather than shut down cleanly), the database is very likely to become corrupted.

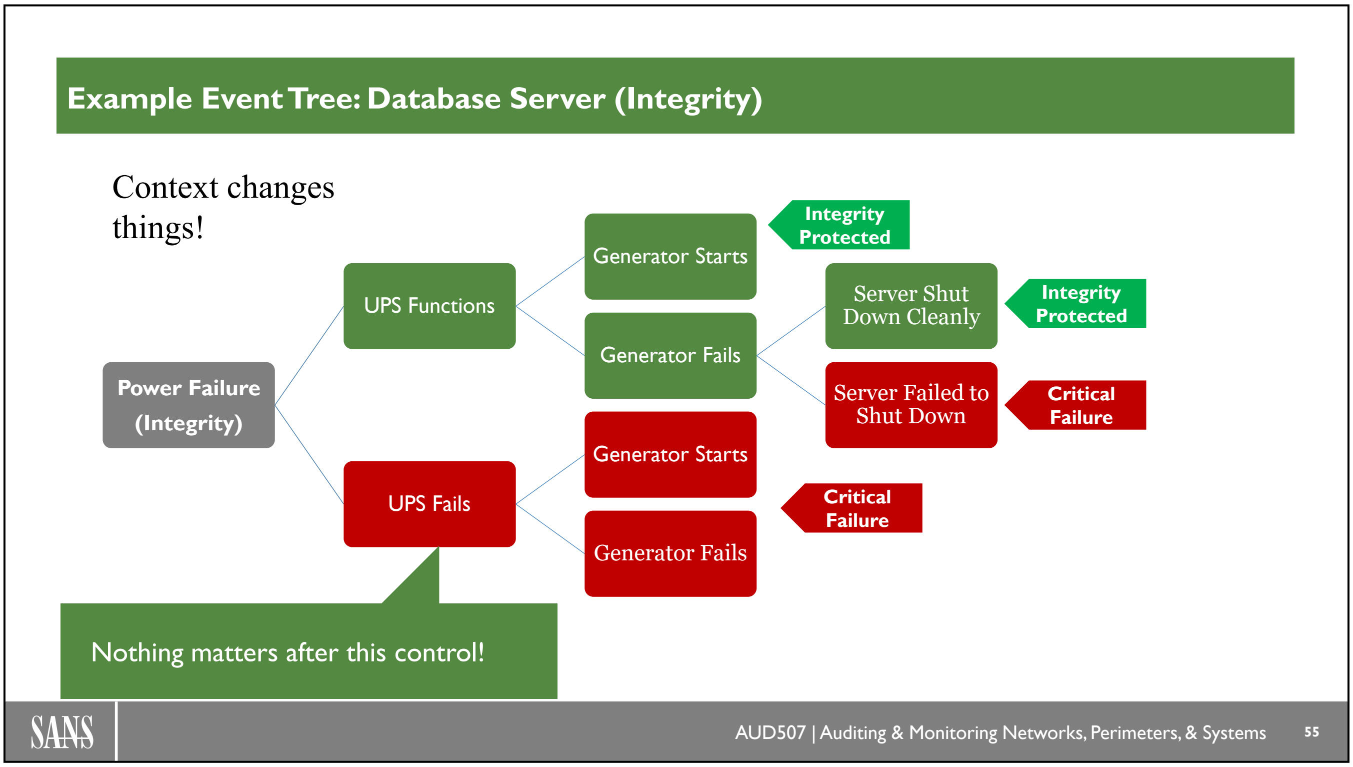

Example Event Tree: Database Server (Integrity)

Context changes things!

Power Failure (Integrity)

UPS Functions

Generator Starts

Integrity Protected

Generator Fails

Server Shut Down Cleanly

Integrity Protected

Server Failed to Shut Down

Critical Failure

UPS Fails

Generator Starts

Generator Fails

Critical Failure

Nothing matters after this control!

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 55

In our discussions with management and after explanations from the technical team, it emerges that once the UPS fails, no other control matters. While we have only illustrated two controls here, even if we had another 30 or 40 compensating controls, once the UPS has failed, nothing else matters.

In all three cases, we have been able to define failure modes that represent an unacceptable risk for the organization. Event trees are especially useful for this. We do not yet know what to do about this risk, though it is clear that additional controls are needed. We will come to that shortly.

How is this better than “Threat × Vulnerability?” This approach provides us with a concrete scenario, or a table-top exercise of our controls, that identifies what happens if the controls fail. Simply using Threat × Vulnerability might lead us to believe that because we have lots of controls to mitigate issues, everything must be OK.

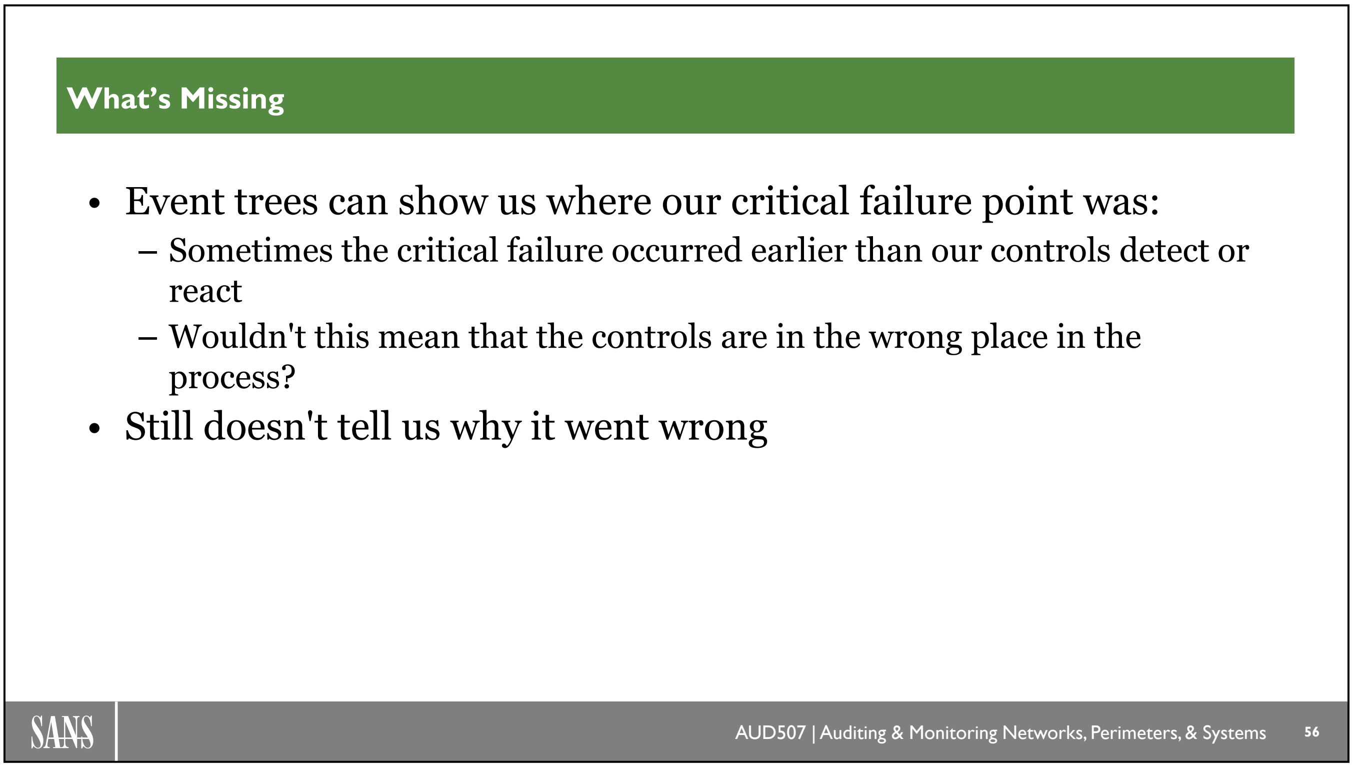

What's Missing

- Event trees can show us where our critical failure point was:
  - Sometimes the critical failure occurred earlier than our controls detect or react
  - Wouldn't this mean that the controls are in the wrong place in the process?
- Still doesn't tell us why it went wrong

The key to using event trees, as simple as they are, is seeing that they can sometimes tell us that the real failure occurred long before we were focusing on the problem. For example, if a nuclear reactor melts down, could it be that detecting a small drop in water pressure in a noncritical portion of the plant much earlier might have given us the information necessary to correct and prevent the meltdown? It's possible that by the time we detect it, we are already on the path to critical failure and can no longer recover.

Even so, this does not tell us what made the problem occur. For instance, how can we figure out that the problem was a leaky pipe in a bathroom? The next assessment explains this.

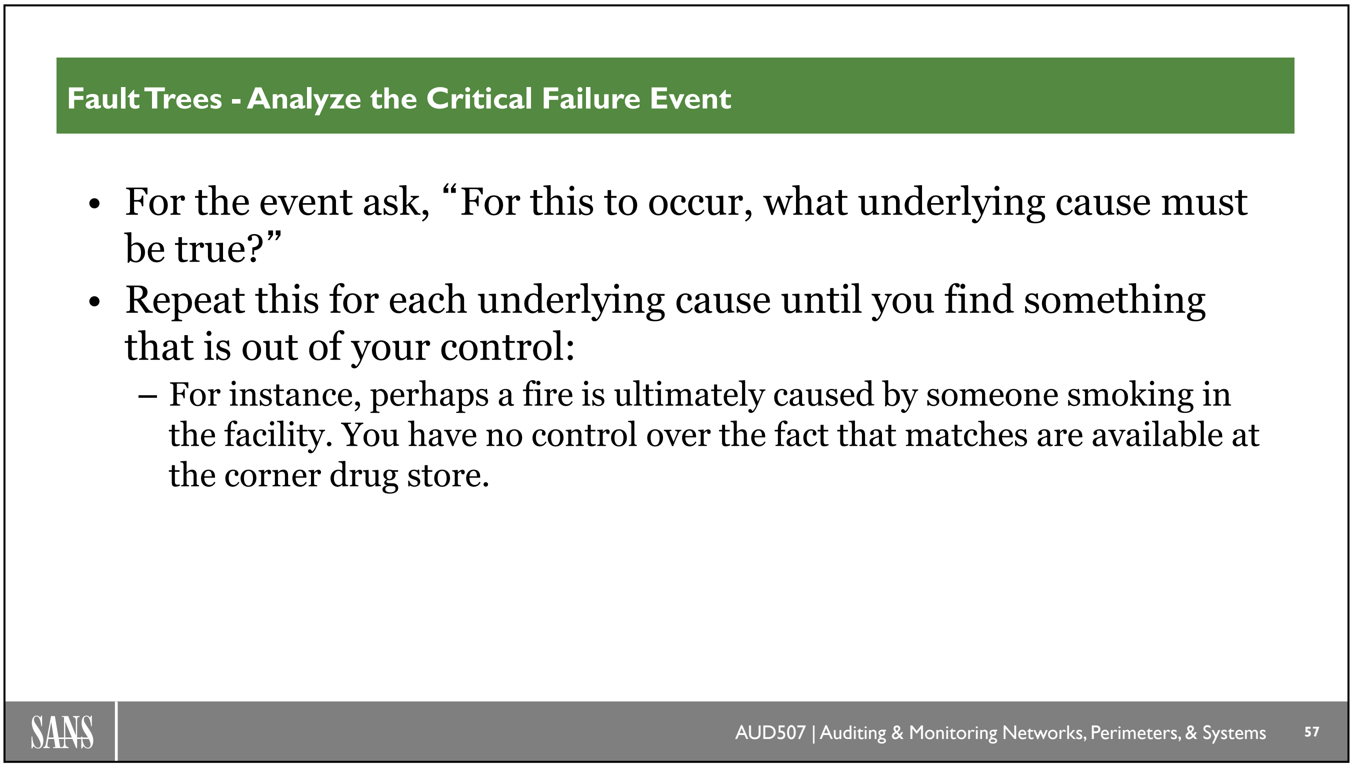

Fault Trees - Analyze the Critical Failure Event

- For the event ask, “For this to occur, what underlying cause must be true?”
- Repeat this for each underlying cause until you find something that is out of your control:
  - For instance, perhaps a fire is ultimately caused by someone smoking in the facility. You have no control over the fact that matches are available at the corner drug store.

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 57

The first step in performing this analysis is to identify a critical event that either has occurred (retrospective, damage, or incident analysis) or might occur (risk analysis). After identifying the critical failure, clearly ask, “For this event to have occurred, what underlying facts must be true?” For example, if a backup tape is blank, what must be true? Here are some possibilities: The backup was never run, the tape was mislabeled, the tape drive is defective, or the backup software is not functioning correctly.

After you find some underlying factors, repeat the question at that level. Actually, keep repeating this question until the underlying cause is out of your control or simply not important. For example, if there has been a fire in a facility, one of the underlying causes might be that someone was smoking. An underlying factor here would be that matches and cigarettes are for sale at the corner drug store, so it is not particularly useful to descend to that level; there’s nothing you can (easily) do to control the corner drug store.

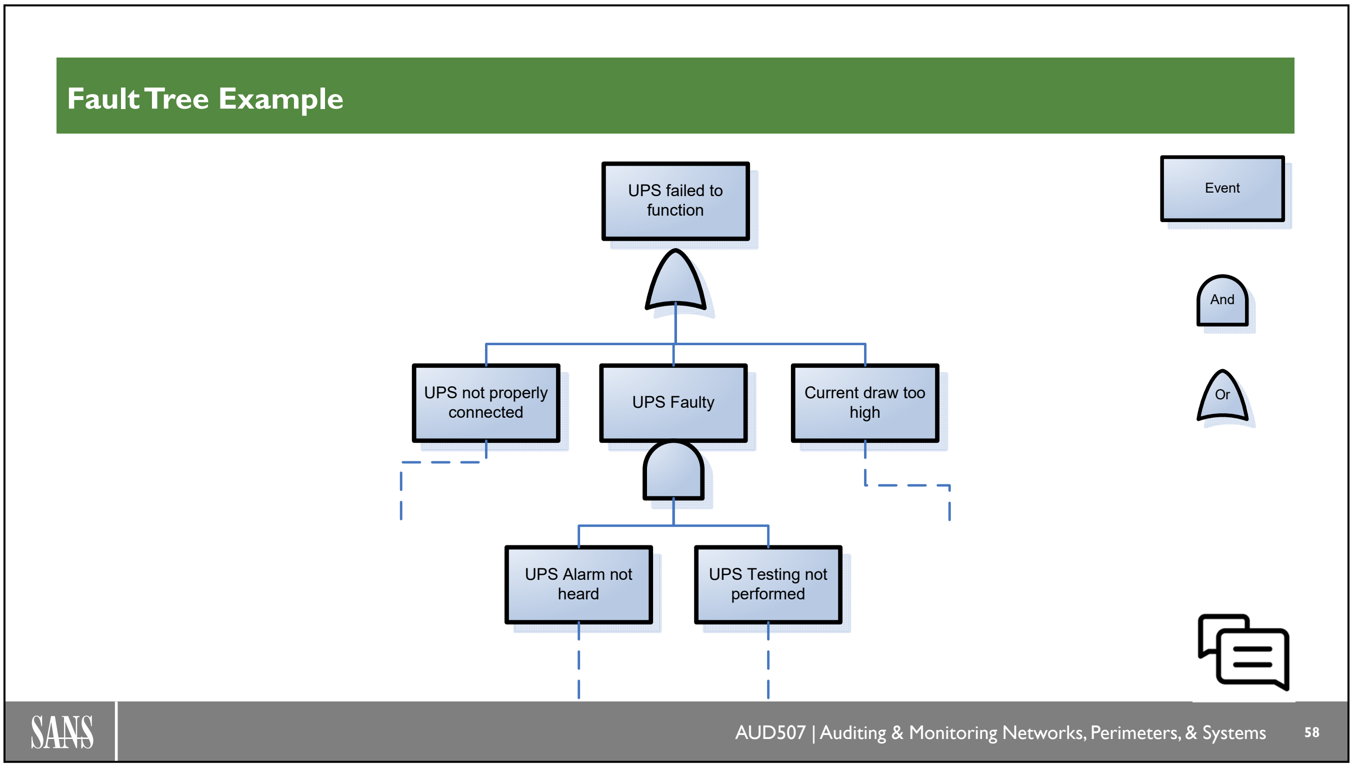

Fault Tree Example

UPS failed to function

UPS not properly connected

UPS Faulty

Current draw too high

UPS Alarm not heard

UPS Testing not performed

Event

And

Or

In this slide, you can see a simple example of a fault tree. To create one, start with a failure (in this case, a UPS failing to function) and then ask, “How could that happen?” Sometimes the causes will be unrelated events that could each occur independently. (For instance, the UPS is not properly connected, the UPS is faulty, or the current draw is too high.) Notice that we separate these with the word “Or.” To represent that, we use the symbol listed as “Or” at the top right of the slide.

In other cases, the problem could be the result of several faults that are dependent on one another. For instance, if the UPS is faulty, that means that both the UPS alarm was not heard, and the periodic UPS testing was not performed frequently enough to discover the faulty equipment. Here we use the word “And,” so we use the appropriate symbol from the upper right of the slide to indicate this.

So, what’s the big deal? The big deal is that we can use this technique to drill down to all the possible causes, each of which might seem trivial on its own but can contribute to a critical failure. Knowing this, we can introduce controls down toward the bottom of the fault tree that can control the risks at the top of the tree!

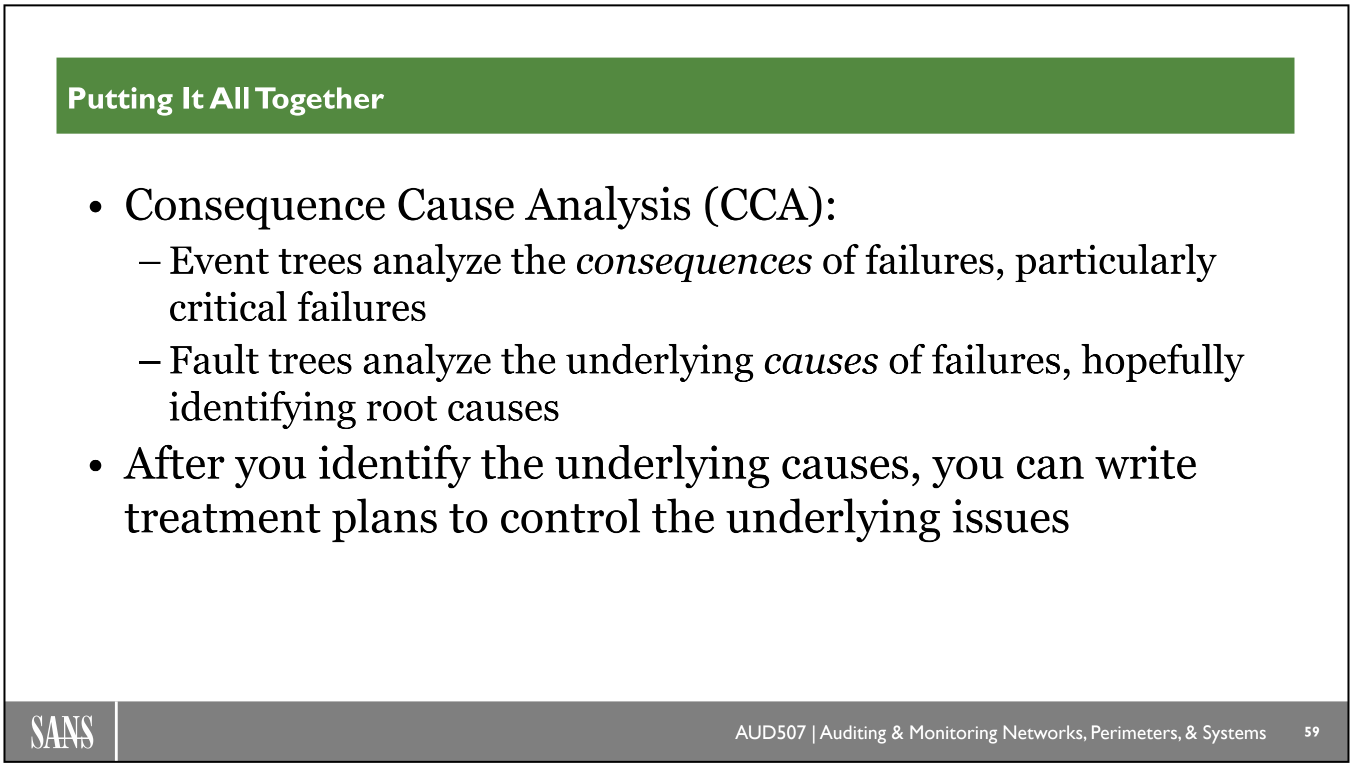

Putting It All Together

- Consequence Cause Analysis (CCA):
  - Event trees analyze the *consequences* of failures, particularly critical failures
  - Fault trees analyze the underlying *causes* of failures, hopefully identifying root causes
- After you identify the underlying causes, you can write treatment plans to control the underlying issues

At this point, we've looked at two separate risk assessment strategies. Likely, you can see where this is heading; each of the methods is useful in its own right but becomes extremely powerful when we join them together (like the Wonder Twins, if you have kids). More important, consider the relationship between the two systems. Event trees are extremely useful for identifying the consequences of failures, particularly in finding critical failures, whereas fault trees are especially good at identifying the underlying causes of the failures, allowing us to find and control the root causes.

When put together, these two strategies make up Consequence Cause Analysis (CCA). Using CCA, the final outcomes are the potential underlying (and uncontrolled) causes of critical failures for which we can now write treatment plans so that detective and corrective controls can be implemented. As auditors, it gives us the opportunity to find the root causes of issues and then to recommend practical fixes at the root of the problem.

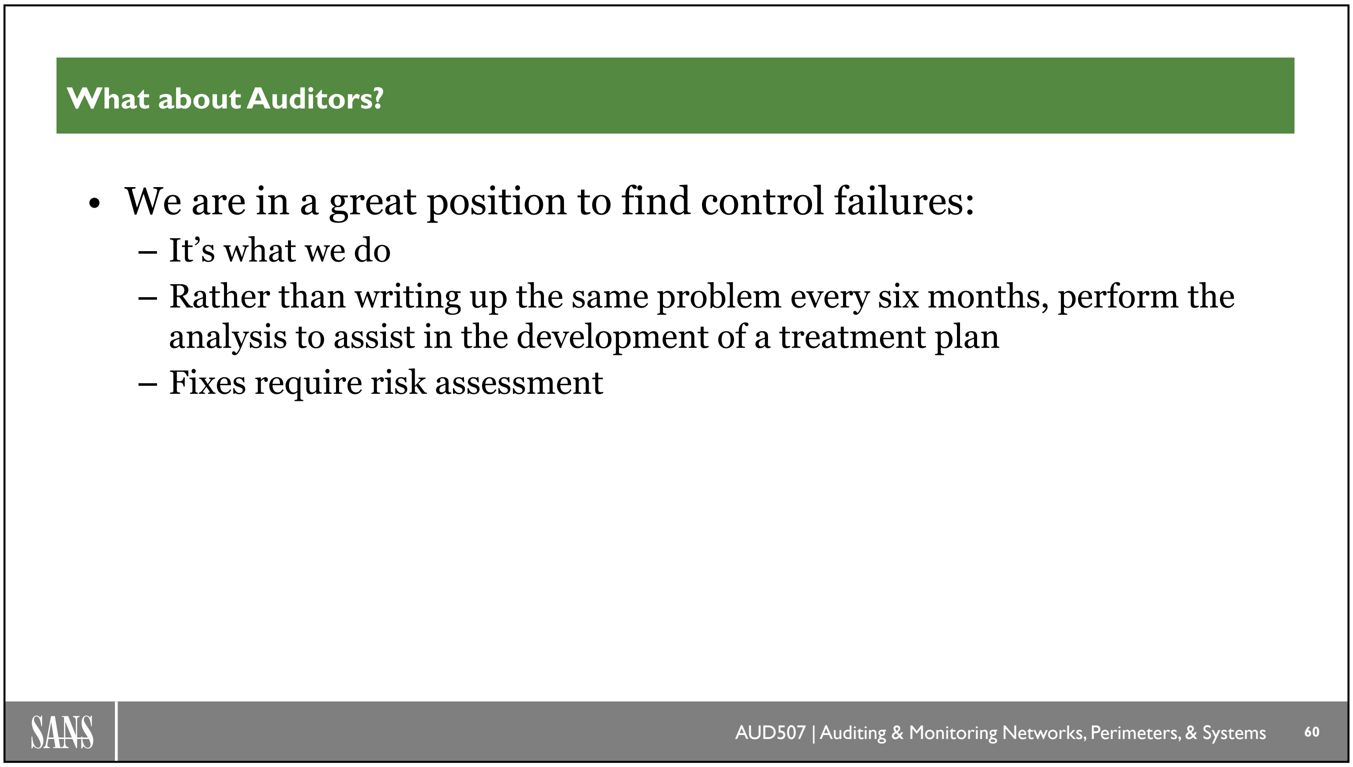

What about Auditors?

- We are in a great position to find control failures:
  - It’s what we do
  - Rather than writing up the same problem every six months, perform the analysis to assist in the development of a treatment plan
  - Fixes require risk assessment

SANS
AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems
60

Auditors are in a wonderful position to find potential (and actual) control failures. It’s what we do for a living! It also allows us to live more satisfying lives as auditors. Rather than writing up the same problems every six months, we can dig underneath, find the root cause, and pull that weed-like problem out of the organization by recommending a fix. The fix that we recommend can now be based on risk assessment, allowing us to treat the actual causes rather than the symptoms. In the end, this is what prevents recurrence.

What if the problem recurs even after applying new controls? Here’s a tip: Audit exceptions are like weeds. If you don’t pull out the root, the weed grows back. The same is true here. If the problem recurs, you might have performed a wonderful risk analysis, but you didn’t successfully uncover the root cause, so it’s time to dig deeper.

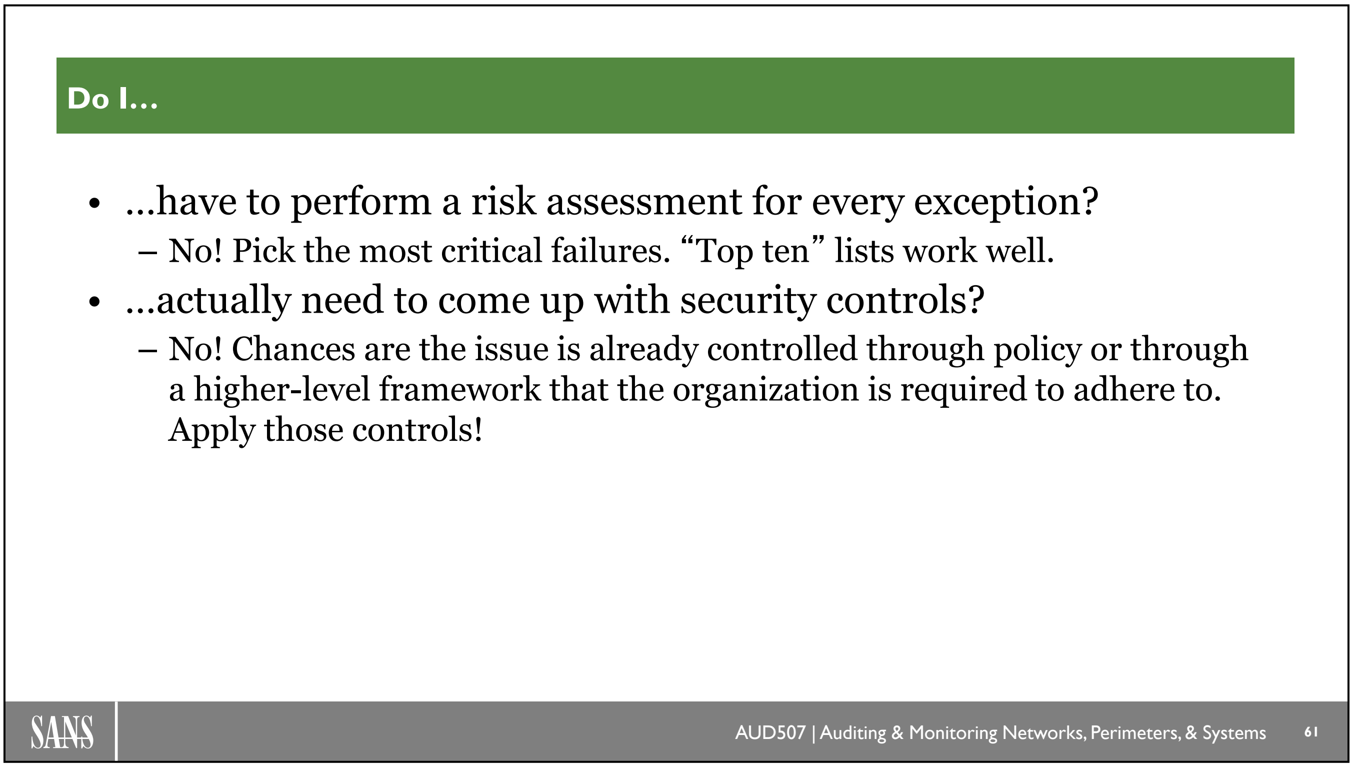

Do I…

- …have to perform a risk assessment for every exception?
  - No! Pick the most critical failures. "Top ten" lists work well.
- …actually need to come up with security controls?
  - No! Chances are the issue is already controlled through policy or through a higher-level framework that the organization is required to adhere to. Apply those controls!

Let's address a few other questions people often have about this. Some people wonder if the auditor has to perform a risk assessment for every single exception that he finds. No. Although it would be ideal to do so, we frankly just do not have the time. You also have to consider how many issues the organization can reasonably remediate. Of course, the organization should remediate all the issues that it discovers, but the more issues that are discovered and documented, the fewer that are remediated. You should report your findings in a useful way by picking the top ten most concerning issues and perform a CCA analysis on these issues. Use this to drive a top ten list of issues that are the highest risk and must be addressed without fail, along with your remediation recommendations from the analysis.

The other question is, "Do I need to create new controls to solve these problems?" Again, the answer is "no." It is quite common to perform analysis and discover that the underlying causes already have controls that are not applied correctly. (Or for this particular problem in the enterprise, the control was created but applied to something else, like putting a hose clamp on an electrical wire. It's still a great hose clamp, but it'll never keep the power from leaking out.)

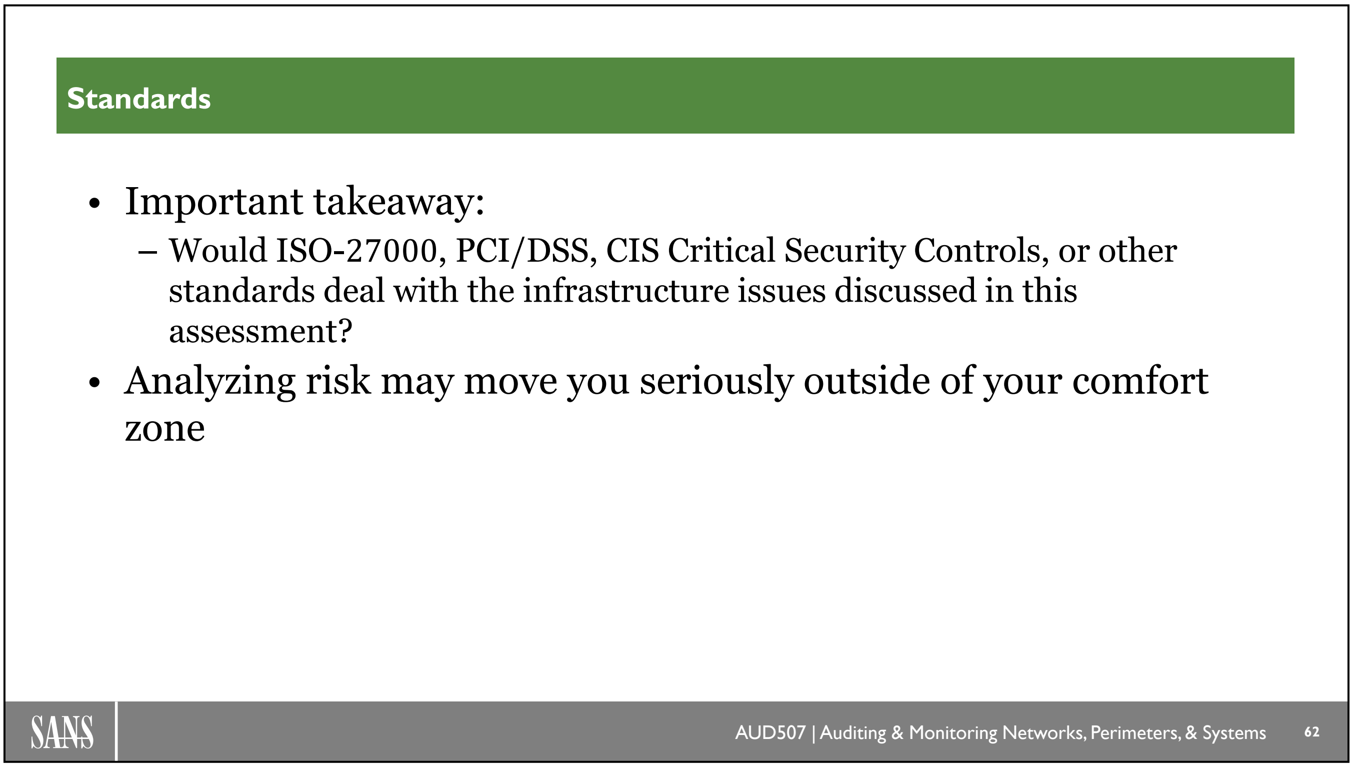

Standards

- Important takeaway:
  - Would ISO-27000, PCI/DSS, CIS Critical Security Controls, or other standards deal with the infrastructure issues discussed in this assessment?
- Analyzing risk may move you seriously outside of your comfort zone

Before moving to the next risk assessment, there's an important takeaway. Working through the event and fault trees, you had the opportunity to discover some serious control gaps within the enterprise, but they weren't what would normally be considered as IT controls. Even so, the infrastructure that supports your systems is critical to the capability of the enterprise to meet its objectives.

Certainly, there was some level of risk assessment performed that allowed you to arrive at a data center with dedicated power systems, including a UPS and a generator. The question is whether your organization consistently applies a formalized approach, forcing the enterprise to identify ongoing risks. Such an approach is extremely important to the delivery of CIA, but there is little to no mention of underlying infrastructure protections in the standards used to control risk in an Information Security setting.

How can you fix this? We encourage you to be brave in asking the critical questions. Doing so may require that you go far outside of your comfort zone to learn about standards for power system design, UPS design, fire suppression systems, and so on. In the end, however, you will not only offer an invaluable service to your enterprise but will also assist your enterprise in making the risk management process surrounding Information Assurance far more robust.

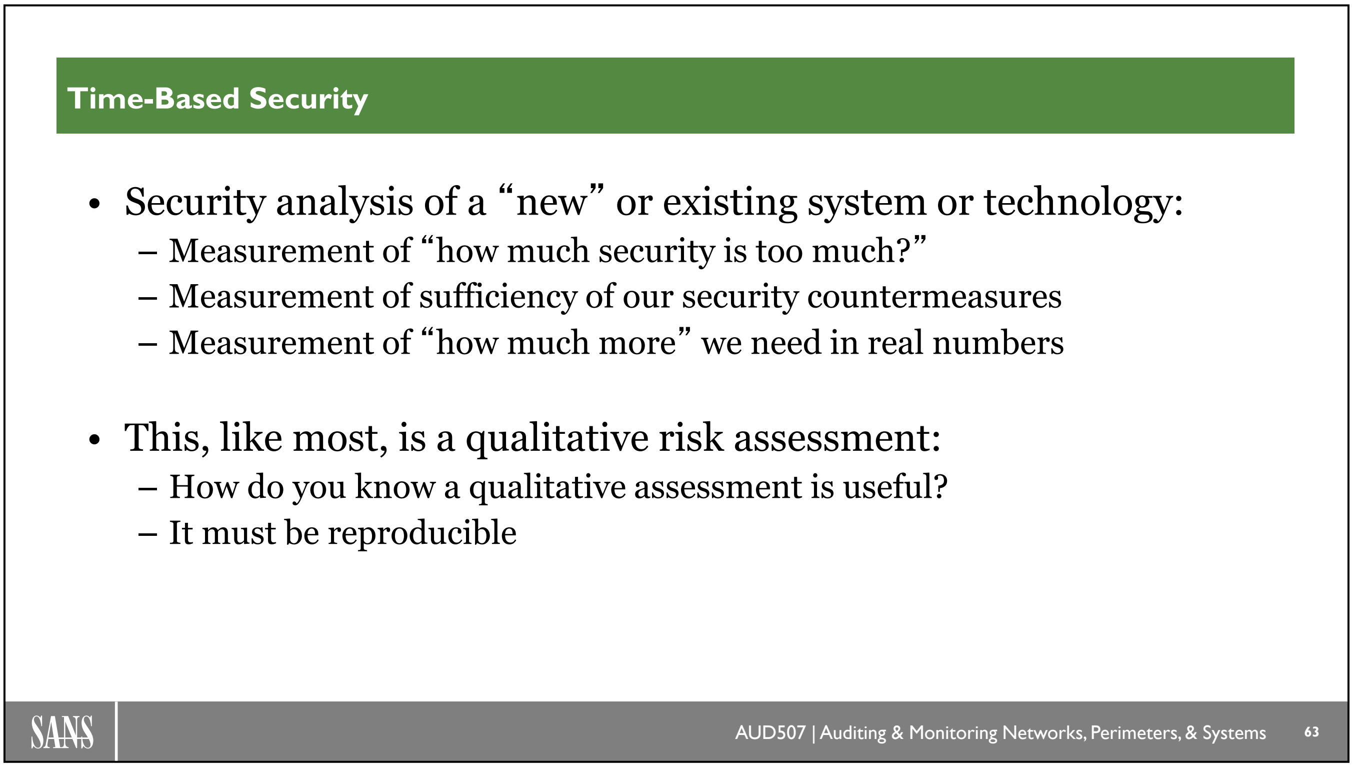

Time-Based Security

- Security analysis of a “new” or existing system or technology:
  - Measurement of “how much security is too much?”
  - Measurement of sufficiency of our security countermeasures
  - Measurement of “how much more” we need in real numbers
- This, like most, is a qualitative risk assessment:
  - How do you know a qualitative assessment is useful?
  - It must be reproducible

SANS AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 63

Whenever we, as security professionals, are presented with a new system or technology, we will inevitably be called upon to make a call as to how secure this new product or method is. Typically, this can become an involved process, requiring us to do all kinds of extra research so that we know this new system inside and out. In reality, however, few of us have the time or luxury to spend the amount of time required to master a new technology or product before there is some organizational requirement for the process or technology to be put into production. As a result, we need a formal, reproducible method to measure the potential risks involved with a new product, to determine how we can best secure this new product, and to communicate the entire process to management, who may not understand all the technical aspects involved.

TBS can enable us to meet these requirements in a general way. Although we can’t eliminate the need to have some knowledge of the new system to be implemented, we can measure that system against generalized criteria to make informed risk management decisions. Part of what this helps us to do is to determine exactly how secure that system is today and what additional measures we must take to secure this new system adequately.

Of course, TBS is fantastic to analyze our existing systems and networks, too! In traditional auditing, we either run through a checklist against a system, measure the current system against baselines, or discard “known good” log entries to drop out anomalies on a system. Using Time-Based Security, we can move outside of the traditional auditing box and measure the systems we have in place in another dimension. Actually, we can measure how intrusion tolerant our network and attached systems are. As with most auditing and security practices, the application of TBS is not restricted to computer and network hardware. Truly, TBS applies to any defensive security application.

Great byproducts of TBS are the determination of how much more security we need to reach an acceptable level of risk, as well as where we can focus our security dollars to produce the best effect on our overall security infrastructure.

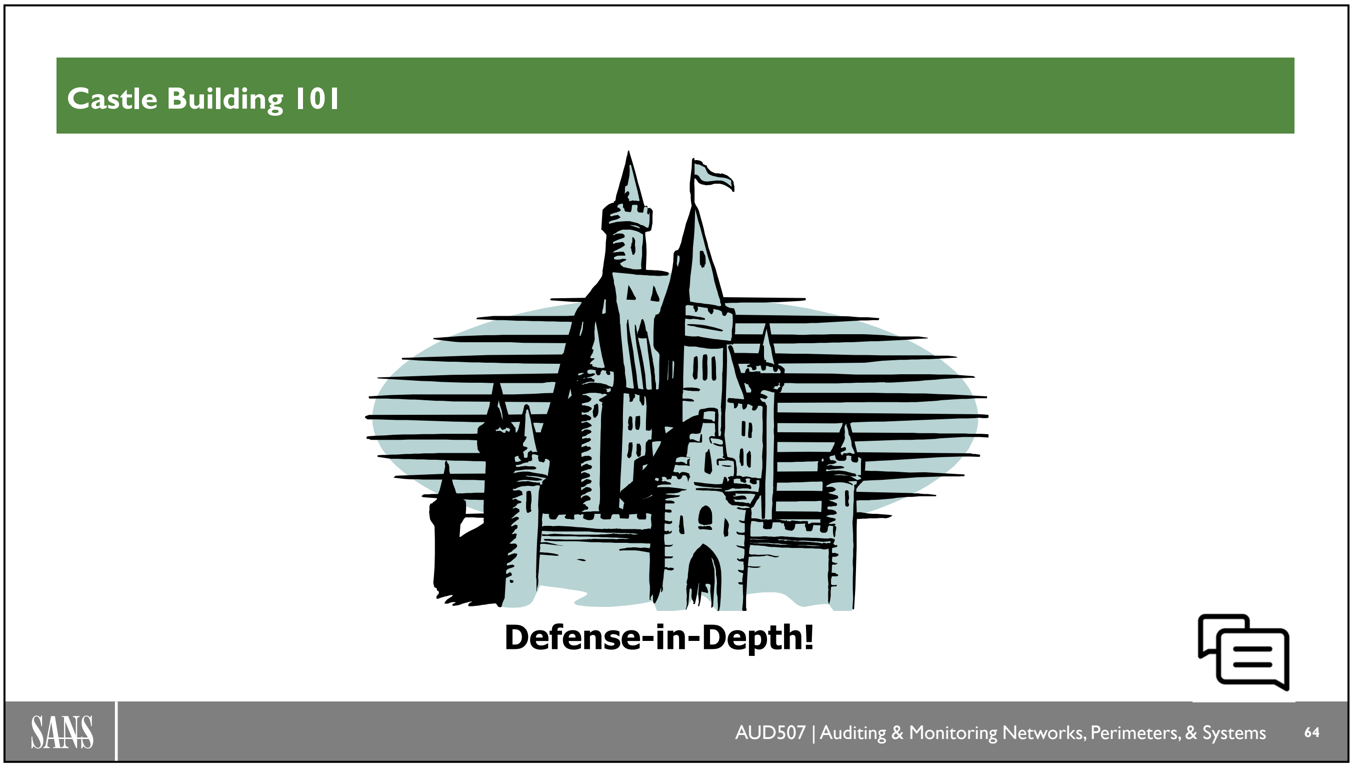

Castle Building 101

Defense-in-Depth!

Security and defense concepts haven't changed a whole lot since the Middle Ages. For those of you who are already familiar with Defense-in-Depth, let's consider the concept. Defense-in-Depth is related to TBS. TBS is a measurement tool for the effectiveness of our level of Defense-in-Depth.

People who built castles knew what they were doing. A good part of their expertise in building castles in secure locations came from the horrible experience of having to rebuild a castle that got knocked down. Take a moment and imagine yourself as a medieval castle contractor. What are some considerations that go into building a new castle for the newest king? Most castles are built either out in the middle of a lake, on the top of a mountain, or on a cliff. Such placement is not a coincidence. There are even more things that most castles have in common. If the castle is on the top of a hill, for instance, nearly all the trees on that entire hill have been cut down. A moat might surround the castle. Around the hill, you'll typically find some low walls, too. All these features are involved with security!

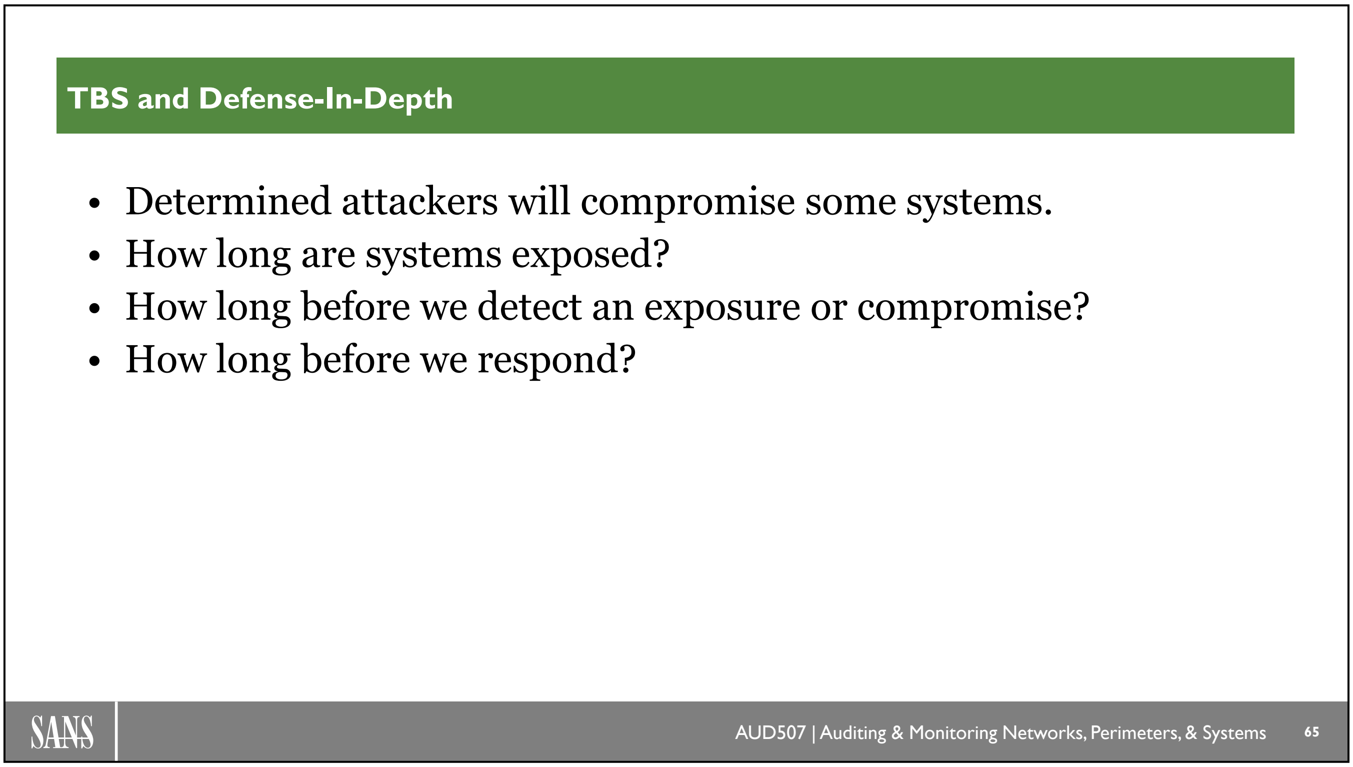

TBS and Defense-In-Depth

- Determined attackers will compromise some systems.
- How long are systems exposed?
- How long before we detect an exposure or compromise?
- How long before we respond?

For instance, those low walls around the hill aren't intended to keep the attackers out but are intended to slow them down just a bit. Also, the trees were cut down, not out of any protest against medieval tree-huggers, but because it gave the guards in the castle a better view of incoming attackers. The moat, too, likely would not stop a determined attacker, but it would slow them down even more. Finally, the big, tall, thick walls of the castle could potentially be breached, but hopefully by the time someone was trying to scale the walls, you already had plenty of notice so you could mobilize your troops.

This is precisely how TBS fits in with Defense-in-Depth. What we've just described with our castle *is* Defense-in-Depth! The motivation behind it is to create an intrusion-tolerant environment in which we recognize that people who don't have our best interests at heart may compromise some portion of our defenses. Defense-in-Depth raises the bar for attackers so that it becomes more difficult to achieve their ultimate goal: Getting into the castle and stealing the crown jewels.

Time-Based Security measures how long it takes under the best and the worst circumstances to respond to an attack. It also helps identify where weaknesses in our Defense-in-Depth strategy lie. It answers the three questions in our slide.

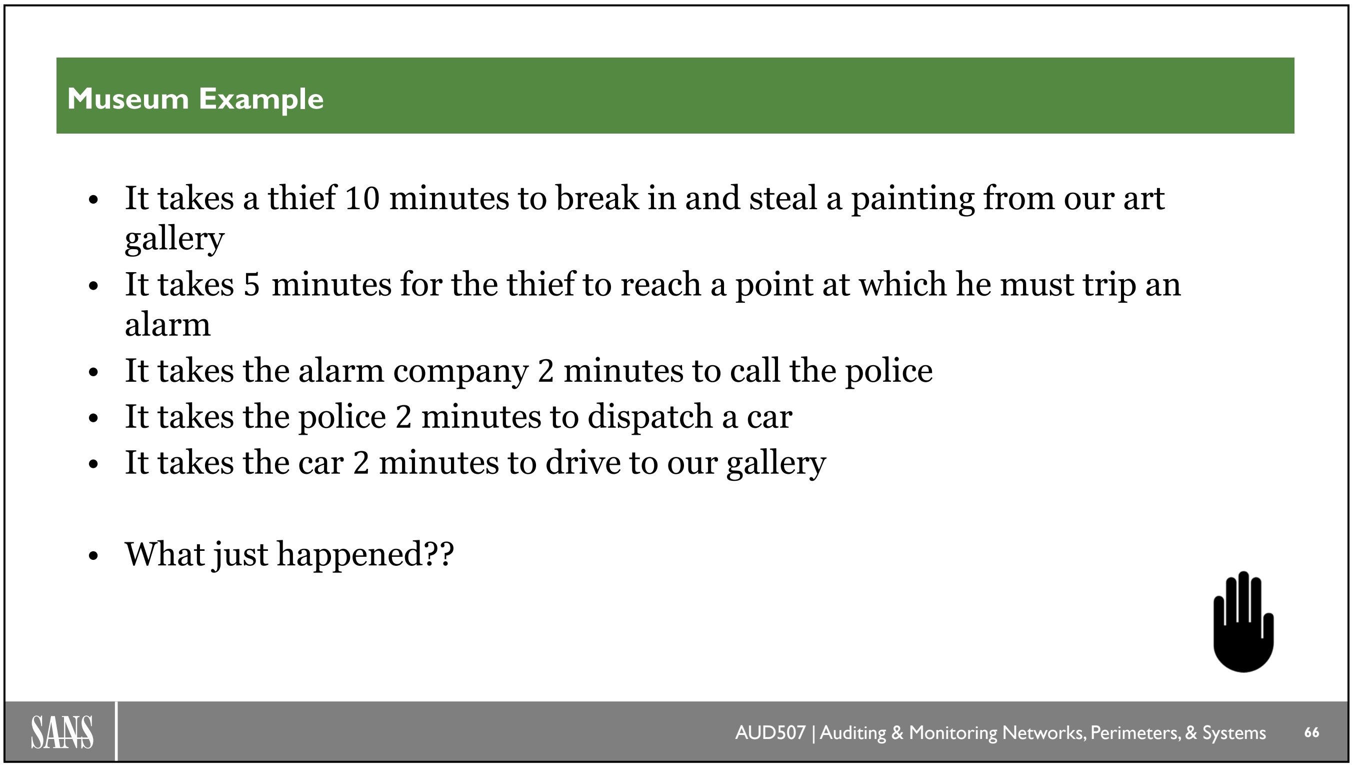

Museum Example

- It takes a thief 10 minutes to break in and steal a painting from our art gallery
- It takes 5 minutes for the thief to reach a point at which he must trip an alarm
- It takes the alarm company 2 minutes to call the police
- It takes the police 2 minutes to dispatch a car
- It takes the car 2 minutes to drive to our gallery

- What just happened??

Think of it this way (don't worry, this isn't going to be one of those word problems from grade school math class):

It takes a thief 10 minutes to break in and steal a painting from our art gallery.

It takes 5 minutes for the thief to reach a point at which he must trip an alarm.

It takes the alarm company 2 minutes to call the police.

It takes the police 2 minutes to dispatch a car.

It takes the car 2 minutes to drive to our gallery.

Guess what? That means we just lost a painting! The alarm system worked fine, the police responded, but we still lost!! Where'd we go wrong??

$(P)10 < 5(D) + (R)6$

Even though everything worked fine, the problem is that when we measure our security measures against the time dimension, we discover that our detection took too long, and the response took too long!

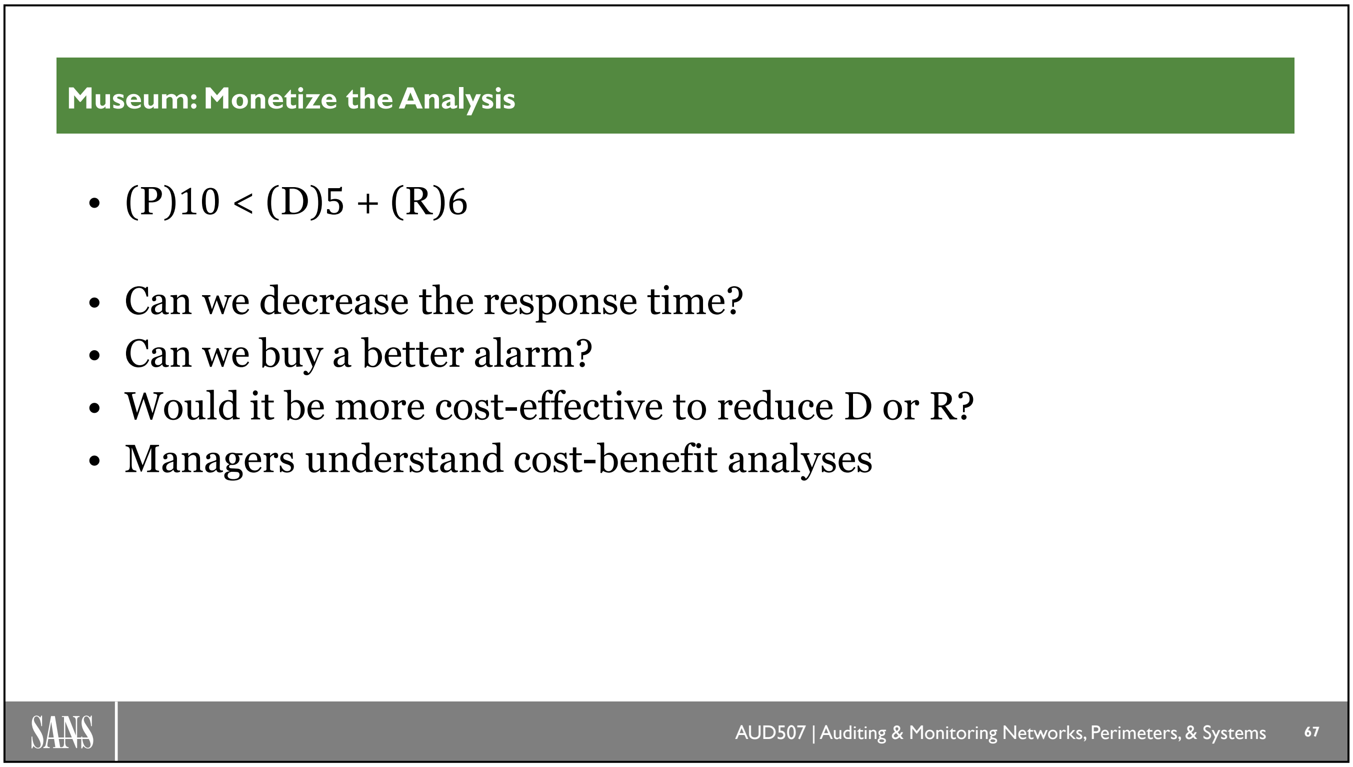

Museum: Monetize the Analysis

- (P)10 < (D)5 + (R)6

- Can we decrease the response time?
- Can we buy a better alarm?
- Would it be more cost-effective to reduce D or R?
- Managers understand cost-benefit analyses

Looking at the formula in the slide, remember that P = Protection, D = Detection, and R = Response.

Although applying TBS to our systems before there’s an incident can have huge payoffs (we’re going to discuss how to go about applying TBS to measure our defenses), let’s continue with the example from the last slide for a moment. Now we’ll look at what TBS quantifies for us when we do a lessons learned after our art heist incident. We can now look objectively at the problem in terms of numbers and determine what we can do better instead of pointing fingers at each other, wondering who left the door unlocked.

Looking at this example, it’s probably clear that there’s not a whole lot that we can do about the response time. The alarm monitoring company followed proper procedures and notified the police quickly, the police dispatched a car quickly, and the car got to the gallery quickly. The problem is, by the time the car arrived, our thief had been gone for more than 1 minute. What about the alarm system? If we were going to invest dollars to fix this problem, TBS shows us clearly that the alarm system is where we need to spend the money. If we detected the intrusion just 1 minute sooner, we would have stood a good chance of catching the intruder.

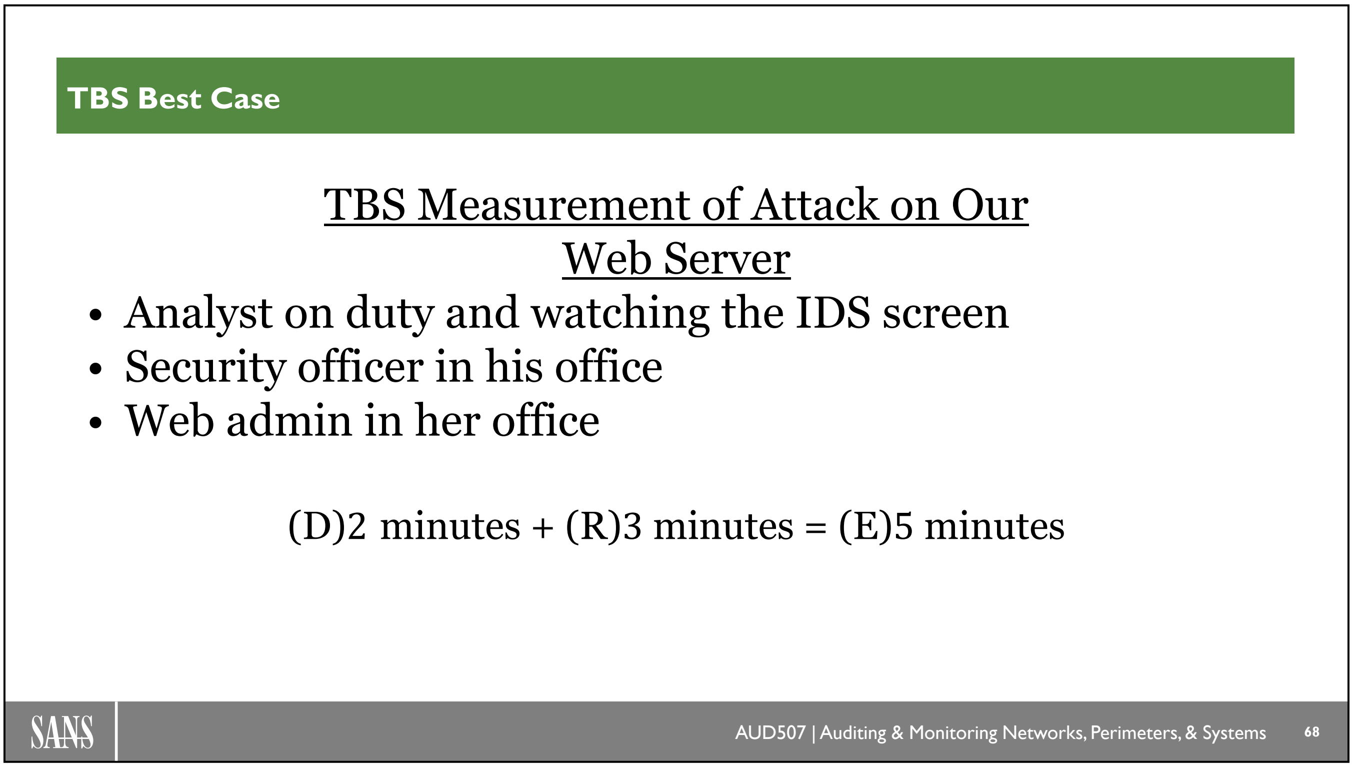

TBS Best Case

TBS Measurement of Attack on Our Web Server

- Analyst on duty and watching the IDS screen
- Security officer in his office
- Web admin in her office

(D)2 minutes + (R)3 minutes = (E)5 minutes

First, get a stopwatch. Next, sit down with an Excel spreadsheet and start measuring data. Begin by firing a sample signature at your web server while the analyst is watching. As soon as the IDS alerts the analyst, the analyst examines the data and validates the incident. He follows his escalation procedure and calls the security officer, who looks at the data and certifies that it does appear to be an event. He immediately calls the web admin, who is also in her office, and they then respond together to handle the incident and effect an appropriate response.

D: Clock starts ticking when the analyst sees an event. It stops when he has validated data and sends an alert.

R: Clock starts when the analyst calls the security officer and stops when he and/or the web administrator close the hole.

So, our total exposure time, if our protection device fails completely, would be 5 minutes in this best-case scenario.

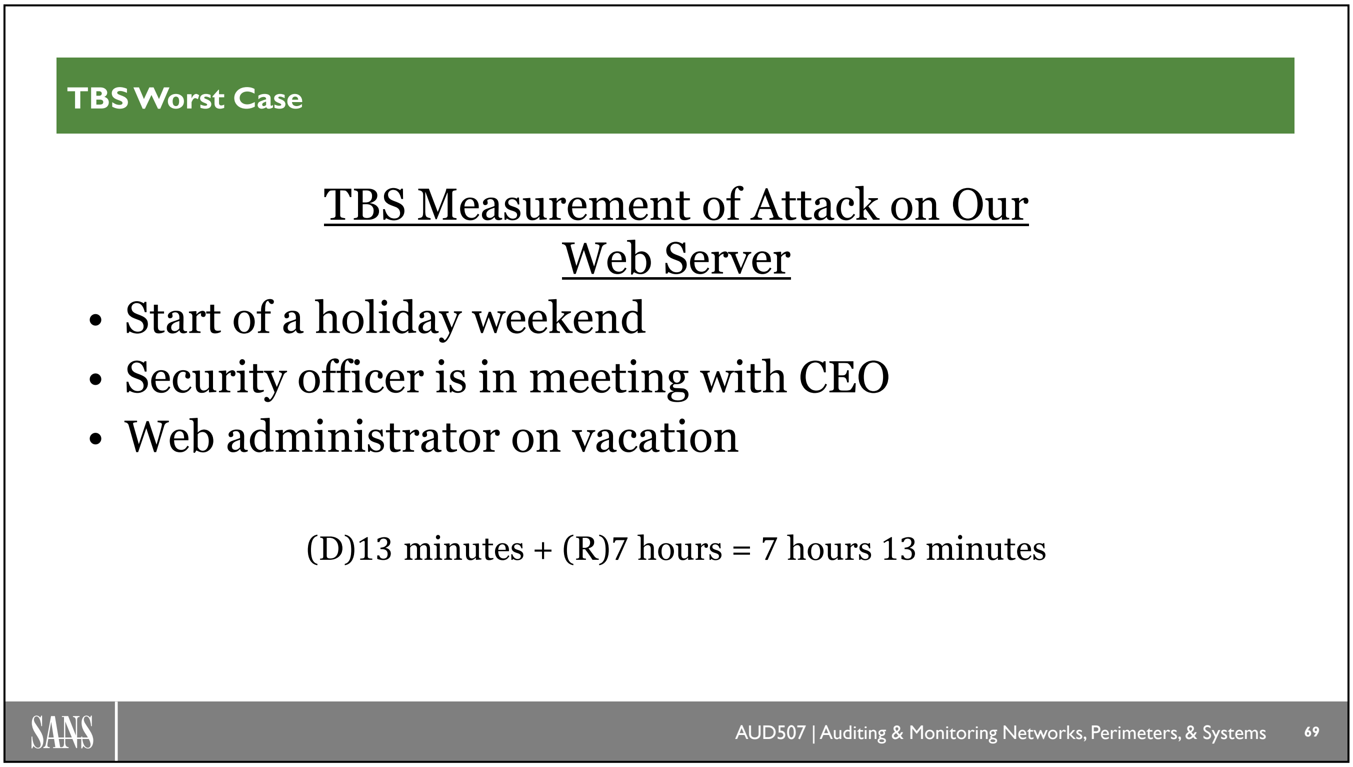

TBS Worst Case

TBS Measurement of Attack on Our Web Server

- Start of a holiday weekend
- Security officer is in meeting with CEO
- Web administrator on vacation

(D)13 minutes + (R)7 hours = 7 hours 13 minutes

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 69

Now, the worst-case scenario is where you want someone with an auditing mentality to take over this whole thing. If you asked the security officer, the analyst, or the web administrator to fill in these values for a worst case, chances are they'd be generous. We want someone who is detached from the actual detection and response flow to objectively estimate values. How can we do this? How about trying to get the security officer on the phone or to respond to an email 15 minutes after he leaves the facility? How about talking to the department administrative assistant to see how long it takes to get the web admin on the phone when she's away on vacation? Of course, the auditor would also check to see if there's an alternative escalations list, which should reduce D and R in the worst case.

Let's imagine that we email the security officer, who manages to get away from his meeting 10 minutes later. It takes another 3 minutes to go over the data with him on the phone, and he gives the okay to the analyst to contact the web admin. That means that D = 13 so far.

Now, the analyst starts trying to contact the web administrator. Let's imagine that the web admin is on the West Coast in Vancouver enjoying some of the best skiing on the face of the Earth up at Blackcomb/Whistler. It's 2 PM there, so by the time the web admin gets off the mountain, goes to the sauna, and then gets back to her room to check her email, 5 hours have passed. She calls the office back and assists over the phone with the evaluation and response on the web server, which takes another 2 hours at this point. That means that R equals 7 hours.

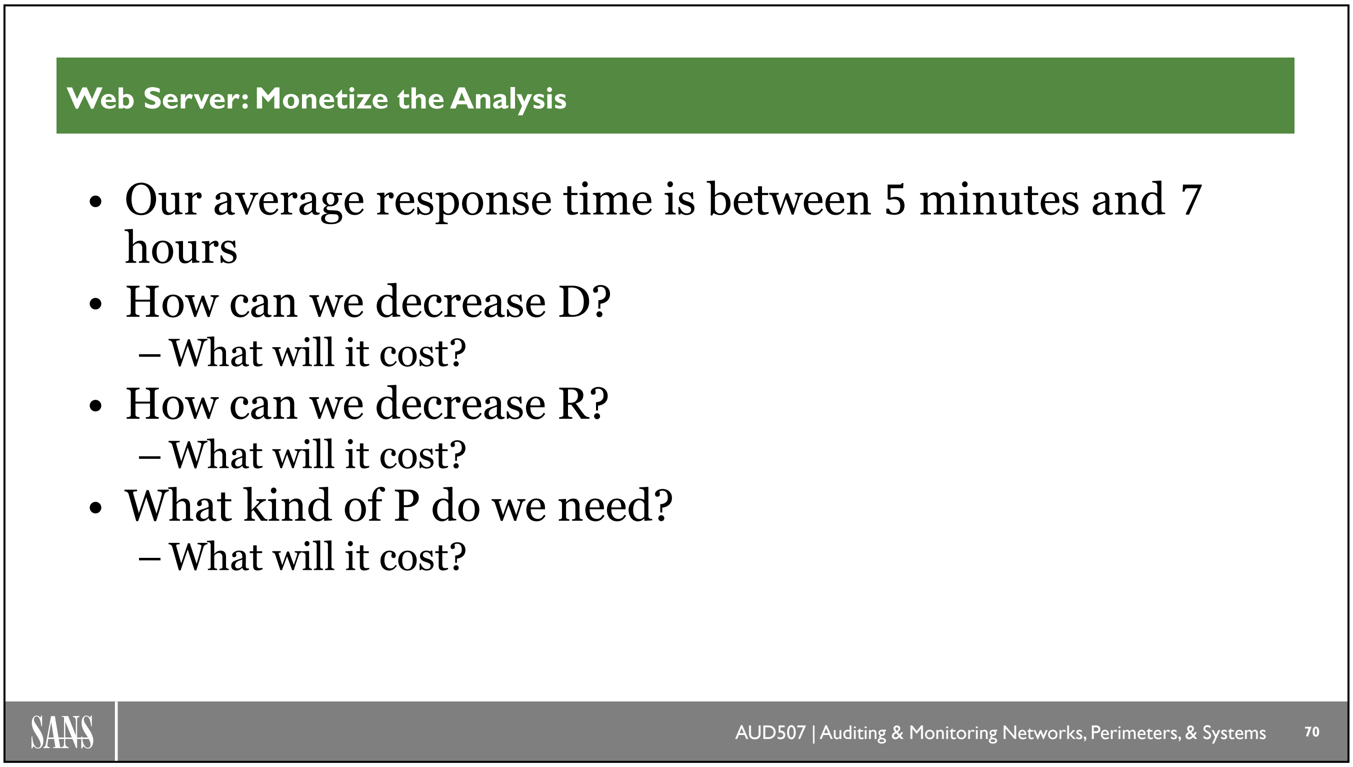

Web Server: Monetize the Analysis

- Our average response time is between 5 minutes and 7 hours
- How can we decrease D?
  - What will it cost?
- How can we decrease R?
  - What will it cost?
- What kind of P do we need?
  - What will it cost?

SANS AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 70

Here's the payoff. Most organizations attempt to reduce their exposure to threats by reducing the detection time. Reducing detection time can be effective, but it quickly becomes costly. However, in our examples, our TBS numbers indicate that our detection was not that much worse (comparatively) in our worst-case scenario. Rather, what it showed was that our reaction and response time was downright awful! Reducing our reaction time is generally far cheaper than reducing detection times. Actually, what was revealed in our worst-case scenario is the need for better policies and procedures for the response portion of incident handling. Perhaps have an alternative escalation contact for web server events?

In addition, we can now derive an actual requirement for P. If we estimate that our average E = D + R is approximately 20 minutes, we now know that we need to purchase a piece of protection equipment that will last for at least that long. Relate it back to our art gallery. What we didn't analyze before was why there was a 5-minute gap between the time the intruder gained access to our building and the time that the alarm system detected him. Buying a new alarm system and rewiring the gallery could be prohibitively expensive. What if we invested in steel curtains to roll down over the doors and windows after closing? What we're doing is increasing P, possibly at a substantially lower cost than reducing D or R.

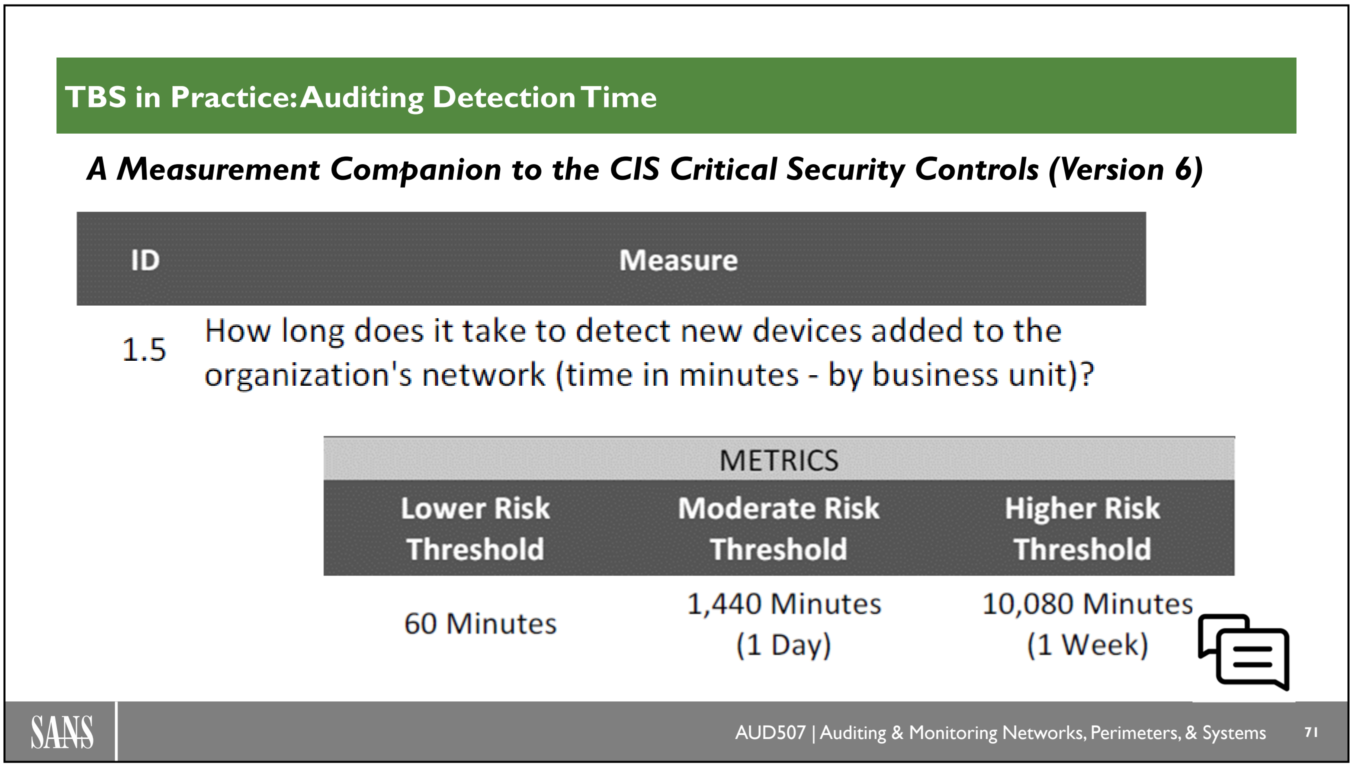

TBS in Practice: Auditing Detection Time

A Measurement Companion to the CIS Critical Security Controls (Version 6)

ID	Measure
1.5	How long does it take to detect new devices added to the organization's network (time in minutes - by business unit)?

METRICS		
Lower Risk Threshold	**Moderate Risk Threshold**	**Higher Risk Threshold**
60 Minutes	1,440 Minutes (1 Day)	10,080 Minutes (1 Week)

The Measures and Metrics document for Version 8 of the Critical Controls rates an organization’s risk exposure based on how many subnets are covered by a host discovery control, rather than using substantive testing of how well the control works. While this is helpful in determining the magnitude of risk, you may find it to be less effective because it ignores the function of the control under consideration.

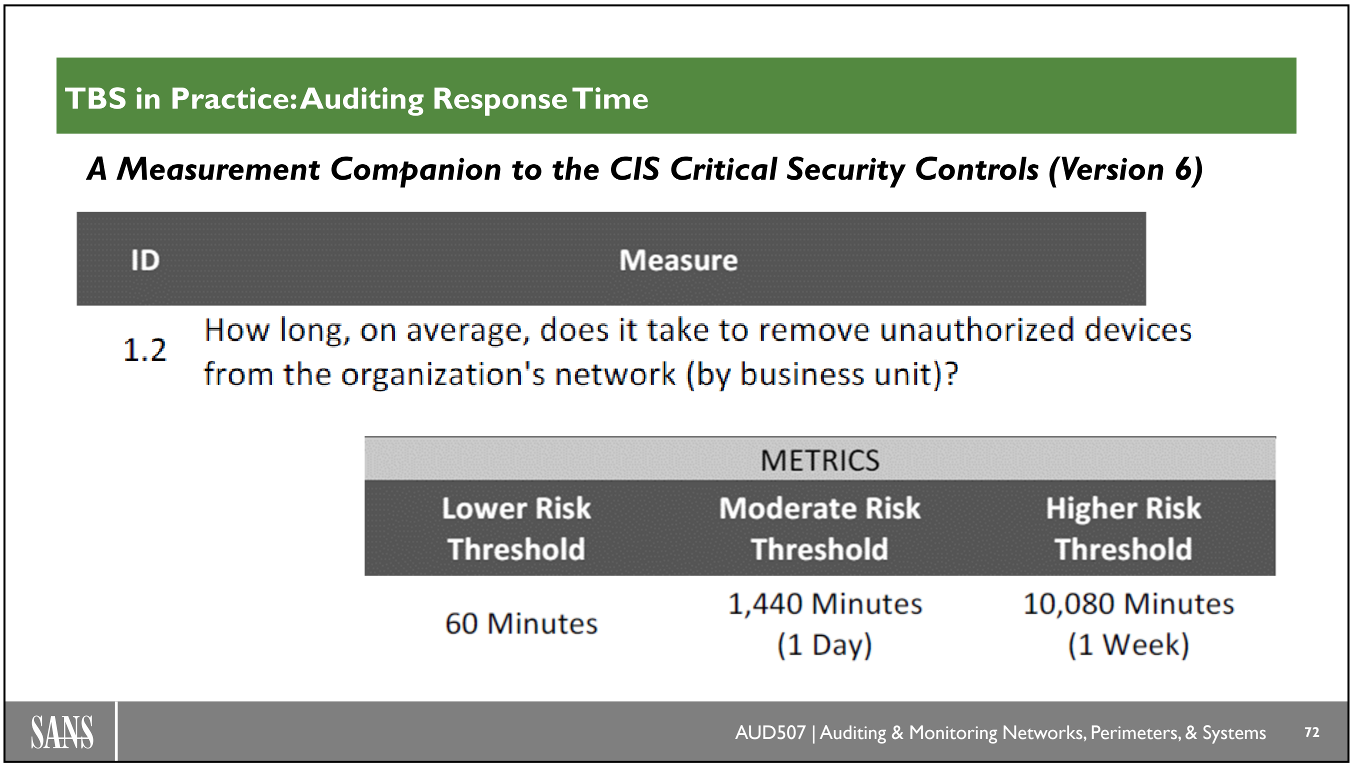

TBS in Practice: Auditing Response Time

A Measurement Companion to the CIS Critical Security Controls (Version 6)

ID	Measure
1.2	How long, on average, does it take to remove unauthorized devices from the organization's network (by business unit)?

METRICS		
Lower Risk Threshold	**Moderate Risk Threshold**	**Higher Risk Threshold**
60 Minutes	1,440 Minutes (1 Day)	10,080 Minutes (1 Week)

Here is an example of TBS used for real-world auditing. The audit guide developed by the Center for Internet Security for Version 6 of the Critical Controls (sadly, the measurement guide for Version 8 uses a much different measurement technique) included many time-based metrics for an auditor to use in verifying the correct function of the controls within an enterprise.

It would be easy to devise a test for Measures 1.2 and 1.5 (on the next page). Simply place a benign device on the subnet to be tested and measure how long it takes to detect (Measure 1.5) and subsequently remove (Measure 1.2) the device.

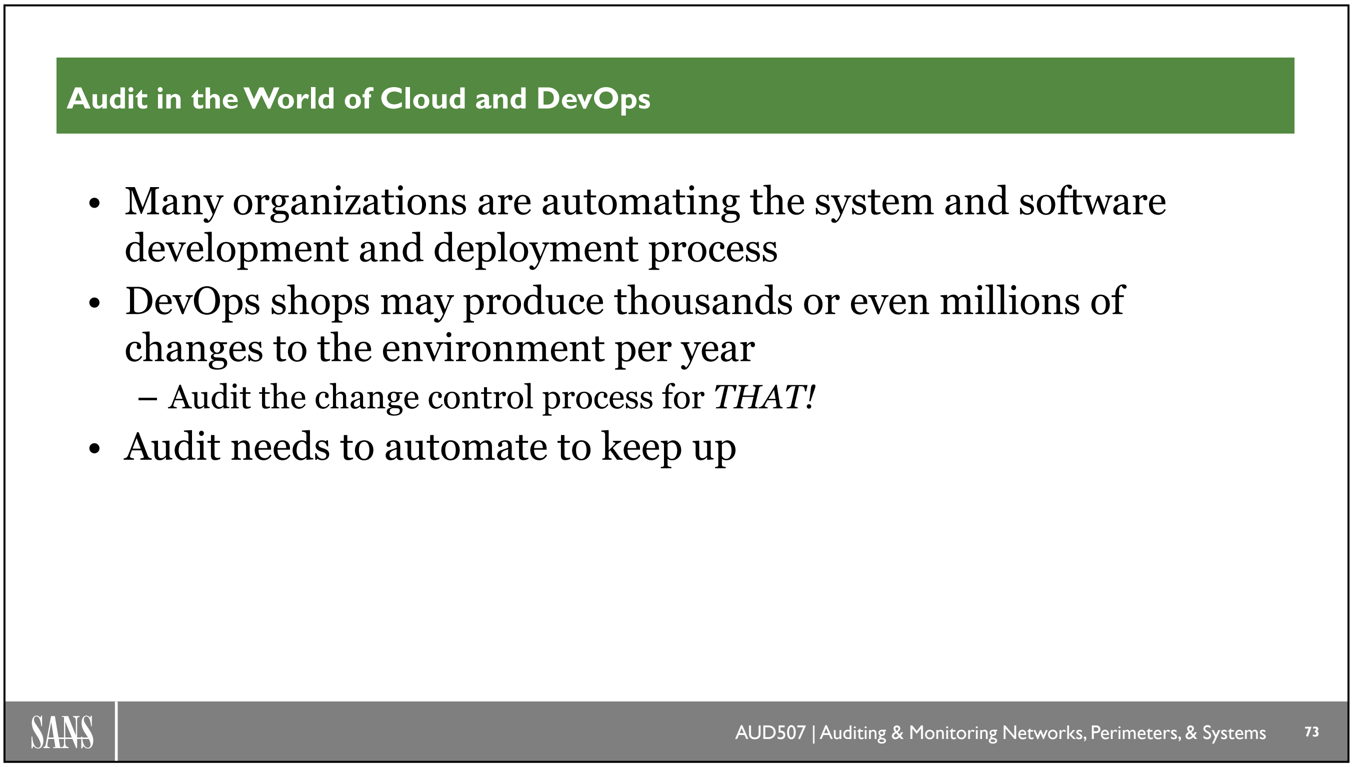

Audit in the World of Cloud and DevOps

- Many organizations are automating the system and software development and deployment process
- DevOps shops may produce thousands or even millions of changes to the environment per year
  - Audit the change control process for *THAT!*
- Audit needs to automate to keep up

No discussion of the role of auditors would be complete without a discussion of emerging technologies.

New development techniques can result in thousands of changes to the production environment in a month. Imagine being asked to audit the change control process for these thousands of changes. It's important that we, as auditors, keep up. We may need to do this by better integrating into the development or deployment environment, using scripts to automate auditing, placing non-intrusive controls into the development pipeline.

It will be critical in coming years that you, as an auditor, keep up with the changing enterprise environment if you want to stay relevant.

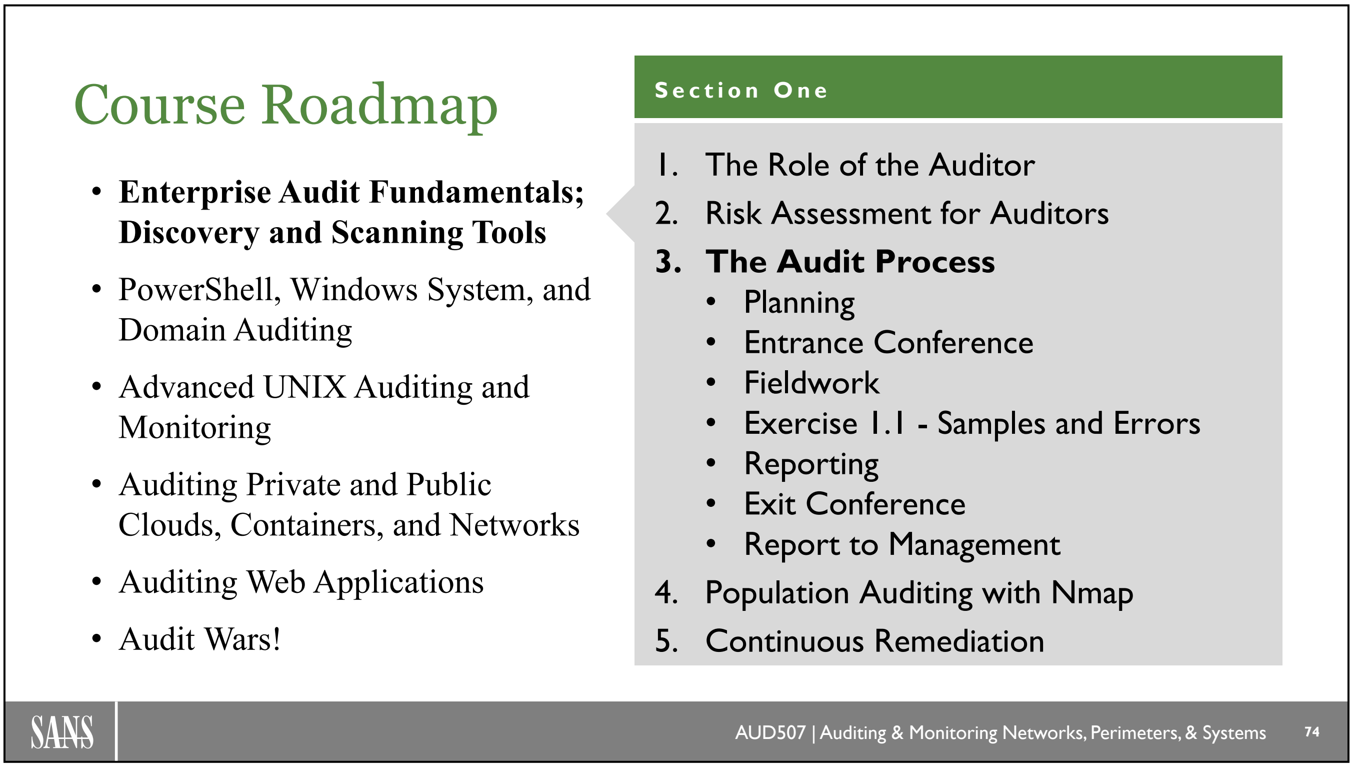

Course Roadmap

- **Enterprise Audit Fundamentals; Discovery and Scanning Tools**
- PowerShell, Windows System, and Domain Auditing
- Advanced UNIX Auditing and Monitoring
- Auditing Private and Public Clouds, Containers, and Networks
- Auditing Web Applications
- Audit Wars!

Section One

1. The Role of the Auditor
2. Risk Assessment for Auditors
3. **The Audit Process**
   - Planning
   - Entrance Conference
   - Fieldwork
   - Exercise 1.1 - Samples and Errors
   - Reporting
   - Exit Conference
   - Report to Management
4. Population Auditing with Nmap
5. Continuous Remediation

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 74

This page intentionally left blank.

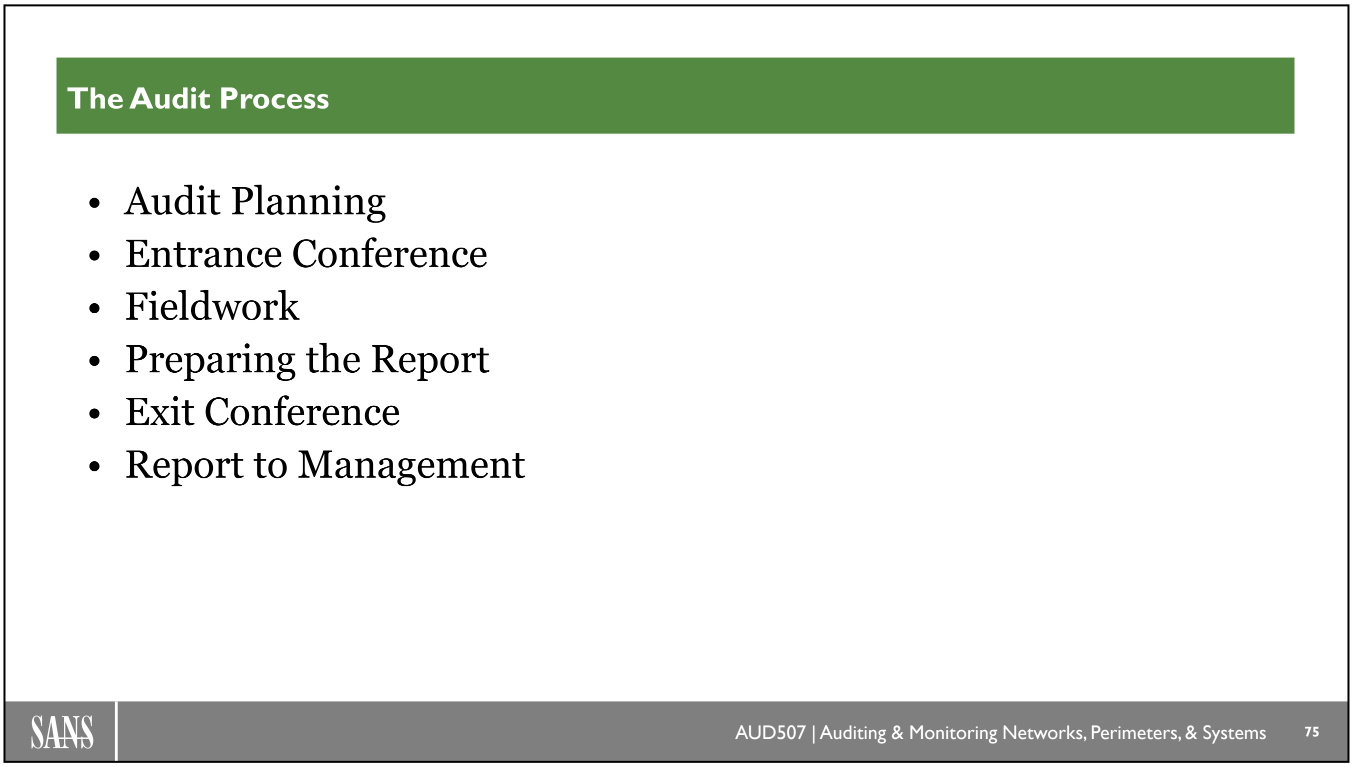

The Audit Process

- Audit Planning
- Entrance Conference
- Fieldwork
- Preparing the Report
- Exit Conference
- Report to Management

This slide shows each of the six steps in the audit process. We’ll take each in turn in the subsequent slides and talk about them in depth. As a brief overview, audit planning is part of what we’ve been doing all day! We reviewed audit strategies that can measure systems. Audit planning is all the activities leading up to performing an actual audit.

The entrance conference is essentially when the auditor meets with the individuals responsible for whatever it is that is audited.

Fieldwork is the actual measurement of the systems.

After completing the fieldwork, the auditor can take the necessary time to prepare the audit report, which will be presented at the exit conference and finally, an executive version of the report will be prepared and communicated to upper management.

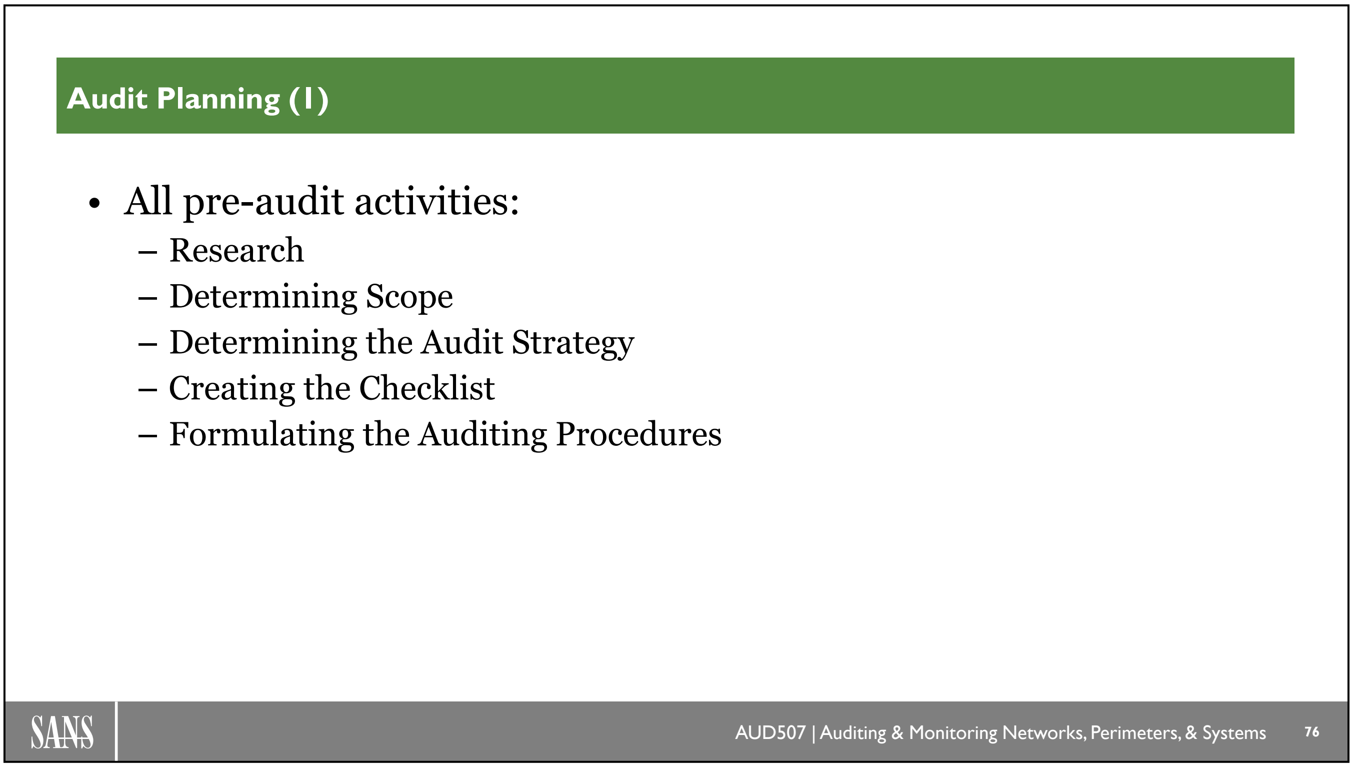

Audit Planning (1)

- All pre-audit activities:
  - Research
  - Determining Scope
  - Determining the Audit Strategy
  - Creating the Checklist
  - Formulating the Auditing Procedures

All the steps in the audit process are crucial, but audit planning truly is the most crucial step. This phase may be kicked off in a number of different ways. Either the auditor is engaged to perform an audit against an organization or, if we are talking about an internal auditor, he becomes aware of the need for an audit to be performed against some aspect of the organization. Either way, the first thing that the auditor will do is research! During this phase, the auditor also spends time determining the scope, either by working for clarification with the contracting individual or manager or perhaps by working with an internal security officer. While this research is being done, the auditor will develop a strategy for auditing the process or system and will formulate a checklist. Hand in hand with the checklist will be the creation of the actual auditing procedures or the "How To" guide for the auditor.

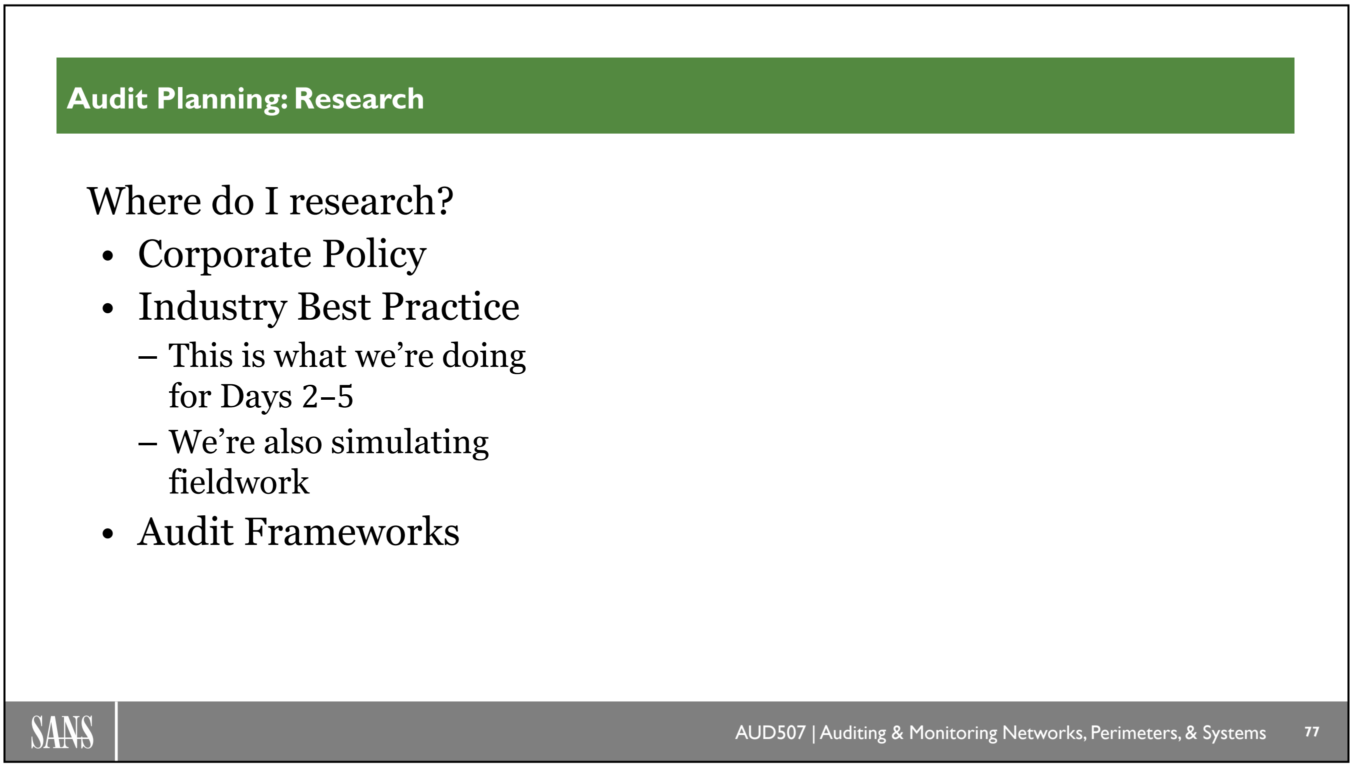

Audit Planning: Research

Where do I research?

- Corporate Policy
- Industry Best Practice
  - This is what we're doing for Days 2–5
  - We're also simulating fieldwork
- Audit Frameworks

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 77

Now let's talk about some of the "How To's." I've talked to auditors who have been through various training programs, and when they leave some say, "This sounds like a great process, but how do I do it?" One of the biggest questions they have is, "Where do I do the research to come up with an audit strategy?" The answer is internal policy documents coupled with industry best practices. An excellent resource for computer security best practice is the Center for Internet Security (https://u.aud507.com/1-3). CIS has taken the time to put out best practice documents and step-by-step guides for securing certain systems. These documents can easily be turned into auditing checklists and procedures. They also should not be overlooked when you're doing the actual audit because they recommend a wide variety of tools that can measure all sorts of aspects of the systems in question.

It's important to repeat something that we said in one of today's earlier slides. As we go through the material for Days 2 through 5, we're not trying to turn you into administrators. Instead, we walk you through the depth and type of research that is necessary to perform an effective audit of an information system. We also simulate some of the fieldwork so that as you design audit programs in your organization, you will know how to perform the tests and how to evaluate the results.

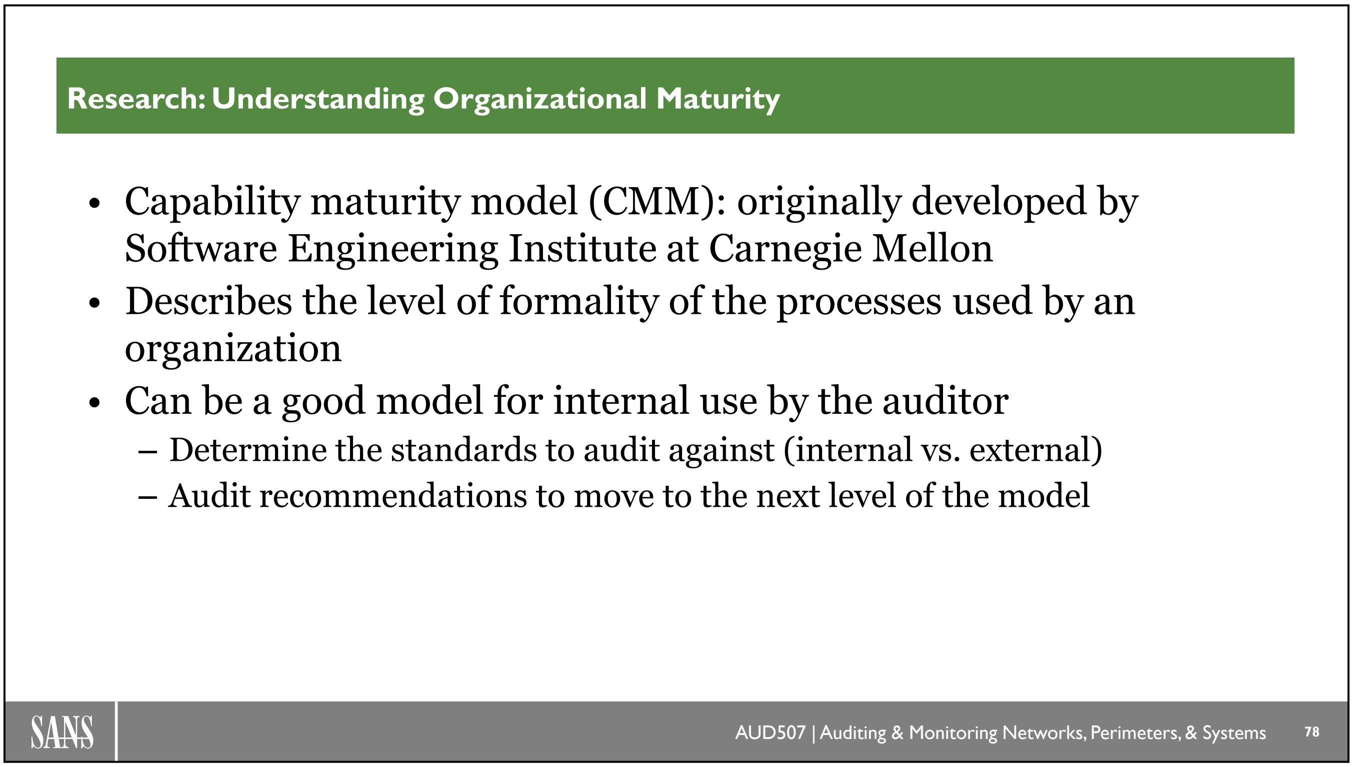

Research: Understanding Organizational Maturity

- Capability maturity model (CMM): originally developed by Software Engineering Institute at Carnegie Mellon
- Describes the level of formality of the processes used by an organization
- Can be a good model for internal use by the auditor
  - Determine the standards to audit against (internal vs. external)
  - Audit recommendations to move to the next level of the model

Back in the 1980's, the US Department of Defense funded research by the Software Engineering Institute into how to measure the organizational maturity of software development shops. The resulting model, called the Capability Maturity Model, has been the basis of a lot of subsequent analysis of how organizations mature in their management functions.

Gaining at least an informal understanding of the auditee's maturity level can be very informative for the auditor. In fact, I'm going to let you in on one of the dirty secrets of auditing: very often your report will consist of recommendations that the organization move from one maturity level to the next.

On the next slide, we see the levels of the original CMM.

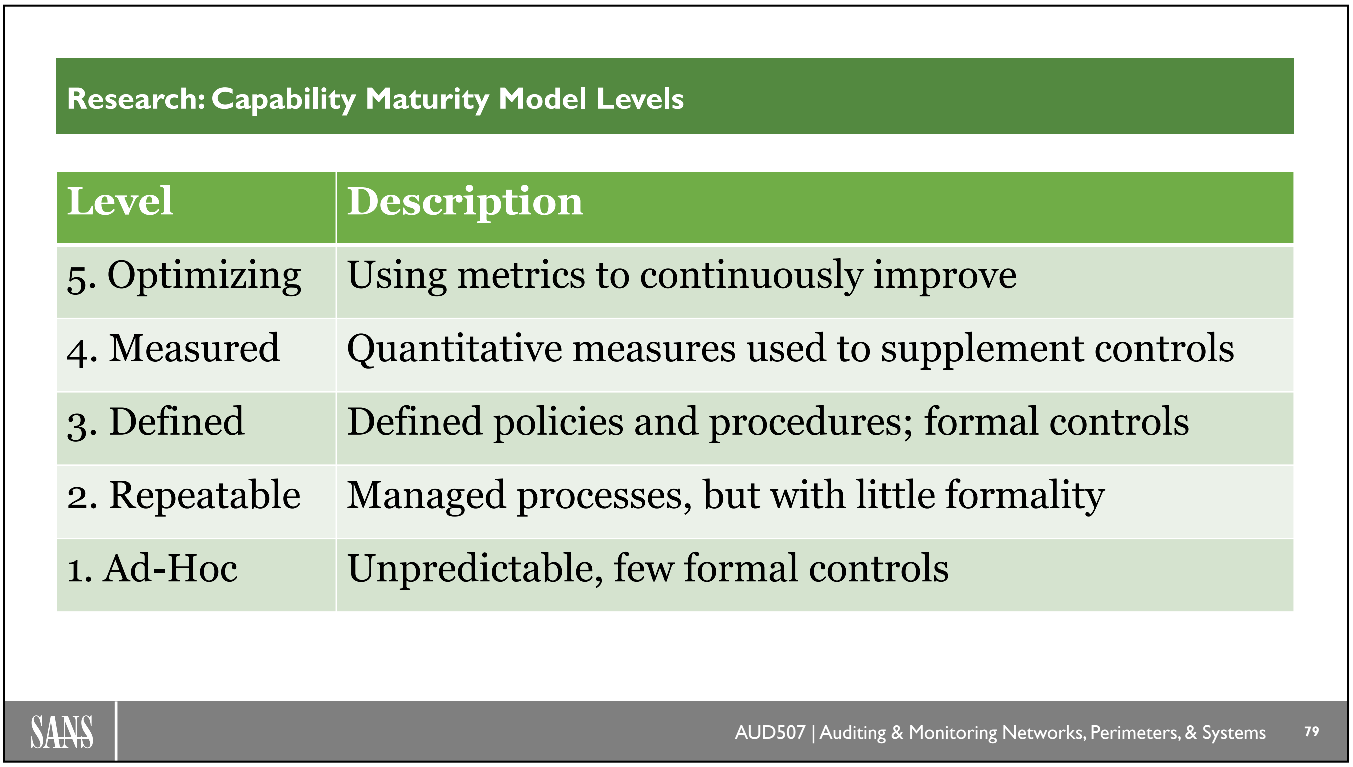

Research: Capability Maturity Model Levels

Level	Description
5. Optimizing	Using metrics to continuously improve
4. Measured	Quantitative measures used to supplement controls
3. Defined	Defined policies and procedures; formal controls
2. Repeatable	Managed processes, but with little formality
1. Ad-Hoc	Unpredictable, few formal controls

The CMM defines five level of maturity for the organization:

1. Ad-hoc: The organization makes decisions as the need arises. There is very little in the way of formal process or controls and usually almost no documentation. Audit recommendations for a Level 1 organization will often be to decide on standardized processes and controls and to strive for consistent implementation. Auditing an organization at this maturity level will usually require that the auditor and management impose an external standard as a basis for the audit.
2. Repeatable: The organization has begun to do things the same way most of the time. There is still little in the way of formal documentation, but processes and controls are easier to identify. Auditing an organization at this level will require the imposition of some sort of formality, either through a documentation effort or by using an external standard. Recommendations will always include the need to document processes and controls formally.
3. Defined: The organization has formal policies, procedures, processes, and controls. At this level, we can audit against the organization's own standards, and overall recommendations will tend to push the organization toward measuring the effectiveness of their controls.
4. Measured: The organization uses some sort of measurement scheme to determine if their controls are functioning correctly. Audits will be performed against the defined policies, procedures, processes, and controls. Recommendations will often tend toward using the control measurements to optimize the effectiveness and efficiency of controls.
5. Optimizing: The organization is actively feeding measurement information into the control planning process. Audit recommendations are usually geared toward making more effective use of the measurement data or better optimizing control processes.

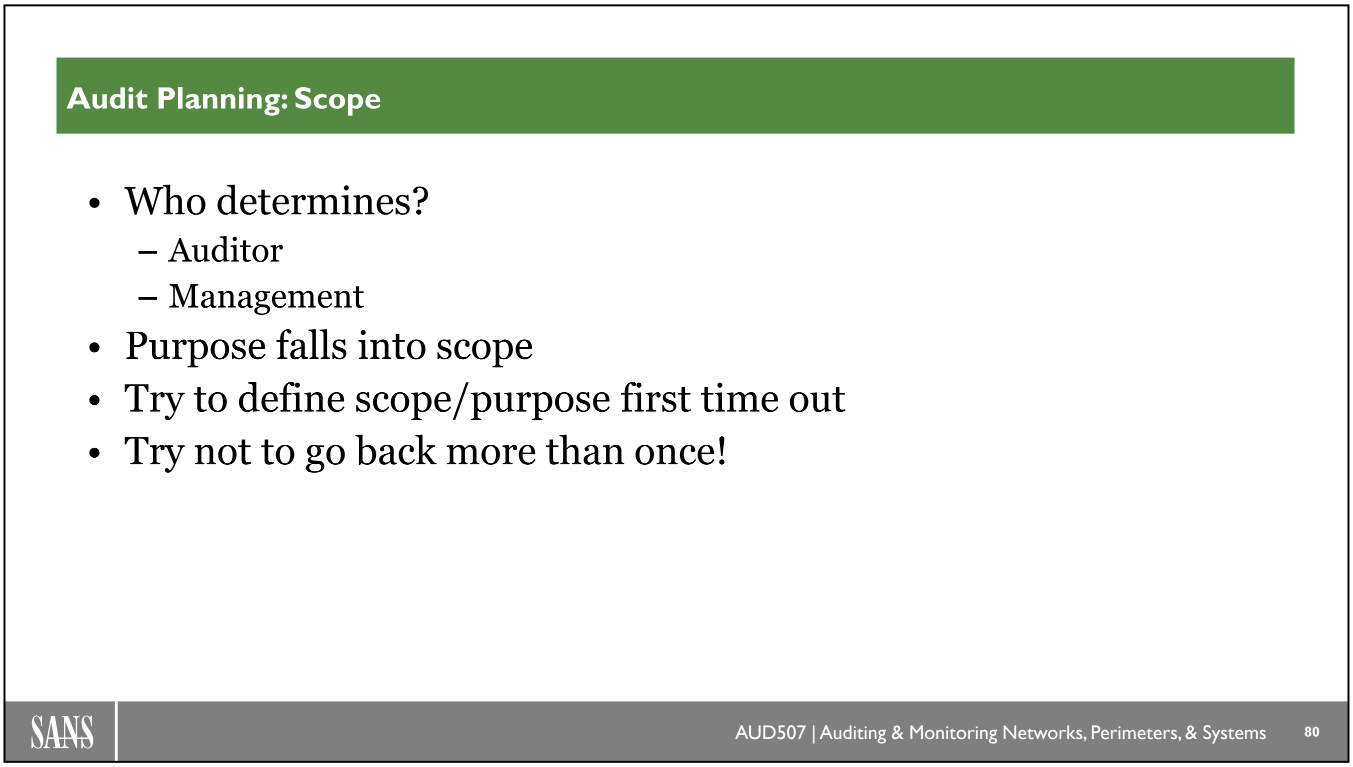

Audit Planning: Scope

- Who determines?
  - Auditor
  - Management
- Purpose falls into scope
- Try to define scope/purpose first time out
- Try not to go back more than once!

The scope of an audit is often the result of a negotiation between the organization and the auditor. Management will have some initial idea of the scope and objectives of the audit, and the auditor will work with them to establish cost-effective parameters for the audit. The auditor will refine the scope and objectives to meet the needs of the organization as efficiently and effectively as possible.

Although strong lines of communication in business are usually considered a good thing, we need to put things into the perspective of reality for just a moment. Remember that you, the auditor, will report to management regarding, quite likely, some technical stuff. Management will rely on your report to make long-term business decisions—decisions that could cost the organization a great deal of money. Imagine yourself to be the CIO of a corporation that has either hired an external auditor or tasked the internal audit team with auditing the organization's defensive posture for internet attacks. You communicate clearly (in your mind at least!) the requirements and the scope as well as the purpose of the audit/assessment. The auditor goes away to do some research and comes back explaining that there is a need to define the scope a little bit more clearly or perhaps to narrow the scope. You spend time with the auditor clarifying the scope and, after the auditor explains the issues, you agree that the scope is too broad, and you narrow it a bit. Everything is fine. Then, two days later, the auditor is back in your office looking to clarify scope again. You spend time with the auditor again, only to find him on your doorstep the next day with the same problem! How would you rate your confidence in the auditor's findings at this point?

If you need to go back, go back. But if you go back two or more times, you begin to look like the character at the bottom of the slide. This will have a serious impact on your credibility with the business.

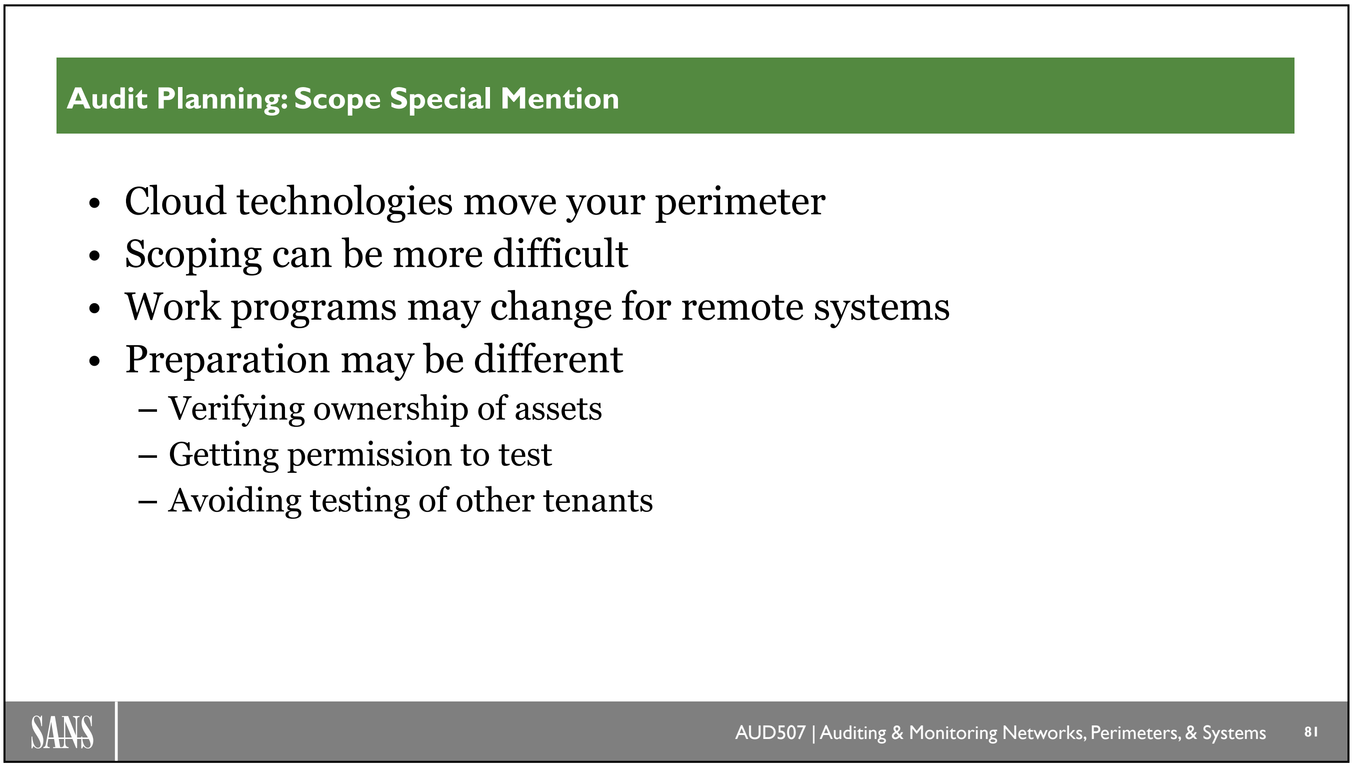

Audit Planning: Scope Special Mention

- Cloud technologies move your perimeter
- Scoping can be more difficult
- Work programs may change for remote systems
- Preparation may be different
  - Verifying ownership of assets
  - Getting permission to test
  - Avoiding testing of other tenants

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 81

Scoping an audit is relatively straightforward in an organization with a “traditional” infrastructure which has a well-defined perimeter and systems which are all located on-premises. As the enterprise adopts cloud technologies, containers, microservices, and the like, scoping becomes more complicated.

The auditor might use one work program for on-premises systems, and something entirely different for cloud-based systems. The permission required to test cloud systems will be different than locally-hosted machines. The IP address used by my company’s cloud application may be shared with other tenants who are outside my scope.

It is so important that we get these things right that we will dedicate a later portion of today’s material to discussing these issues in more detail.

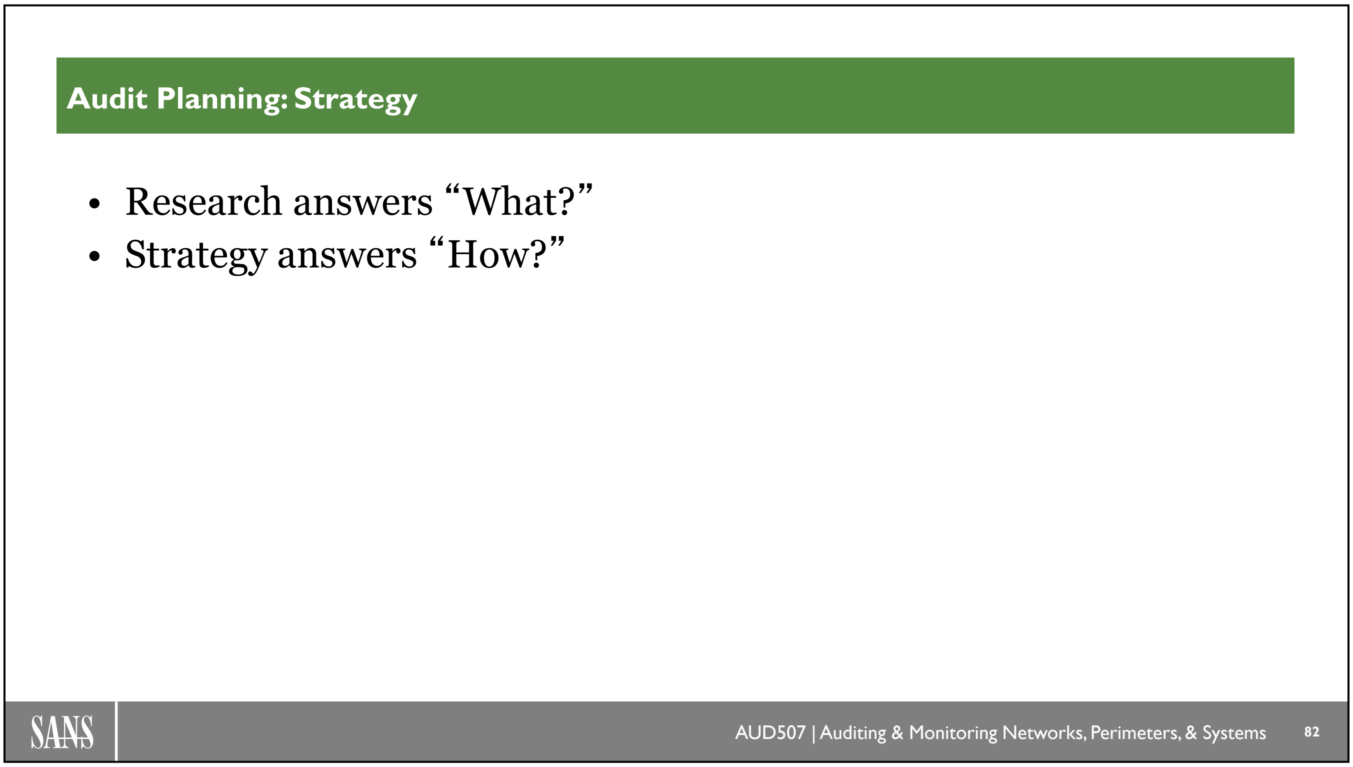

Audit Planning: Strategy

- Research answers "What?"
- Strategy answers "How?"

After you finish all the research—or sometimes, while you're still researching—you will begin to develop a feel for what strategy will work well for the audit you are attempting to perform. What are you looking for in an audit strategy? It's not necessarily the easiest to perform. What you're looking for is the strategy that will best achieve our purpose in auditing the system. If you audit how well a system conforms to corporate policy, chances are you would use a strategy of comparing the system to the approved baseline. If you want to determine how effective the corporate firewall is, a baseline of the firewall, while interesting, won't tell you how effective it is. In this case, perhaps you would use an assessment tool of some sort. Perhaps you could pull in incident reports involving network security. Perhaps you would measure the firewall based on the ruleset and how well that ruleset conforms to proven best practices. You might need to incorporate multiple strategies into your audit, and you will likely include both auditing and assessing as a part of the engagement.

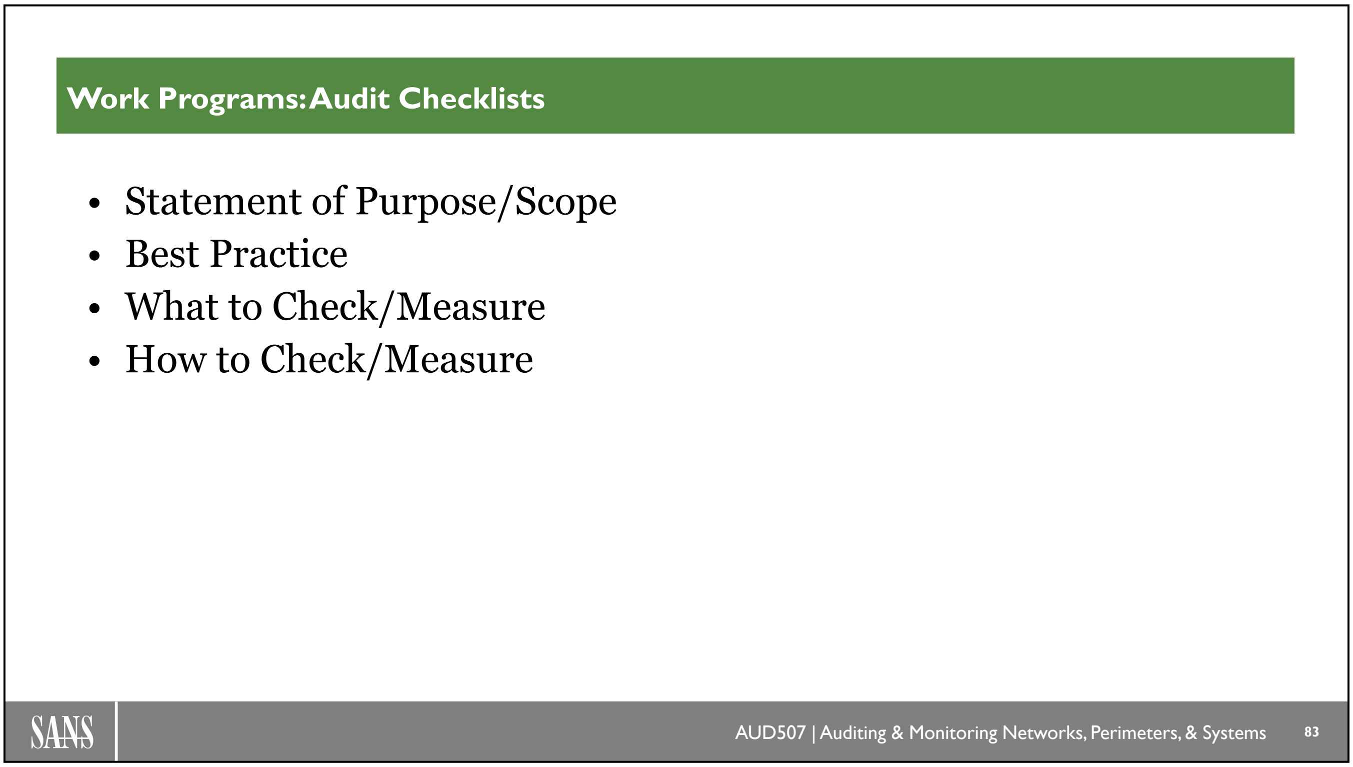

Work Programs: Audit Checklists

- Statement of Purpose/Scope
- Best Practice
- What to Check/Measure
- How to Check/Measure

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 83

One of the primary tools of an auditor is the checklist. One of the key things we plan to teach you this week is how to write a good checklist. One of the problems when assessing organizations is that the audit checklists, if they exist, have the same problems that the security policies have. They are, at times, too vague, lack scope, and do not cite industry references. We spend some time in this section of the course moving from policy statements to checklist line items, and from checklist line items to policies. One thing we do not do in class is our best practice research. Before you start writing a checklist, you should research whatever it is you are seeking to audit, assess, or secure. These research sources should be listed at the outset of your checklist to provide the framework upon which your auditing criteria are based. Why is this important? Because when someone says to you, "I don't understand why we need to turn the widgets clockwise before we connect them to the whoziwhats – in fact, I think that we should remove that from the security requirements completely," you can quickly point them to a reference work that backs this up as industry best practice. It makes the entire audit process far less subjective and can serve to allow the auditor and the sysadmin to examine the checklist together as equals, rather than the sysadmin viewing the audit checklist as the auditor's cattle prod!

Included in a good checklist will be a statement of scope, a list of what needs to be checked (and sometimes why), as well as how to check or measure compliance. Your checklist and audit procedures should be so well documented that a trained monkey can follow it—but we don't want to imply in any way that auditors are trained monkeys. Like good policy, almost anyone should be able to pick up the document, read it, and understand how to use it and why it's important.

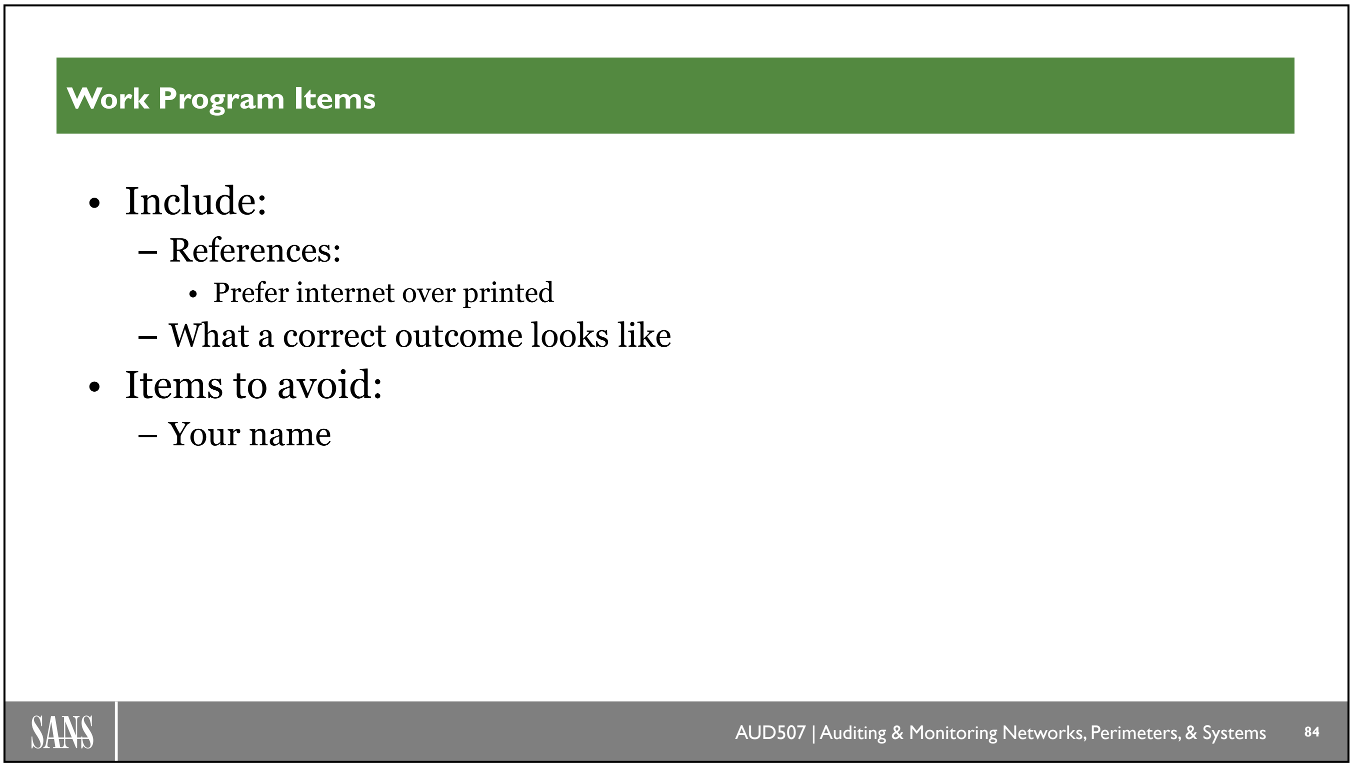

Work Program Items

- Include:
  - References:
    - Prefer internet over printed
  - What a correct outcome looks like
- Items to avoid:
  - Your name

When creating a checklist, we encourage you to include a reference for everything you are inquiring about, regardless of how detailed or high level your checklist might be. These references should be pointers to the internet or possibly corporate policy. Why the internet? You may have noticed that administrators have fewer and fewer reference volumes these days. Many administrator offices have books that are more than ten years old. Given the effectiveness of internet search engines, most administrators perform research exclusively on the internet. You may have something valuable that you've found in a reference volume, but unless you can find it on the internet as well, the administrator has no way to validate and further research the standard you are putting forth.

We also recommend that you include not just how to test but also an example of what an acceptable outcome looks like. This allows the administrator to perform the test himself without the auditor. It's also quite valuable to you as an auditor because it may be a year or more before you reuse this checklist. Will you remember what things are supposed to look like?

We strongly recommend that you never put your name on a checklist, however. Your organization's name is fine, but don't claim authorship. During your engagement with an administrator, you want him to freely express his opinion about the process, including objecting to some points on the checklist. If your name is on the checklist, it instantly creates conflict because it is now the employee challenging the auditor. If your name isn't present, it allows you to ask the employee why he objects. If the objection is reasonable, you might adjust your process. If not, you can suggest that both the auditor and the administrator review the included reference, at which point the administrator either can provide a conflicting reference (which the auditor will need to research) or will learn something that he didn't know and cooperate with the request.

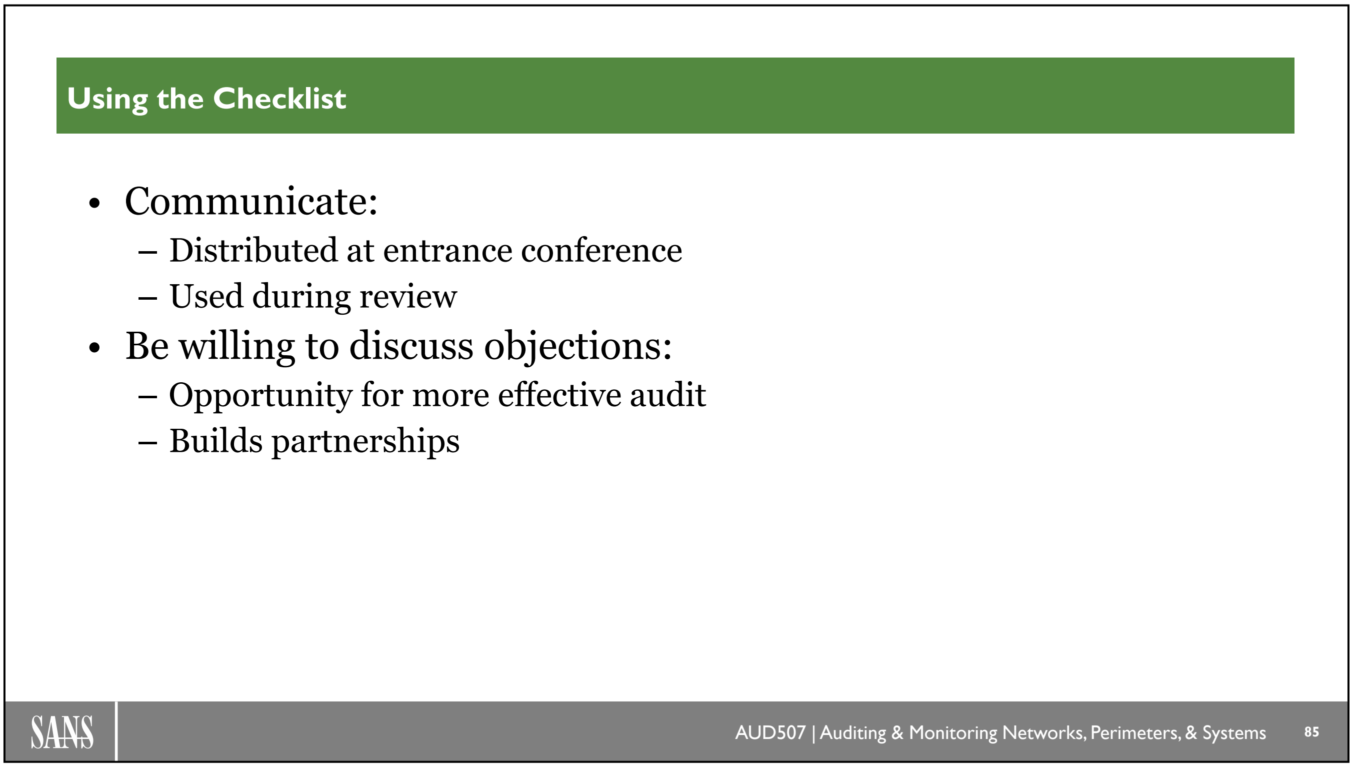

Using the Checklist

- Communicate:
  - Distributed at entrance conference
  - Used during review
- Be willing to discuss objections:
  - Opportunity for more effective audit
  - Builds partnerships

Because our goal is to measure and report on risk rather than catch people doing things wrong, it's okay for us to hand out the checklist to the people being audited at the beginning of the process. Along with the rest of the audit program documentation, the checklist acts as a working document to guide all the activities that take place during the audit.

Providing the checklist to the administrators before beginning the fieldwork is quite valuable. This can be true whether we perform a process review, where we ask questions and look for documentation, or we perform a detailed technical audit. Providing a list of questions that form the basis of the discussion for a process review allows the individuals to formulate meaningful responses and gather documents before the meeting. This preparation usually results in more focused meetings and allows us to identify any additional questions to answer. It also means that we leave the meeting with most of the validation documentation that we need rather than leaving the meeting waiting for documents. In my experience, people are far more likely to provide a large amount of documentation when *preparing* for this meeting rather than *following* this meeting. It seems to be a psychological effect. Think of it this way: When do you study? Before the test or after the test?

Providing the checklist for a technical audit during the entrance conference also allows administrators to begin to prepare for the audit, in addition to allowing them to identify any exceptions or objections they may have. Along the way, we should always be sure to inquire about initiatives that the administrators individually or collectively have been trying to move forward. It is not unusual to find that the administrators have an innate sense for issues that we may identify. If they are trying to fix an operational issue but failing to gain traction with management, and that fix satisfies a recommendation that we have, why not partner with the administrators?

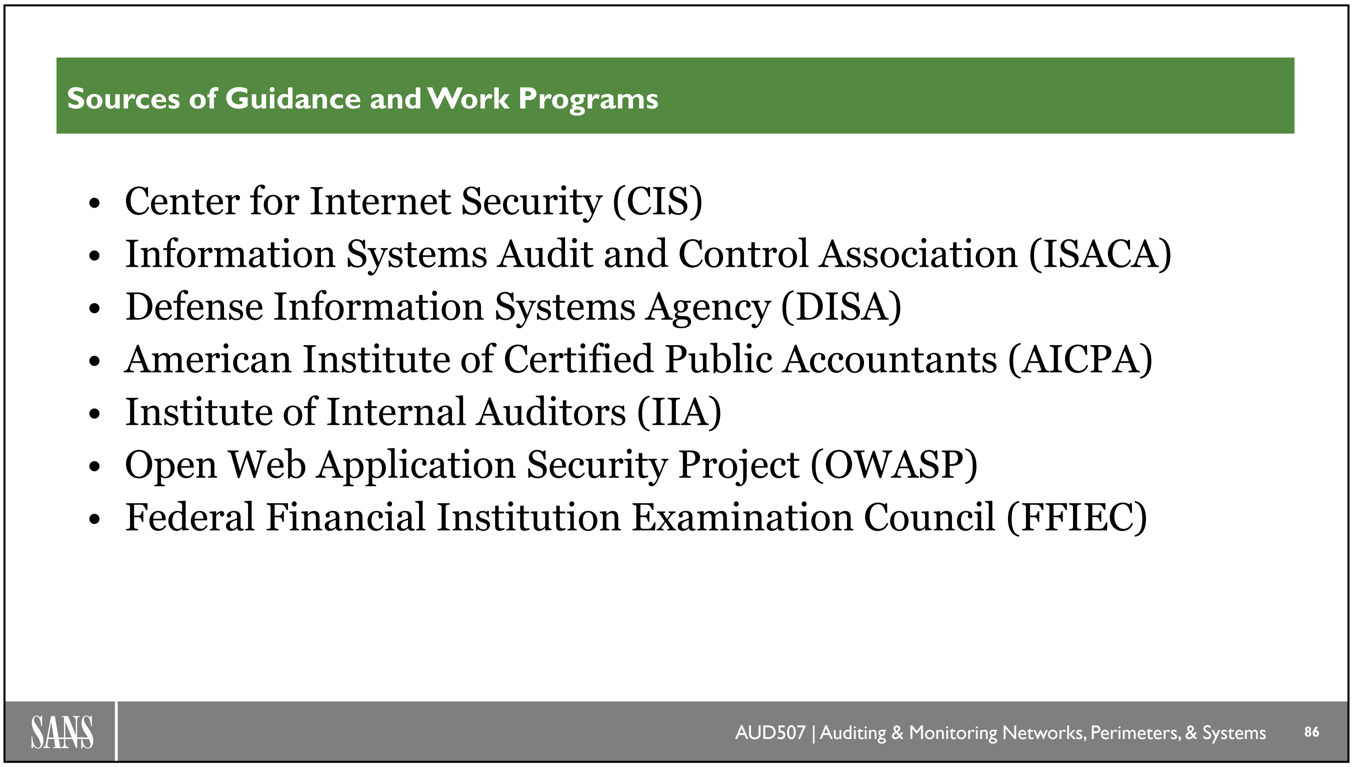

Sources of Guidance and Work Programs

- Center for Internet Security (CIS)
- Information Systems Audit and Control Association (ISACA)
- Defense Information Systems Agency (DISA)
- American Institute of Certified Public Accountants (AICPA)
- Institute of Internal Auditors (IIA)
- Open Web Application Security Project (OWASP)
- Federal Financial Institution Examination Council (FFIEC)

There are a number of good sources for work programs or checklists that are external to the organization under audit:

- (Center for Internet Security (CIS) (cisecurity.org) has industry consensus guidelines for securing many technologies.
- Information Systems Audit and Control Association (ISACA) (isaca.org) has many free and paid resources for IT auditors and security professionals.
- Defense Information Systems Agency (DISA) (disa.mil) develops the STIGS (security technical implementation guides), some of which are freely available to the public.
- American Institute of Certified Public Accountants (AICPA) (aicpa.org) publishes many of the standards used in financial and operational audits of all types.
- Institute of Internal Auditors (IIA) (theiia.org) has material for use by internal auditors, including work programs to audit your internal audit function!
- Open Web Application Security Project (OWASP) (owasp.org) produces a HUGE number or guides, control, and verification standards for securing web applications.
- Federal Financial Institution Examination Council (FFIEC) (ffiec.gov) produces the "IT Handbook" work programs for auditing IT in banks, but in general enough terms to be used by virtually anyone.

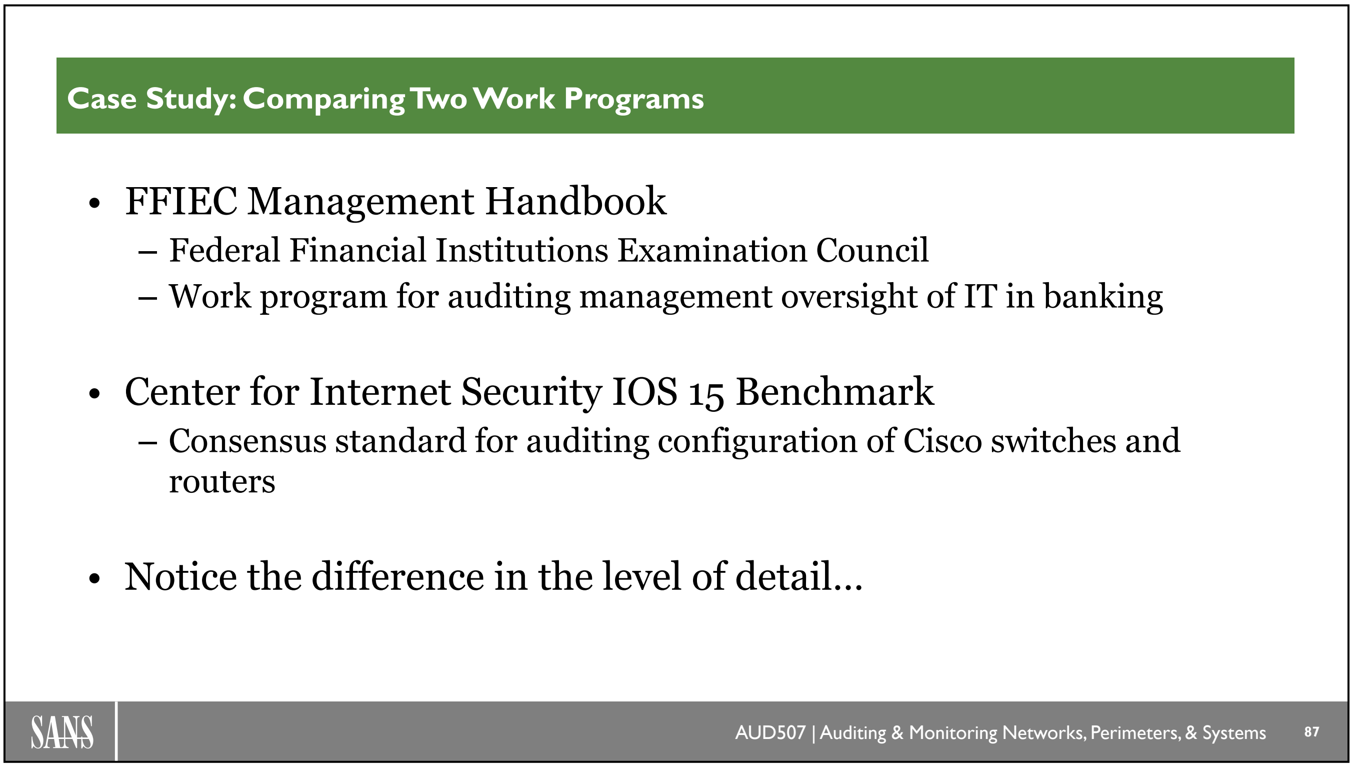

Case Study: Comparing Two Work Programs

- FFIEC Management Handbook
  - Federal Financial Institutions Examination Council
  - Work program for auditing management oversight of IT in banking
- Center for Internet Security IOS 15 Benchmark
  - Consensus standard for auditing configuration of Cisco switches and routers
- Notice the difference in the level of detail…

SANS AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 87

Let's compare two widely available work programs and discuss the differences in apparent scope and objectives, and the level of detail given in the instructions. Our first example is from the Federal Financial Institution Examination Council (FFIEC) work program on management. This checklist is part of a multi-part FFIEC handbook on managing and auditing Information Security in banks and other financial institutions. It instructs the auditor how to determine if a bank's management is doing the things it should to assist with Information Security. This and the other FFIEC work programs are available from:

https://u.aud507.com/1-2

The second example is from the Center for Internet Security's (CIS) Cisco IOS Benchmark. It gives very detailed instructions which include the commands to have an administrator run on the system and specific recommendations for correcting deficiencies. This benchmark and the others offered by CIS are available from:

https://u.aud507.com/1-3

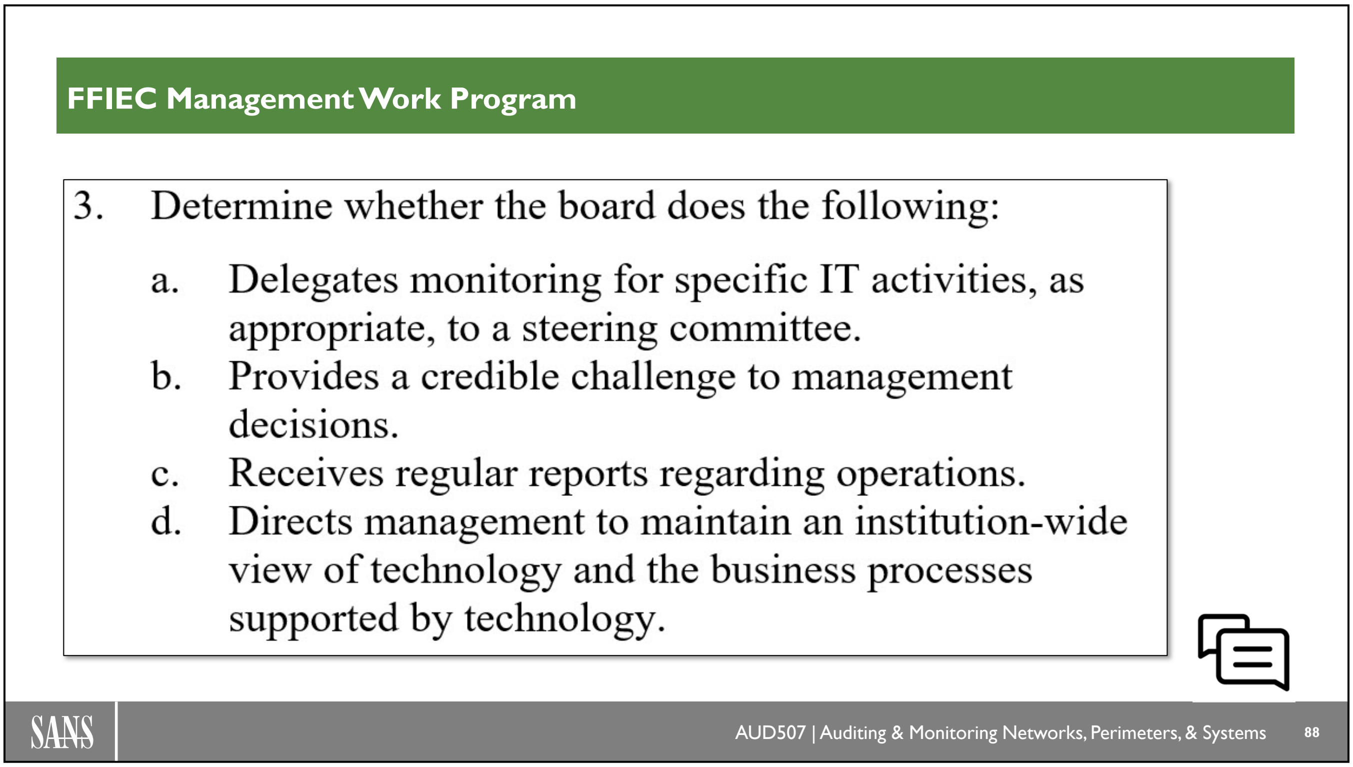

FFIEC Management Work Program

3. Determine whether the board does the following:
   a. Delegates monitoring for specific IT activities, as appropriate, to a steering committee.
   b. Provides a credible challenge to management decisions.
   c. Receives regular reports regarding operations.
   d. Directs management to maintain an institution-wide view of technology and the business processes supported by technology.

The FFIEC work program gives the auditor VERY high-level instructions about how to make the required determinations. As you read through the example on the slide, ask yourself what audit techniques you would use for this part of the audit. Would this be better handled by running an automated script or by conducting face-to-face interviews with members of management and the IT staff? Could the auditor obtain the minutes of board meetings to see whether these items are being regularly covered?

If you are taking this class live, your instructor will lead a discussion comparing this work program to the one on the next page.

CIS IOS Benchmark

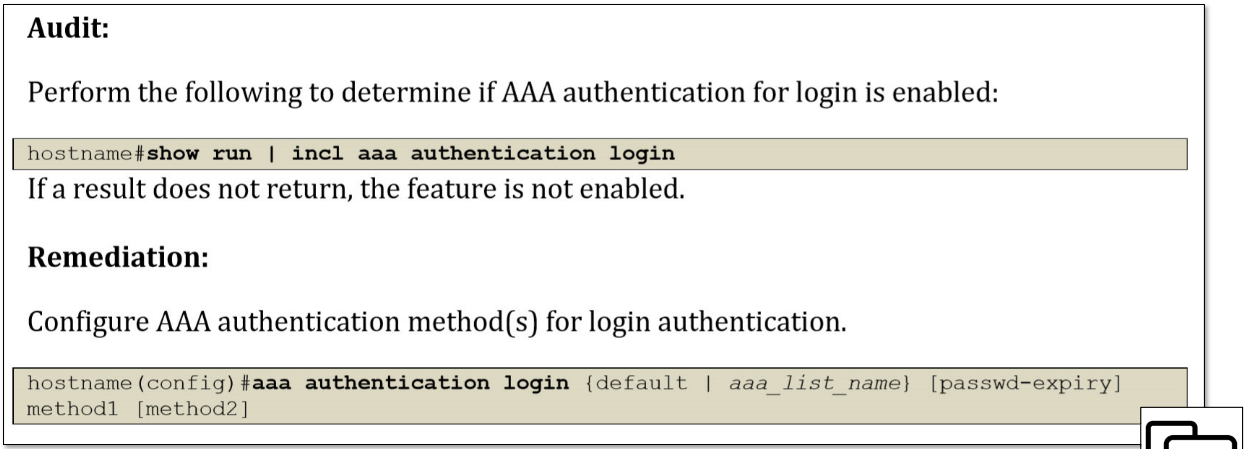

Audit:

Perform the following to determine if AAA authentication for login is enabled:

```
hostname#show run | incl aaa authentication login
```

If a result does not return, the feature is not enabled.

Remediation:

Configure AAA authentication method(s) for login authentication.

```
hostname(config)#aaa authentication login {default | aaa_list_name} [passwd-expiry]
method1 [method2]
```

In this example from the CIS Cisco IOS Benchmark, the instructions to the auditor are extremely specific. The administrator will run a specified command, and the auditor will examine the output to determine a pass or fail on this checklist item.

Would this checklist be well suited to automate with a script? Would interviews be required? How easily could this procedure be carried out by an inexperienced auditor?

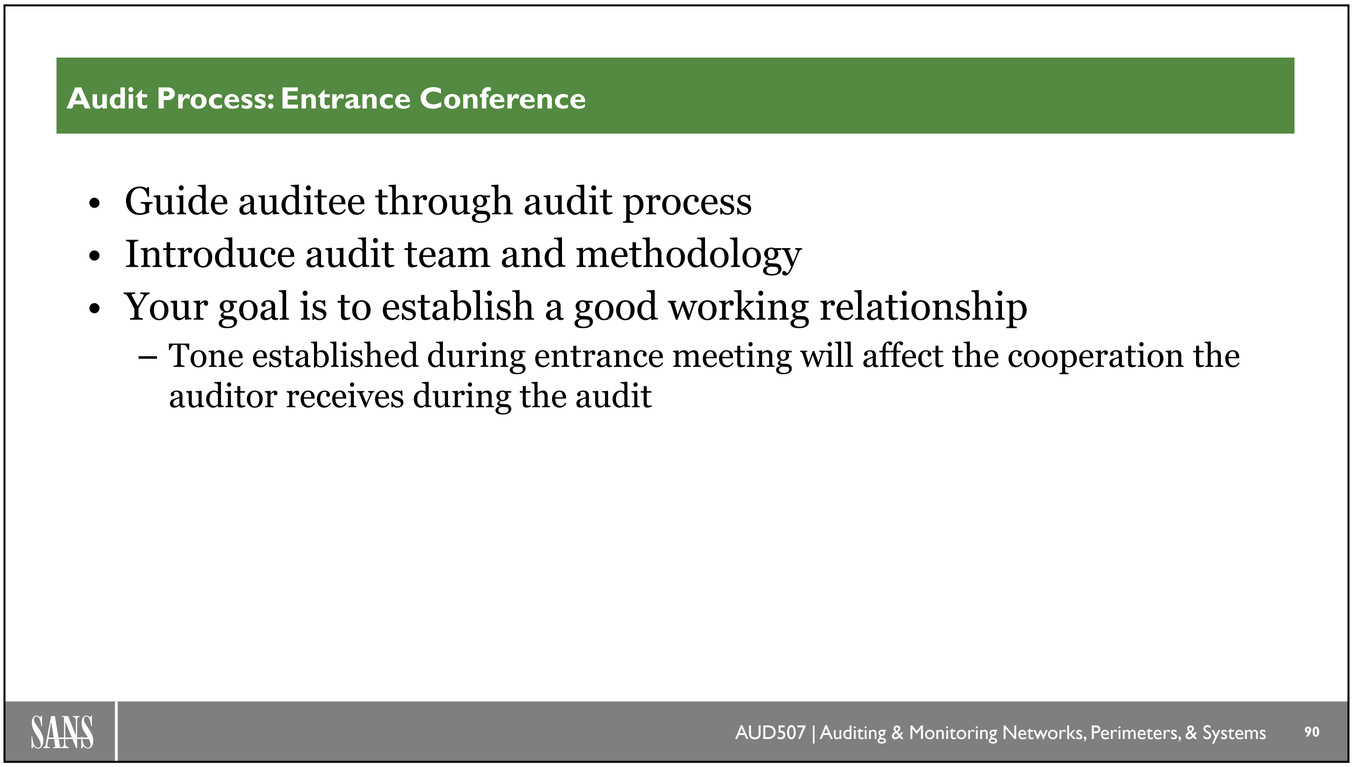

Audit Process: Entrance Conference

- Guide auditee through audit process
- Introduce audit team and methodology
- Your goal is to establish a good working relationship
  - Tone established during entrance meeting will affect the cooperation the auditor receives during the audit

We move on now to the next phase of the audit process, the entrance conference. This is the auditor's chance to establish a good working relationship with the management and staff of the auditee.

The goal of the auditor during this meeting is to explain to the auditee how the audit process will work and what they can expect. This is also the time to introduce the audit staff to the auditee's staff and establish the credibility of the audit team.

The following slides will cover who should be invited and give tips on what should be on the agenda.

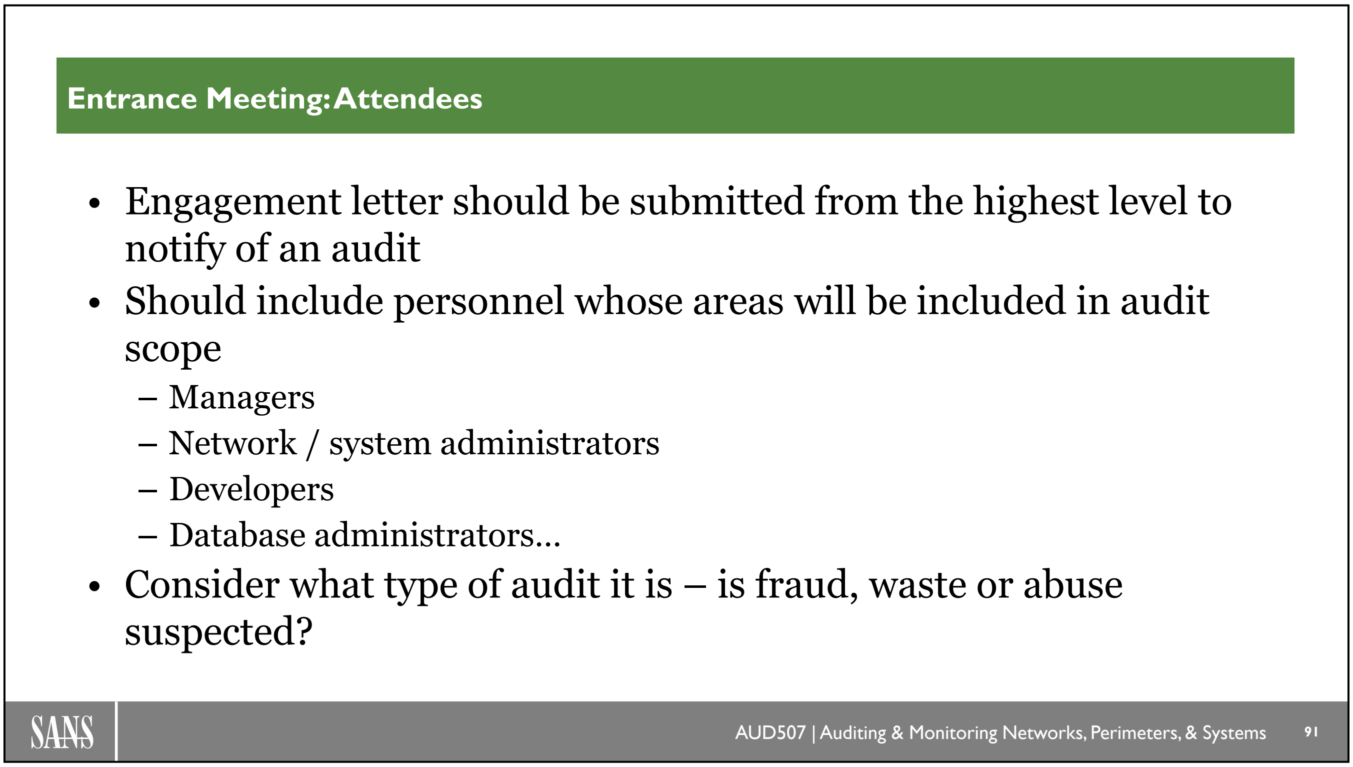

Entrance Meeting: Attendees

- Engagement letter should be submitted from the highest level to notify of an audit
- Should include personnel whose areas will be included in audit scope
  - Managers
  - Network / system administrators
  - Developers
  - Database administrators…
- Consider what type of audit it is – is fraud, waste or abuse suspected?

The engagement letter should be submitted from the highest level to notify that an audit will be taking place and to assist the auditor with gaining information for the entrance meeting. The auditor will need an introduction to begin gathering information prior to the entrance meeting, their way can be cleared by the senior level management sending out an engagement letter stating an audit will be taking place and to please provide needed assistance and cooperation as planning, execution and wrap-up takes place.

Who should be invited? Anyone who will be affected by the audit or its fieldwork can be invited. However, management may want to have their key people present and then delegate tasks. IF you have too many people present, you may not be productive. Consider the following types of individuals:

- Managers
- Network / system administrators
- Developers
- Database administrators
- Help desk
- System users
- Security professionals
- Anyone else whose area will be included in the audit scope

When considering who to invite, also consider what type of audit it is when making this decision. For example, is fraud, waste or abuse suspected?

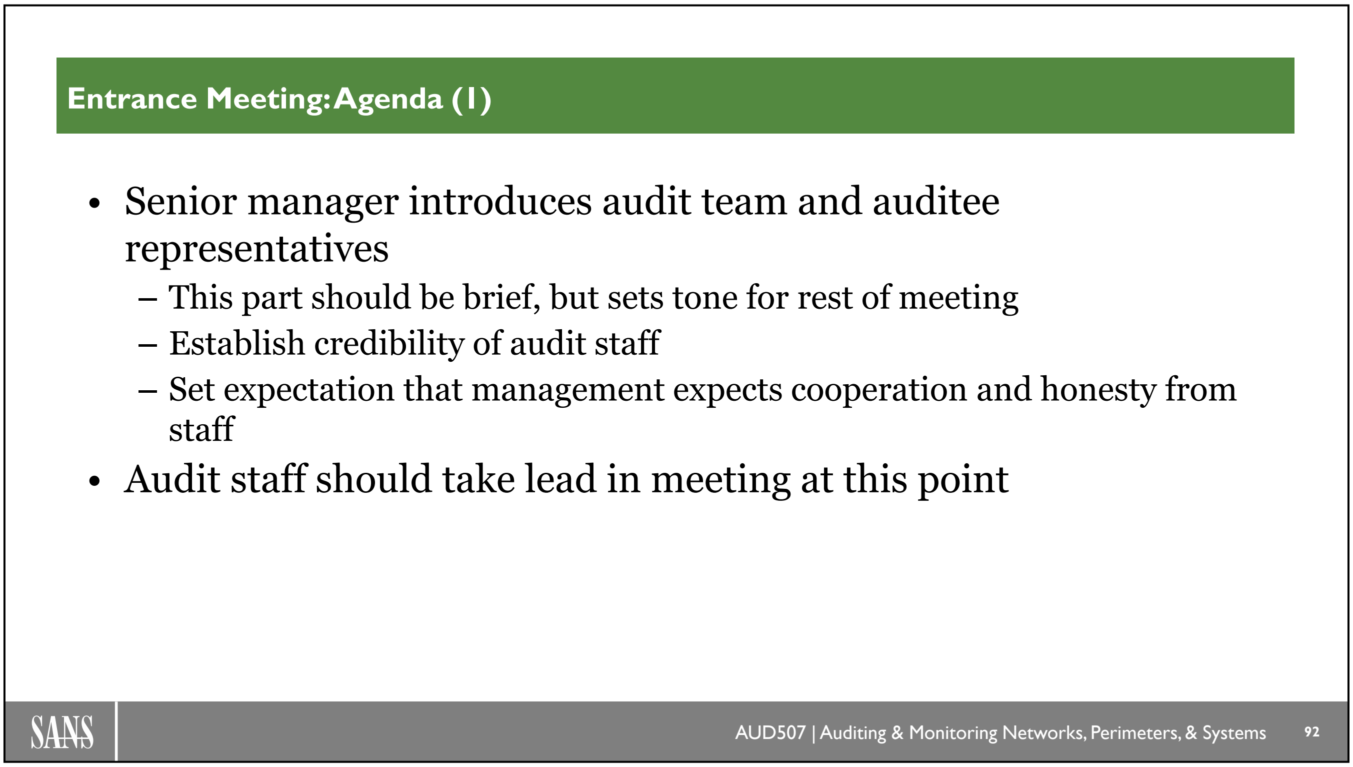

Entrance Meeting: Agenda (1)

- Senior manager introduces audit team and auditee representatives
  - This part should be brief, but sets tone for rest of meeting
  - Establish credibility of audit staff
  - Set expectation that management expects cooperation and honesty from staff
- Audit staff should take lead in meeting at this point

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 92

The senior manager who called the meeting should kick it off by introducing the audit team and explaining to staff that the auditors are all friendly, helpful people who don't want to hurt them or scare them in any way. The tone of cooperation set by the manager will go a long way toward making the staff feel more comfortable with the auditors and the audit process. The manager should pass control of the meeting to the lead auditor after the introductions and "tone-setting."

The lead auditor should be prepared to lead the rest of the meeting.

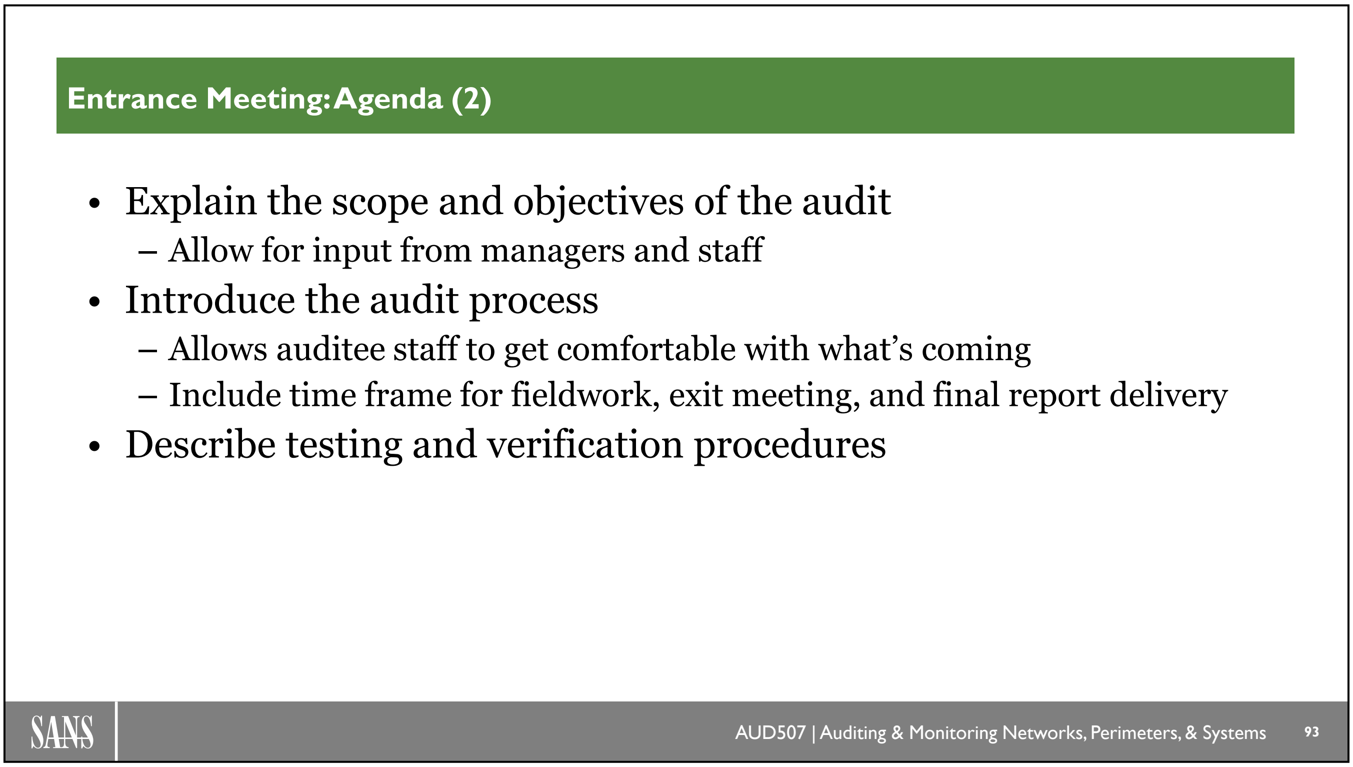

Entrance Meeting: Agenda (2)

- Explain the scope and objectives of the audit
  - Allow for input from managers and staff
- Introduce the audit process
  - Allows auditee staff to get comfortable with what's coming
  - Include time frame for fieldwork, exit meeting, and final report delivery
- Describe testing and verification procedures

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 93

The lead auditor should explain the objectives and scope which have been set for the audit. It's common that managers and staff from the in-scope organizations will have input or questions regarding the scope. While being ever mindful of scope creep, the audit staff can use their input to further refine the established scope.

Introduce the audit process and the audit work program. You may or may not want to distribute a copy of the audit work program. If you do, this helps to eliminate some of the mystery surrounding the audit process and may allow the staff to gather needed documentation in a timelier manner. However, there are times when the audit work program may be proprietary, the auditee may manipulate the results, or the auditor may spend more time trying to explain the audit work program.

The auditors should also detail the timeline for conducting fieldwork, the schedule for the exit meeting, and at least a tentative schedule for delivery of the final audit report to management.

When specific testing procedures will be used, let the auditee staff know what these are. This will also make them more comfortable, as they will know what to expect when you show up at their office.

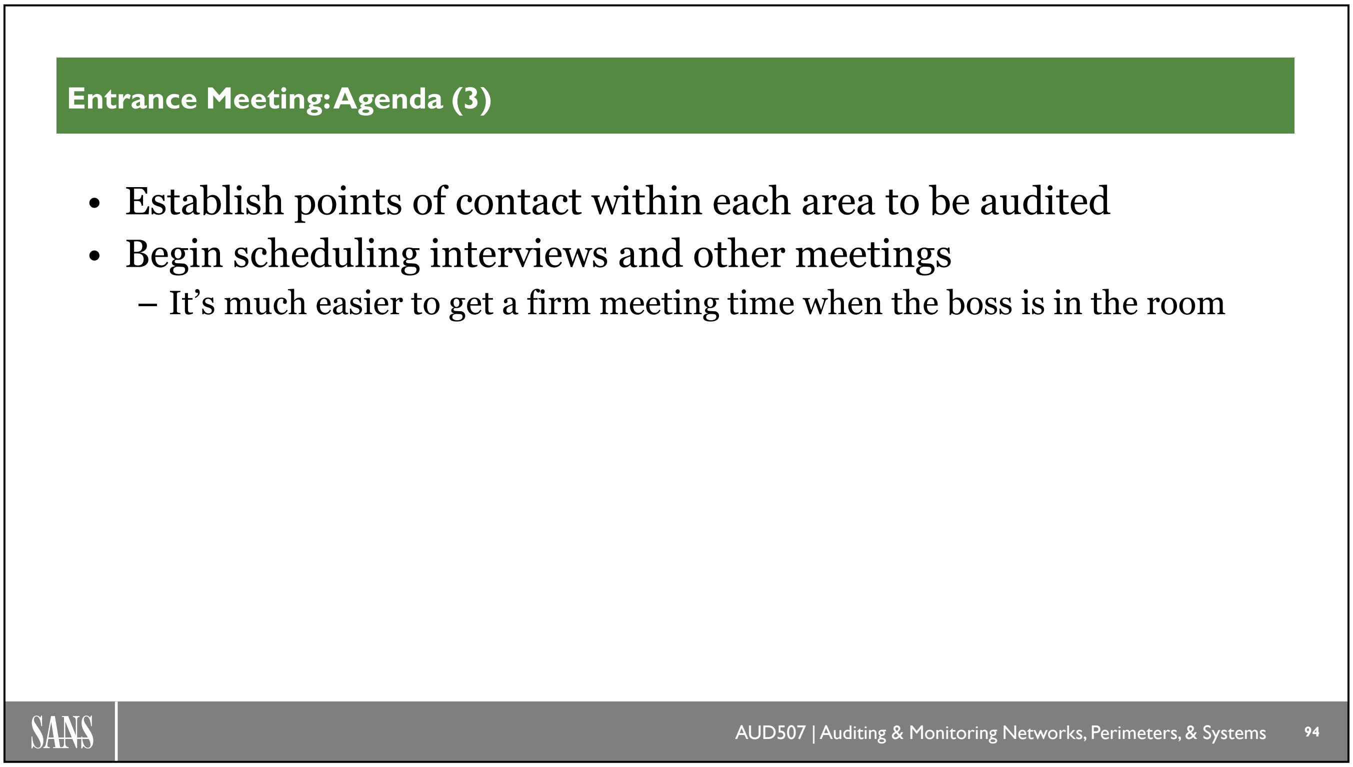

Entrance Meeting: Agenda (3)

- Establish points of contact within each area to be audited
- Begin scheduling interviews and other meetings
  - It's much easier to get a firm meeting time when the boss is in the room

The entrance meeting should be used as a time to establish points of contact within each area included in the audit scope. The auditors should establish whether they have the freedom to schedule meetings directly with staff, or if they should go through a manager or administrative assistant to do so.

It's a very good idea to begin scheduling staff interviews and other meetings during the entrance meeting. You will be surprised how cooperative people are when their boss is still in the room. Service-oriented departments, including the help desk, network infrastructure and software developers tend to hate meetings and are masters of being too busy when you need to talk to them. The audit staff should probably focus on scheduling time with these people during the entrance meeting.

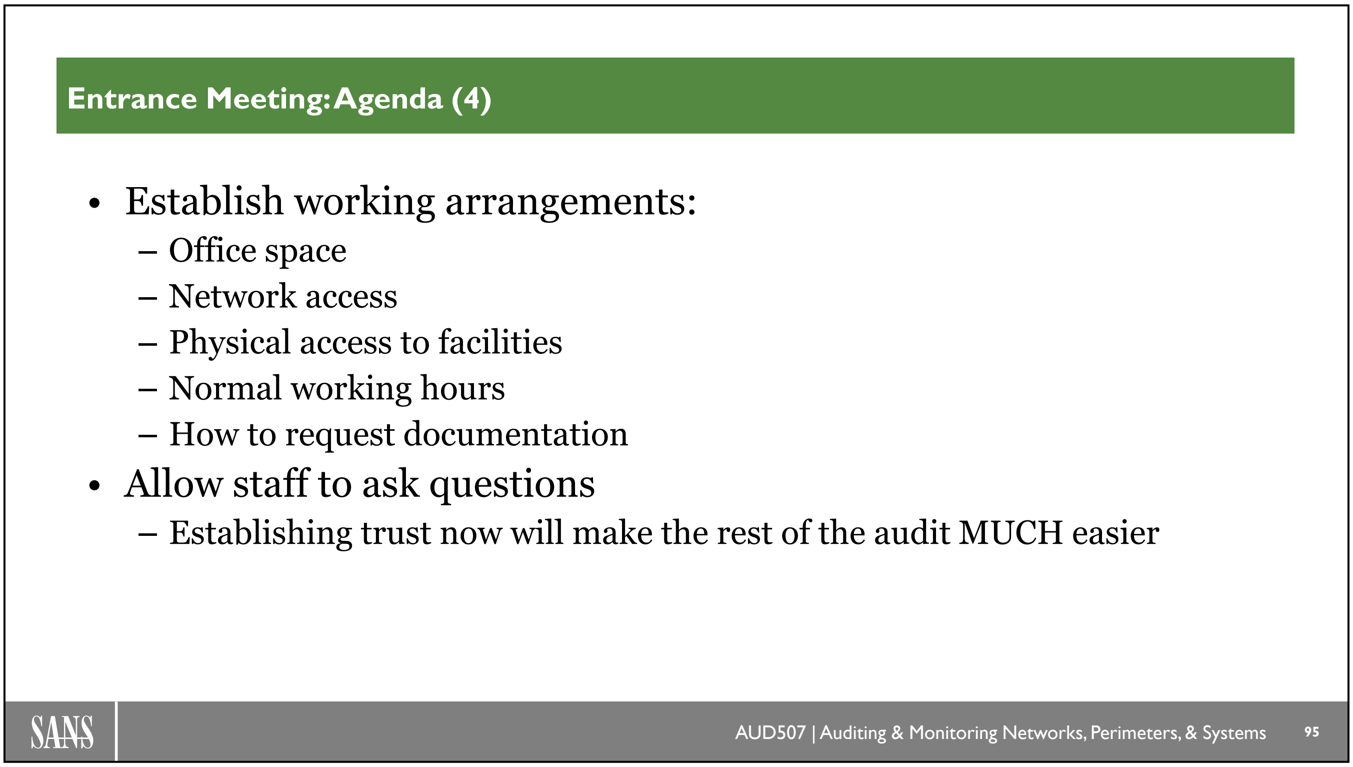

Entrance Meeting: Agenda (4)

- Establish working arrangements:
  - Office space
  - Network access
  - Physical access to facilities
  - Normal working hours
  - How to request documentation
- Allow staff to ask questions
  - Establishing trust now will make the rest of the audit MUCH easier

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 95

Finally, the entrance meeting is a good time to work out logistics concerns for the fieldwork portion of the audit. Items to discuss can include:

- Office or working space for the auditors
- Network access
- Physical access to facilities. Are identification cards required? Will the auditors require an escort?
- Parking provisions for audit staff
- Normal working hours. During what hours should meetings or testing be scheduled?
- How to request documentation

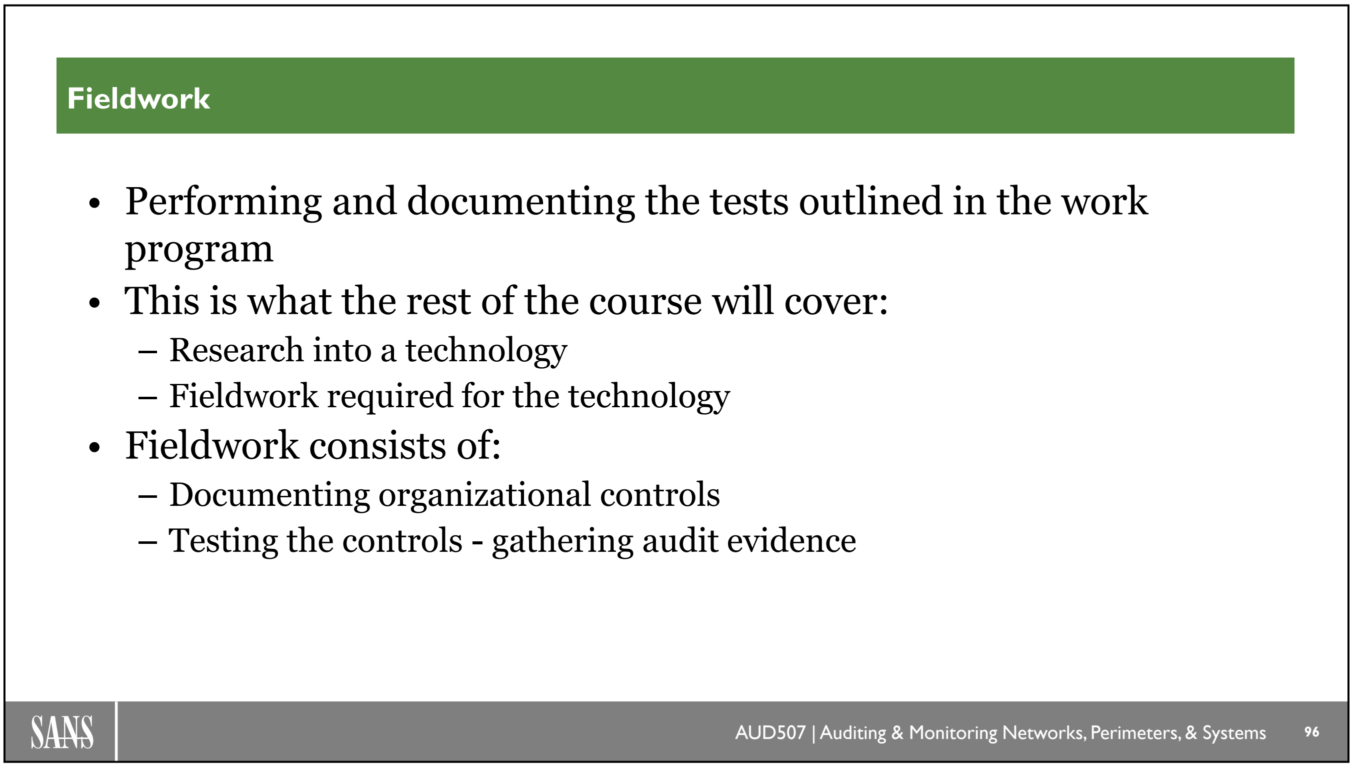

Fieldwork

- Performing and documenting the tests outlined in the work program
- This is what the rest of the course will cover:
  - Research into a technology
  - Fieldwork required for the technology
- Fieldwork consists of:
  - Documenting organizational controls
  - Testing the controls - gathering audit evidence

Here's the bulk of the course material in two simple instructions:

- Identify and document the control environment under review, and then
- Test those controls

Most of the course will revolve around how to do these two things for various technologies and control types.

The auditor should identify the controls in use in the organization which are covered by the audit scope and objective. These should include any administrative controls, such as policies, standards, or procedures, as well as the technical controls associated with the subject of the audit.

During the second part of fieldwork, the auditor will gather evidence to test the effectiveness of the identified controls. The evidence may come in many forms, such as:

- Output from a technical audit tool
- Text file containing device configuration
- Logs generated by detective controls
- Reports generated during application of a control

It is critical that all audit evidence be maintained during the audit. If a finding is called into question, the auditors working papers will be needed to defend the auditor's conclusions.

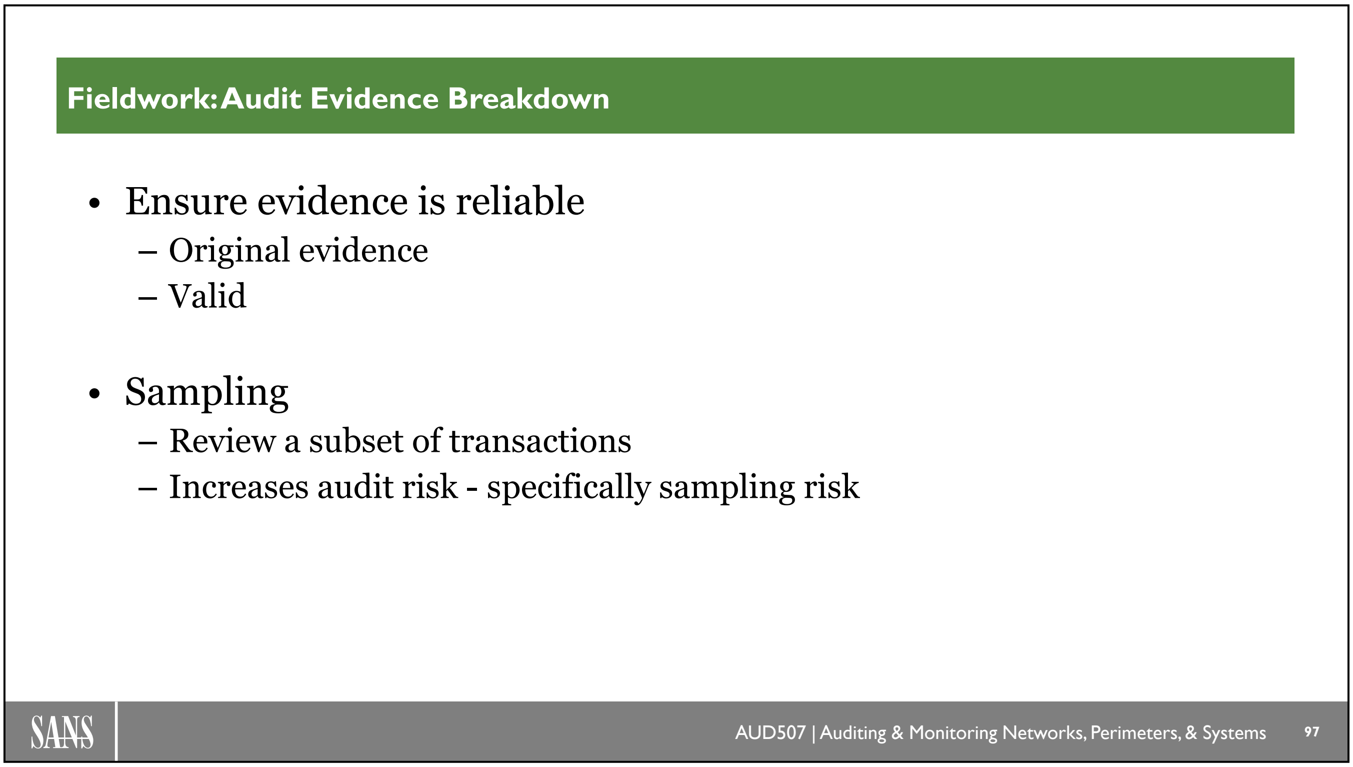

Fieldwork: Audit Evidence Breakdown

- Ensure evidence is reliable
  - Original evidence
  - Valid
- Sampling
  - Review a subset of transactions
  - Increases audit risk - specifically sampling risk

The auditor needs to make sure the evidence has not been "pre-audited" or tampered with, and that the evidence is valid. The auditor must confirm the validity of any evidence or sources.

Sometimes it is not feasible for the auditor to test every instance of a control in use. For example, it may not be possible to review all 1,000,000 transactions processed by a pharmacy during a year to ensure that patient information was properly protected in every case. In these circumstances, the auditor may perform tests on a sample of the transactions. This allows the auditor to complete the work within a reasonable amount of time but increases the risk that the auditor may not detect a material problem in the system.

If an error is found, an auditor needs to determine if they should call this a finding or do further sampling to determine how prevalent the problem is in the system. It may mean the auditor needs to do further investigation to identify the root cause.

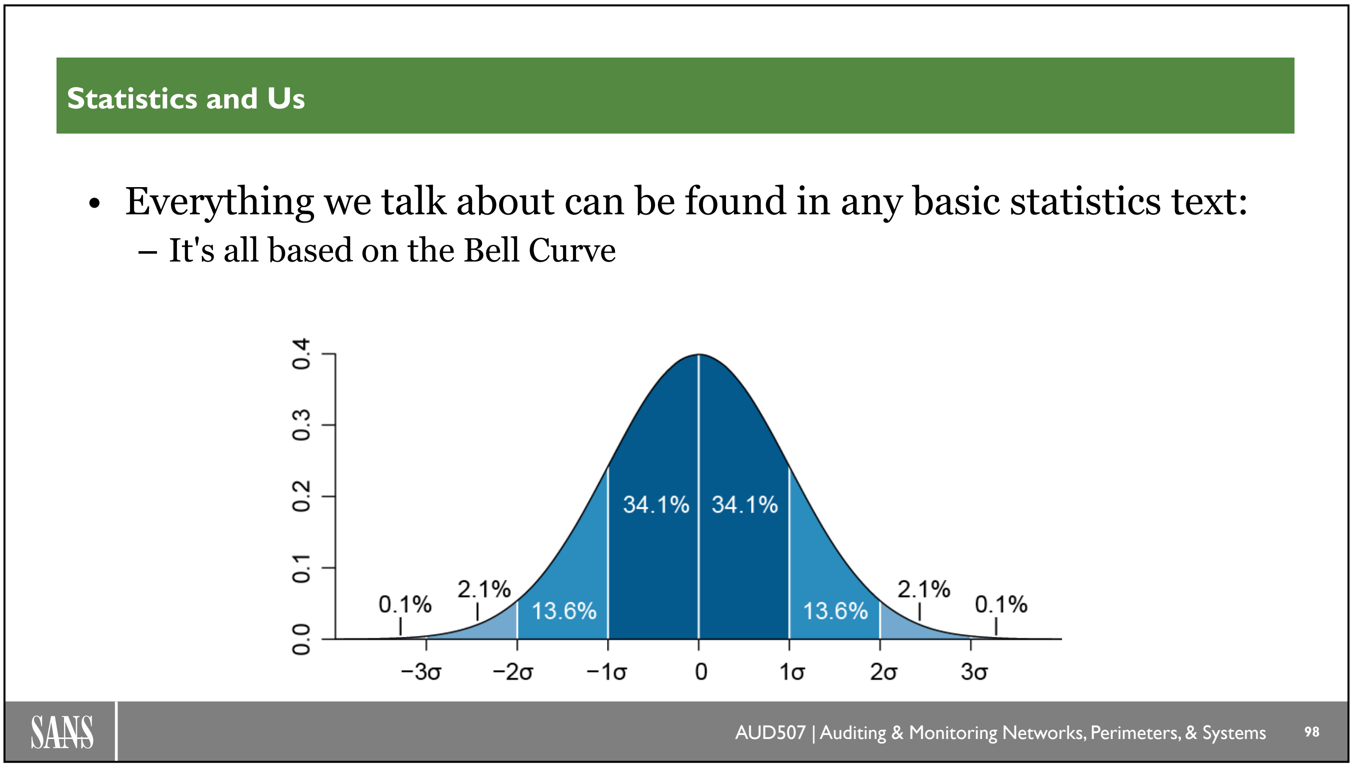

Statistics and Us

- Everything we talk about can be found in any basic statistics text:
  - It's all based on the Bell Curve

Before we go any further, let me assure you that this isn't about to turn into a mathematics or statistics class! However, there are some useful mathematical principles we can draw upon from the field of statistics. Rather than simply guessing what an appropriate sample size would be, we would like to have a repeatable, justifiable method for the selection of a sample size. Statistics provides us with just such a mechanism.

Everything that we discuss that involves statistics is based on the fundamental structure of a great deal of probability theory and basic statistics: The Gaussian Distribution. This is also known as the Normal Curve or, more colloquially, the Bell Curve. Within the distribution, the majority of all values represented by an average population falls within one standard deviation (for our purposes, the average variance from the average of the entire data set) of the average or normal. In this graph, you can see that 68.2% falls within one sigma (or standard deviation) of the normal. The further away from the center of the graph, the more unusual (statistically) the value. So how do we use this? Let's take a look at using it as the basis for calculating sample size and margin of error based on a sample.

Illustration attribution: *"Standard deviation diagram" by Mwtoews. Own work, based (in concept) on figure by Jeremy Kemp, on 2005-02-09. Licensed under CC BY 2.5 via Wikimedia Commons, https://commons.wikimedia.org/wiki/File:Standard_deviation_diagram.svg#/media/File:Standard_deviation_diagram.svg.*

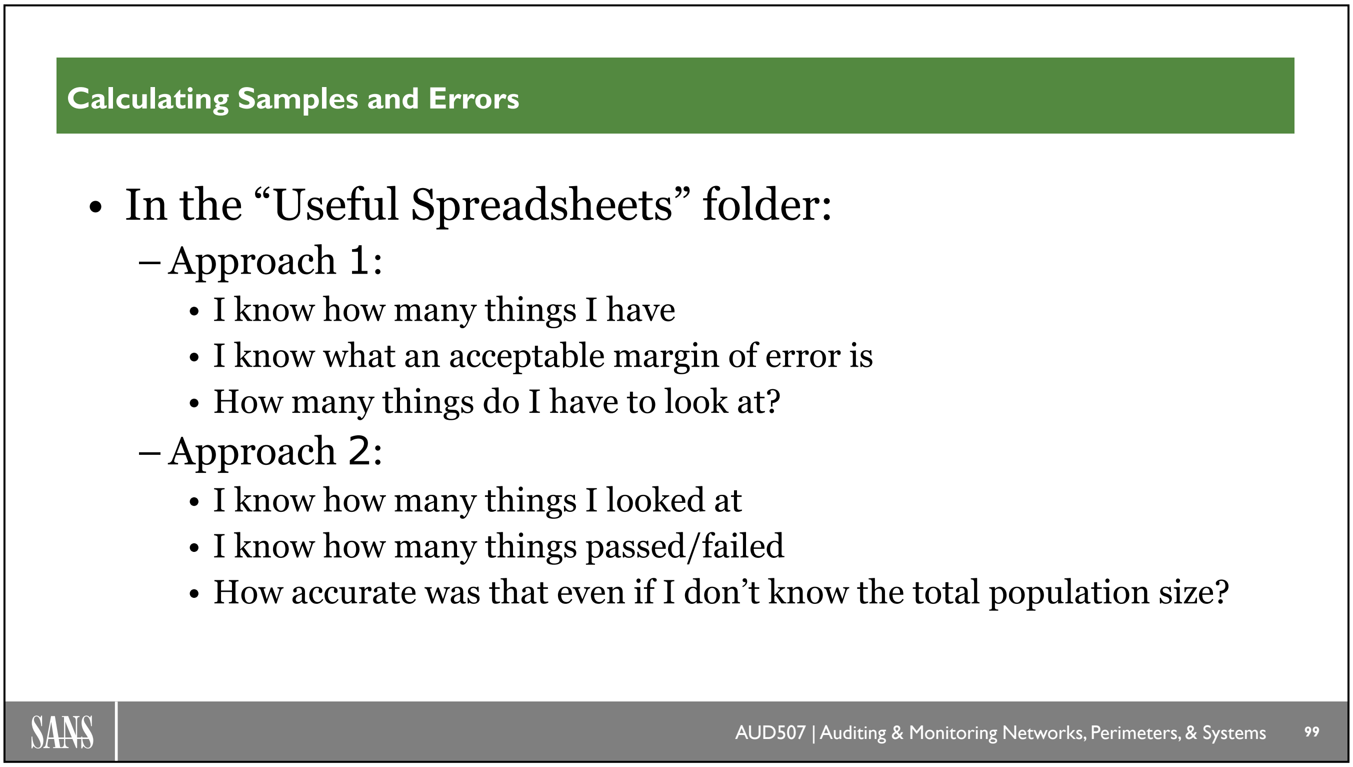

Calculating Samples and Errors

- In the “Useful Spreadsheets” folder:
  - Approach 1:
    - I know how many things I have
    - I know what an acceptable margin of error is
    - How many things do I have to look at?
  - Approach 2:
    - I know how many things I looked at
    - I know how many things passed/failed
    - How accurate was that even if I don’t know the total population size?

SANS AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 99

Rather than work our way through tedious mathematics, we have provided a spreadsheet for you in the Useful Spreadsheets folder, which is broken into two sections.

The top section enables you to take an arbitrary population size, define a margin of error, and automatically calculate how many “things” you would need in your sample to produce results at that accuracy level. This particular piece of the spreadsheet is designed to be used iteratively (more on that soon).

The second part of the spreadsheet enables you to take results that you already have and figure out how accurate they are, even if you don’t know exactly how many “things” there are in total. This is a margin-of-error calculation.

Let's look at each of these in more depth.

Calculate Sample Size

- Fill in the orange boxes:
  - Population size
  - Expected occurrence
  - Margin of error

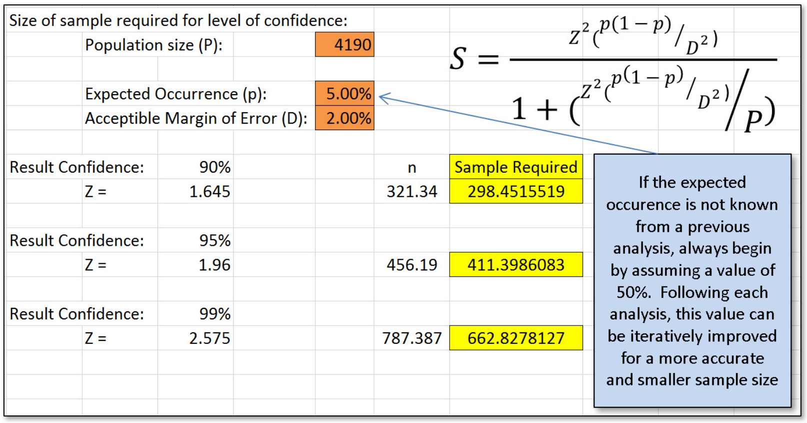

Size of sample required for level of confidence:				
Population size (P):		4190		
Expected Occurrence (p):		5.00%		
Acceptible Margin of Error (D):		2.00%		
Result Confidence:	90%		n	Sample Required
Z =	1.645		321.34	298.4515519
Result Confidence:	95%			
Z =	1.96		456.19	411.3986083
Result Confidence:	99%			
Z =	2.575		787.387	662.8278127

$$S = \frac{Z^2 \left(\frac{p(1-p)}{D^2} \right)}{1 + \left(\frac{Z^2 \left(\frac{p(1-p)}{D^2} \right)}{P} \right)}$$

If the expected occurence is not known from a previous analysis, always begin by assuming a value of 50%. Following each analysis, this value can be iteratively improved for a more accurate and smaller sample size

The margin of error enables you to define the amount of "wiggle room" you can tolerate in your results. To achieve a 0% margin of error, you would have to look at every single element in the population. The margin of error and confidence interval are used to express your overall confidence that the sample accurately depicts the entire population.

Imagine a sample where you decided to use a 95% confidence interval (the most common) with a 2% margin of error. You found that 8% of the systems sampled were not compliant with the tested policy. You could state in your report that, "based on a statistical sample, we are 95% confident that the percentage of systems which are out of compliance falls within the range from 6% to 10% of systems.

In other words, you are 95% sure that your results accurately reflect what is happening in the entire population, even though you are looking only at a subset. Of course, higher confidence intervals require that you examine more items within the total population size.

It sounds complicated, but it isn't. The only field that's a little tricky is the expected occurrence field. This formula is designed to be used iteratively. This means that as you analyze a sample and come up with results, you can feed the data back into the spreadsheet to refine the selection of samples for future audits. When you use the spreadsheet for the first time, and you have no data on which to base a hypothesis of how many systems are or are not in compliance, you must assume a pass/fail rate of 50%, which is a worst-case assumption that will yield the largest sample size.

Typically, for a first pass, you can assume 50% with a margin of error somewhere between 5% and 10%. After conducting that test, use the results to refine your sample, also reducing your margin of error. Such an approach enables you to answer questions about large numbers of systems while only analyzing a small portion of them!

Margin of Error

- Size of Sample
- Number Matching

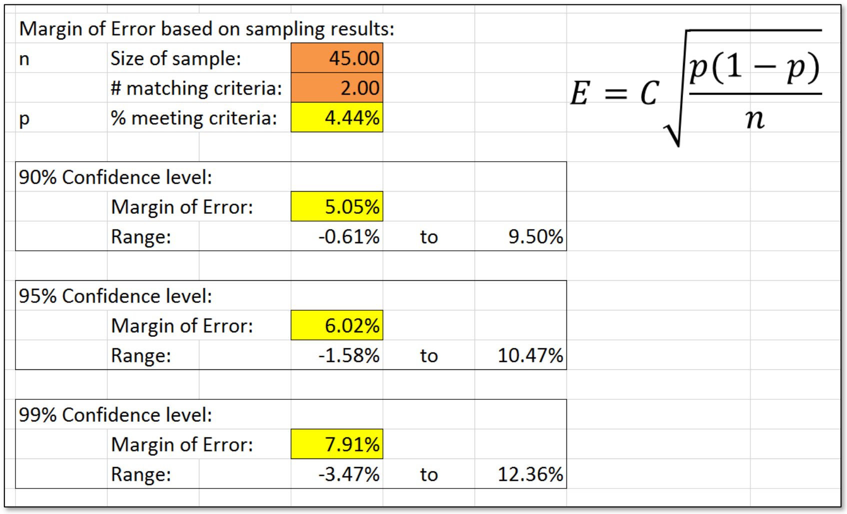

Margin of Error based on sampling results:				
n	Size of sample:	45.00		
	# matching criteria:	2.00		
p	% meeting criteria:	4.44%		
90% Confidence level:				
	Margin of Error:	5.05%		
	Range:	-0.61%	to	9.50%
95% Confidence level:				
	Margin of Error:	6.02%		
	Range:	-1.58%	to	10.47%
99% Confidence level:				
	Margin of Error:	7.91%		
	Range:	-3.47%	to	12.36%

$$E = C\sqrt{\frac{p(1-p)}{n}}$$

The second portion of the spreadsheet is much easier to understand. In this case, we have an arbitrary number of items that have been analyzed, and we'd like to know how accurately that will reflect an entire population. The trick, though, is that we don't know how big the entire population is!

All that we need to do is to fill in the orange boxes. We enter the number of things we looked at and the number that passed or failed a particular test. Below, we see what the percentage is, and then we can see the margin of error involved at different confidence intervals. Again, the confidence interval is the level of assurance you have that those results are accurate. This is why you see that the size of the margin of error increases dramatically as you move toward the 99% level of confidence.

To wrap up, you should use sampling. Rather than arbitrarily choosing sample sizes, however, you should have some statistical foundation for your choices. More than that, when you produce results, it can be extremely valuable to couple those results with statistics that show current compliance, how pervasive a problem is, and any improvement or decay over time.

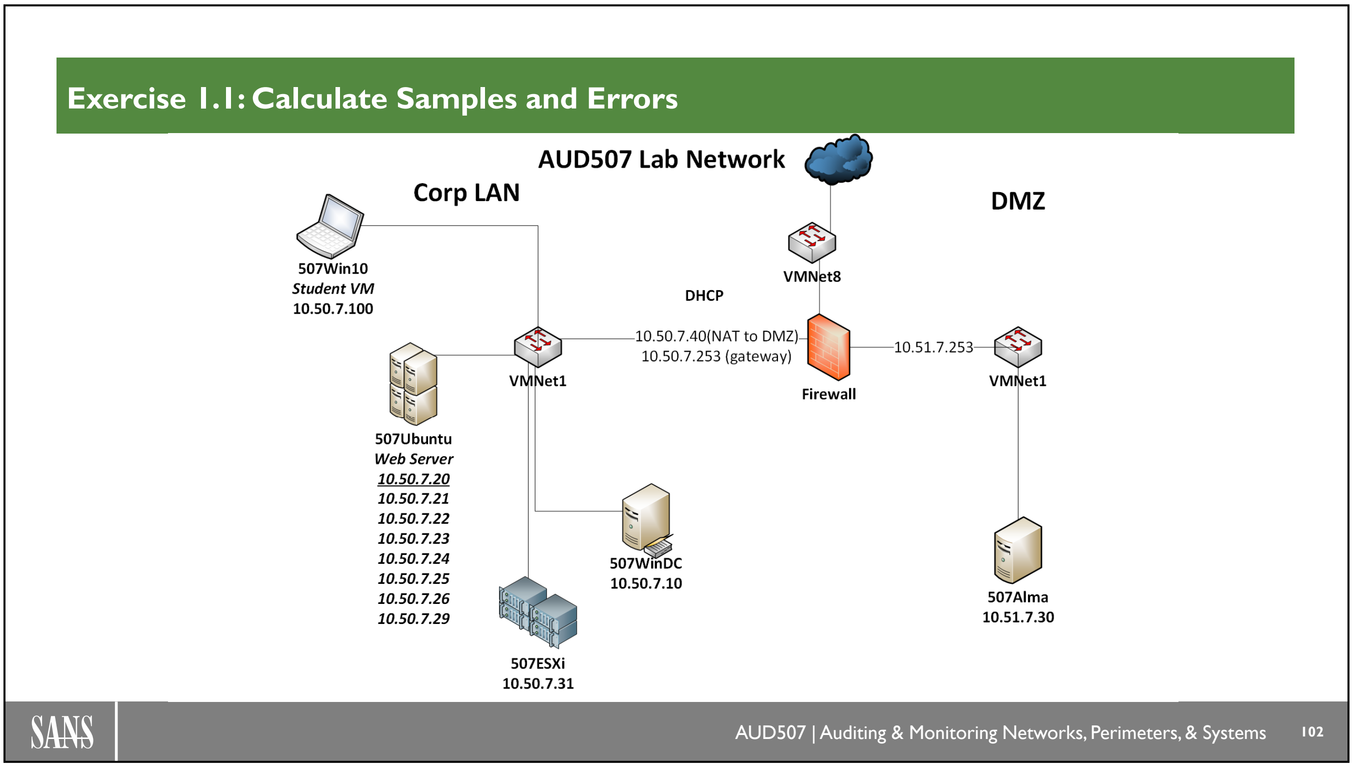

Exercise 1.1: Calculate Samples and Errors

Now turn to the "Exercise 1.1: Calculating Samples and Errors" exercise in your workbook. If you are attending a live conference or a vLive, your instructor will likely do this lab interactively with the class!

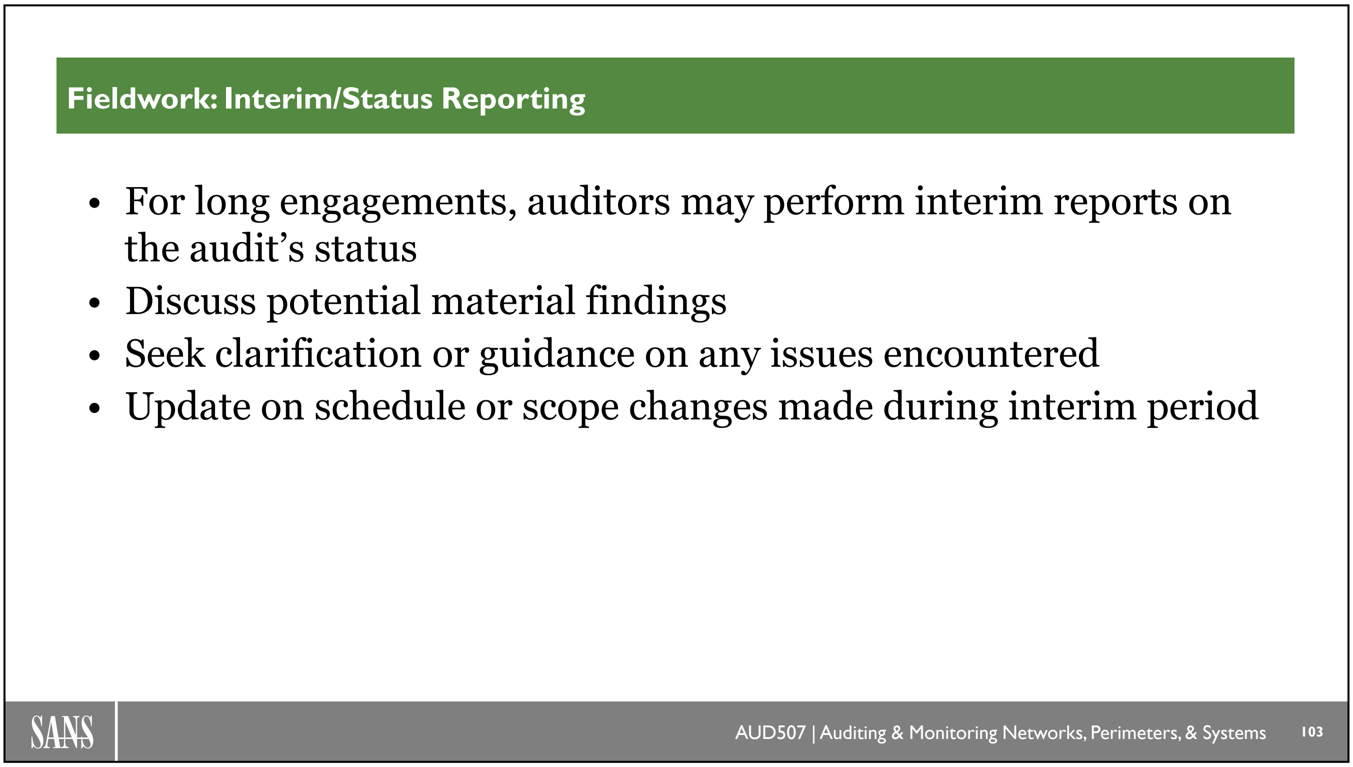

Fieldwork: Interim/Status Reporting

- For long engagements, auditors may perform interim reports on the audit's status
- Discuss potential material findings
- Seek clarification or guidance on any issues encountered
- Update on schedule or scope changes made during interim period

For very long audit engagements, those lasting months, rather than days or weeks, the lead auditor may opt to hold regular status updates with the auditee and/or the executive who is sponsoring the audit. This gives the auditor the chance to share critical information gathered during the audit process and to check that their understanding of potentially material issues is correct before drafting the report.

In some ways, these meetings can serve as a mini exit conference (more about that shortly), allowing the auditor to short-circuit misunderstandings and issues before they cause mistakes that could make it into the final report.

If you are taking the course live, your instructor will likely simulate a status meeting at the end of each class day. If you are taking the source via OnDemand, you can consider the issues raised on the slide at the end of each course day's material.

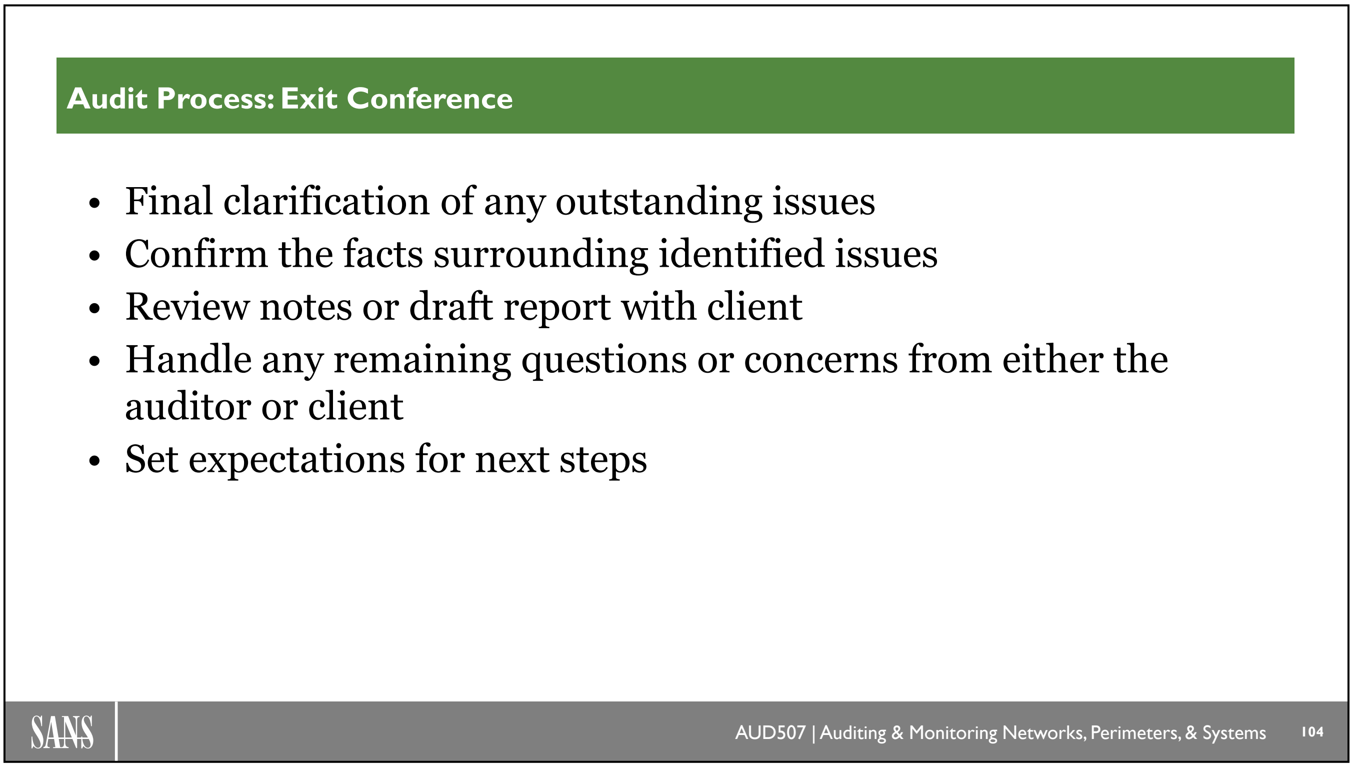

Audit Process: Exit Conference

- Final clarification of any outstanding issues
- Confirm the facts surrounding identified issues
- Review notes or draft report with client
- Handle any remaining questions or concerns from either the auditor or client
- Set expectations for next steps

The exit meeting is used to wrap-up the fieldwork portion of the audit and transition into the reporting phase. During this meeting, the auditors present management with either a draft of the audit report (preferred) or their notes on significant findings (much less preferred).

The goal of the meeting is to ensure that:

- The findings in the final report do not come as a surprise to the auditee
- The auditors are correct in their understanding of relevant facts surrounding any issues identified
- Any remaining questions or concerns from either the auditors or auditee are addressed before the final report is delivered
- Set expectations for next steps. During some audits, the process is followed up by getting management responses to the findings, providing additional information, the auditor issuing the final report, and a follow-up audit. Information about and timing for all of these items can be communicated during the exit meeting.

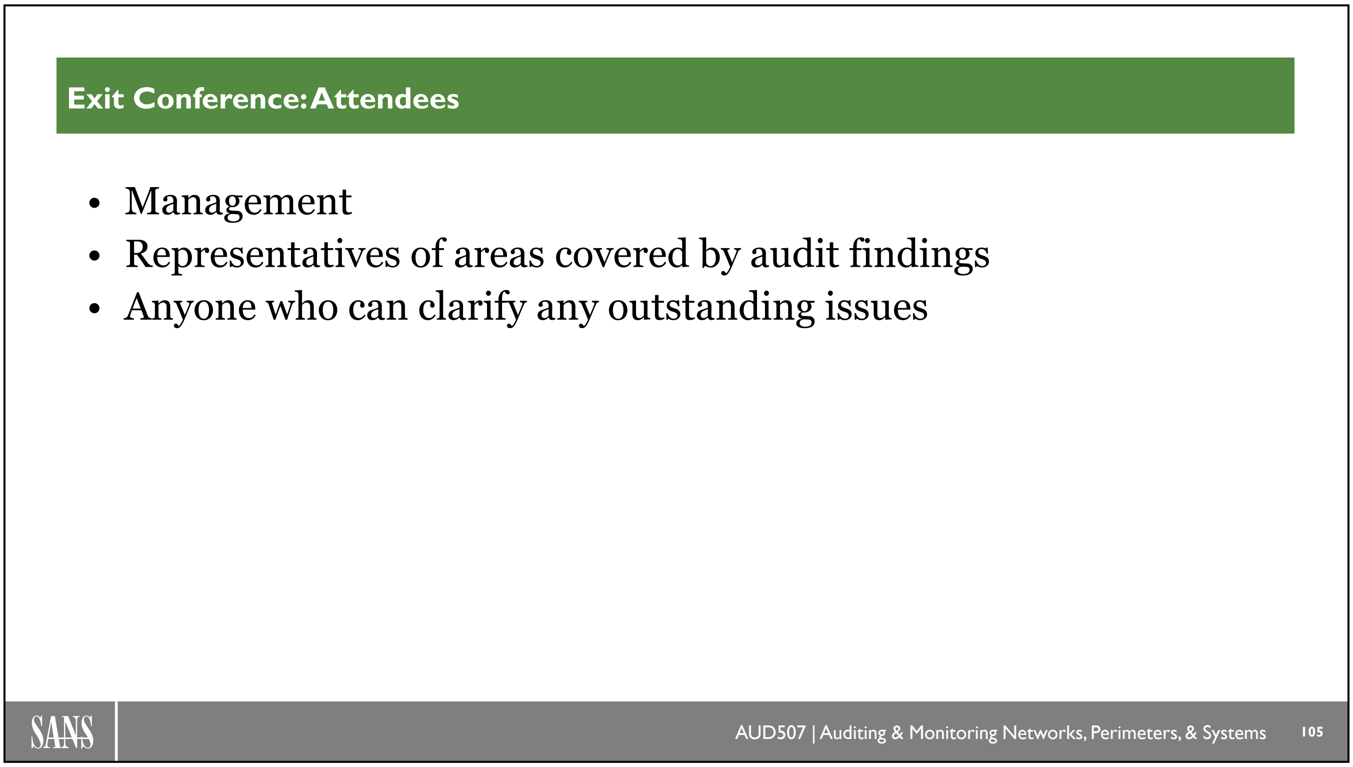

Exit Conference: Attendees

- Management
- Representatives of areas covered by audit findings
- Anyone who can clarify any outstanding issues

The attendee list for this meeting need not be quite as broad as for the entrance meeting. Usually, the auditee will know which people are concerned with the eventual content of the report, and those people will be invited. If the auditors still have specific unresolved questions, the appropriate people should be invited to help deal with them.

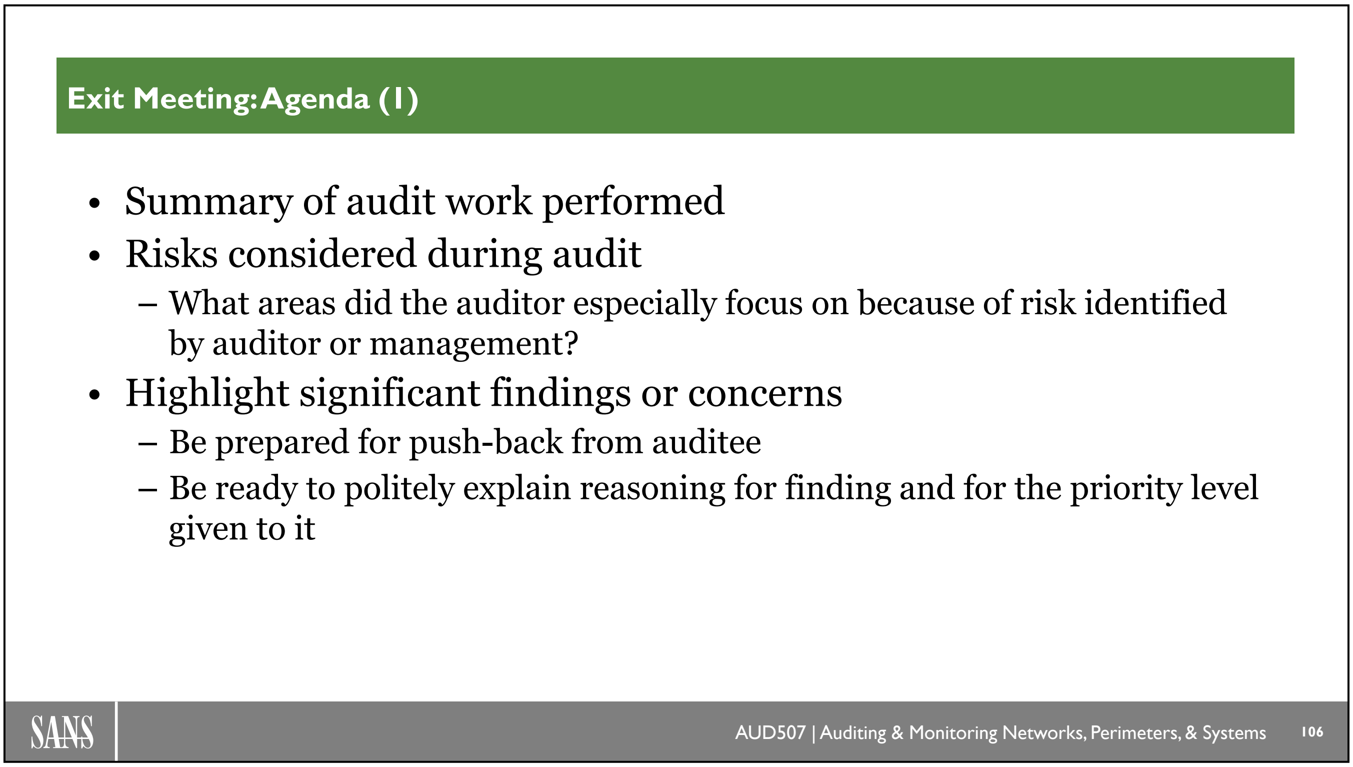

Exit Meeting: Agenda (1)

- Summary of audit work performed
- Risks considered during audit
  - What areas did the auditor especially focus on because of risk identified by auditor or management?
- Highlight significant findings or concerns
  - Be prepared for push-back from auditee
  - Be ready to politely explain reasoning for finding and for the priority level given to it

The lead auditor should be prepared to lead this entire meeting. The agenda should include a description of the work performed during the audit, including brief descriptions of the fieldwork performed and the results of any testing.

Areas of particularly high risk which were discovered during the audit should be addressed, along with a description of how the auditor addressed those risks.

The lead auditor should highlight any significant findings or areas of heightened concern. The auditor should be prepared to explain how conclusions were reached and why the associated priority levels were chosen. The ability to back up these conclusions with the audit evidence collected during the engagement is critical to the auditor's credibility.

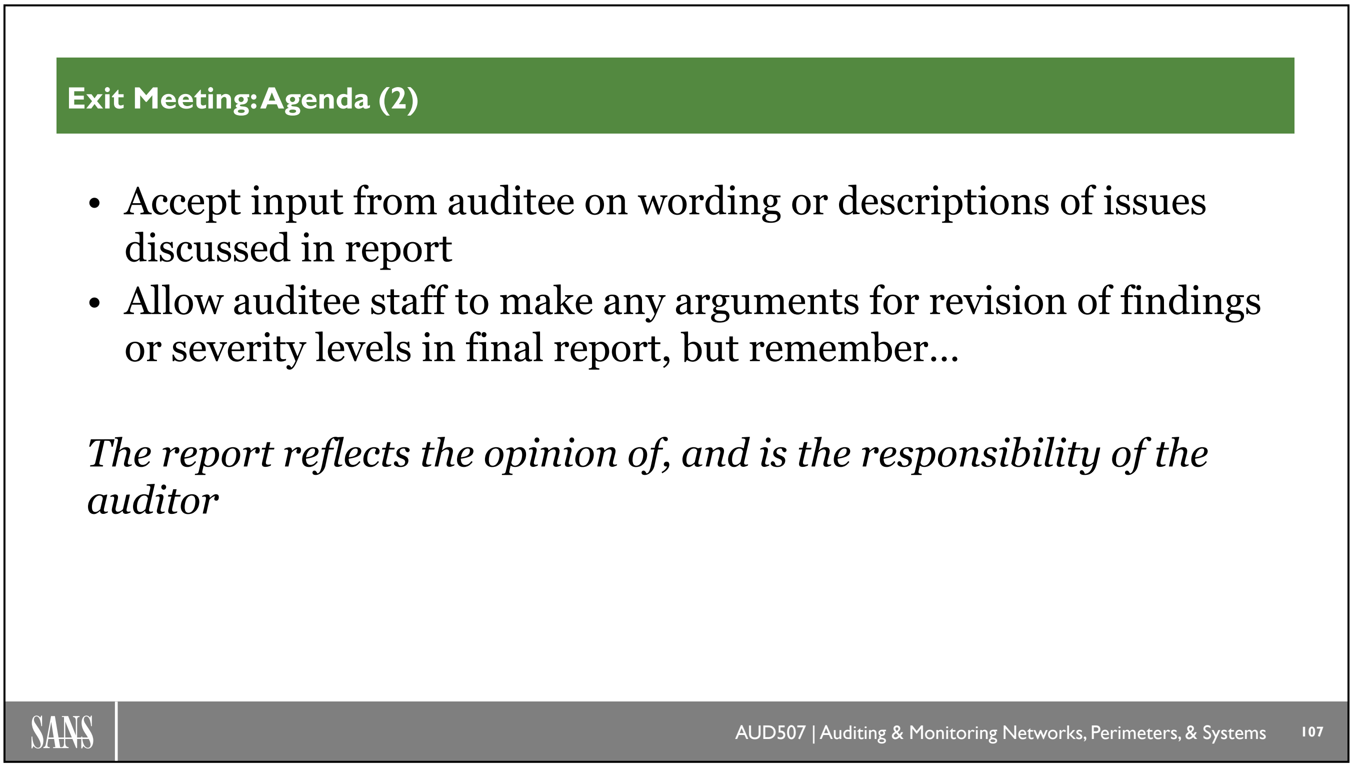

Exit Meeting: Agenda (2)

- Accept input from auditee on wording or descriptions of issues discussed in report
- Allow auditee staff to make any arguments for revision of findings or severity levels in final report, but remember…

The report reflects the opinion of, and is the responsibility of the auditor

The auditor should be prepared for "pushback" or resistance from the auditee regarding some or all of the audit findings. It is perfectly acceptable for the auditor to accept input from auditee staff on wording or the descriptions of findings. It is also acceptable for an auditor to change their opinion about an issue when *new* facts are brought to light by the auditee. However, it is critical that the auditor retain their professional independence and credibility by remembering that the audit report is solely their responsibility and should reflect only their opinion.

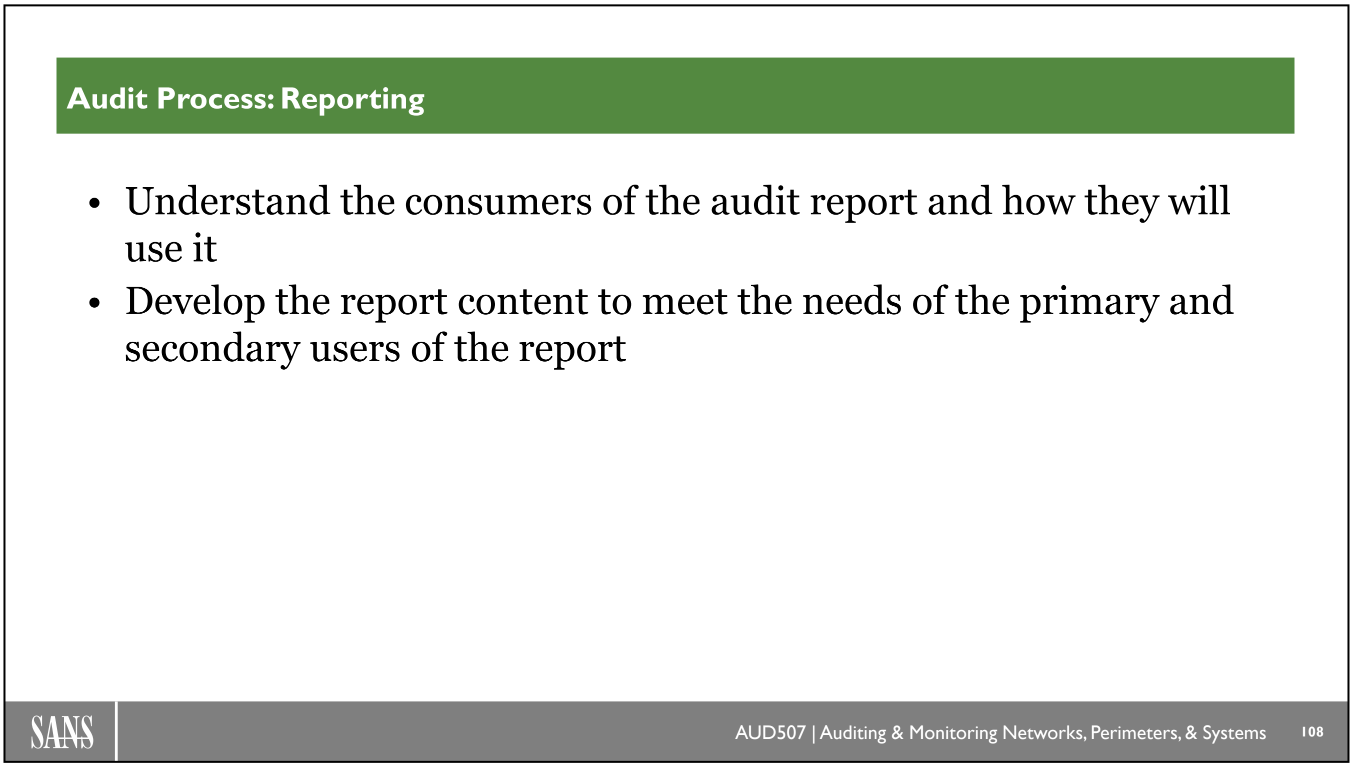

Audit Process: Reporting

- Understand the consumers of the audit report and how they will use it
- Develop the report content to meet the needs of the primary and secondary users of the report

While the audit report will be addressed to management, some other people will read and use the report as part of their responsibility to the organization. As you write the report, you need to remember who these users are and what they need to receive from the report contents. This will help to keep the content clear and focused.

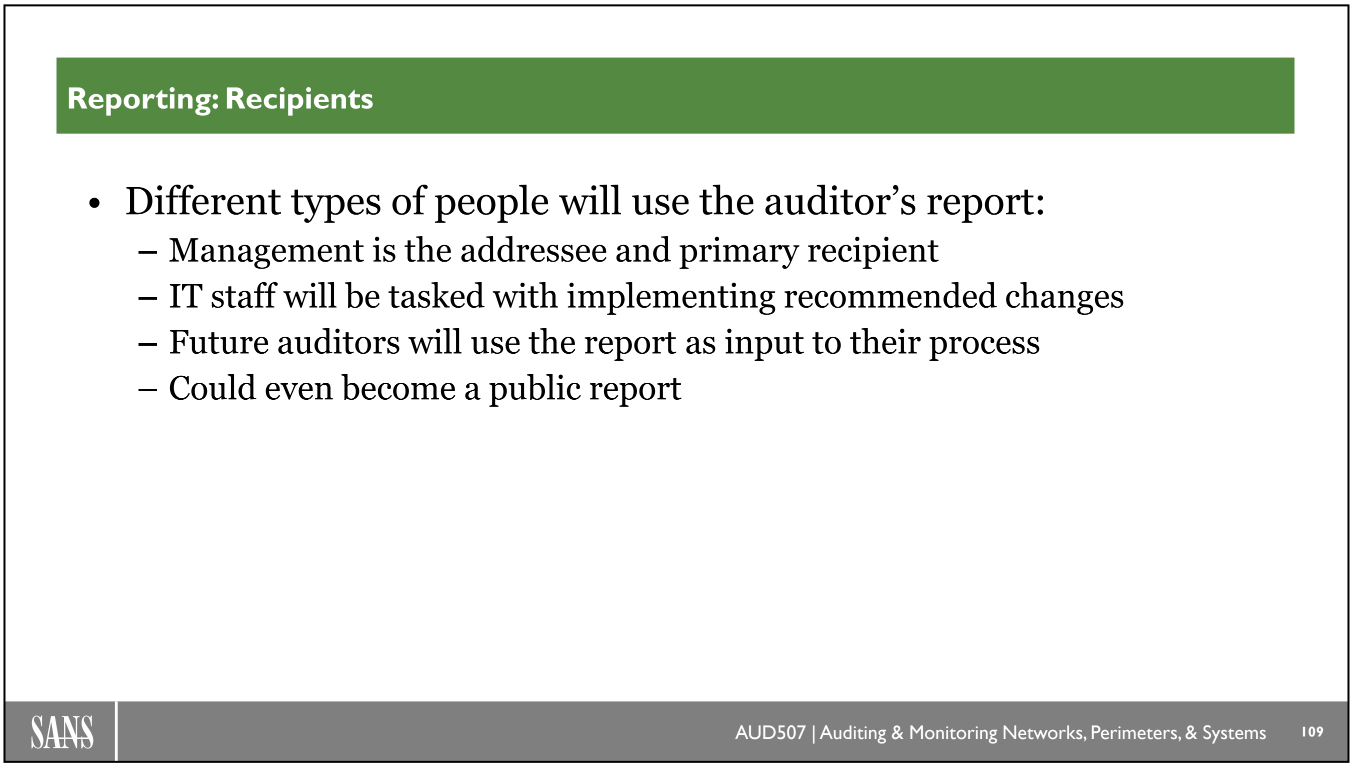

Reporting: Recipients

- Different types of people will use the auditor's report:
  - Management is the addressee and primary recipient
  - IT staff will be tasked with implementing recommended changes
  - Future auditors will use the report as input to their process
  - Could even become a public report

Multiple types of people are likely to read your reports.

Management is the primary recipient, and the report will be addressed to them. The auditor's first job is to make the report useful to management as they seek to guide the organization.

Technical staff will be very concerned with the details of the report and may ask very pointed questions after reading the report. The auditor should be prepared to discuss specifics of testing and findings with the technical staff, although this ideally would have happened during the exit meeting. The technical staff will also be tasked with designing and implementing any recommended changes to the environment. The report should be detailed enough to allow them to do this.

Next, the report will likely be used by future auditors as they perform their duties. The testing descriptions and justification for findings should be sufficiently detailed for these auditors to understand each item included in the report.

Finally, some audit reports may become public. Therefore, the report needs to be written in a format so that anyone can understand the content.

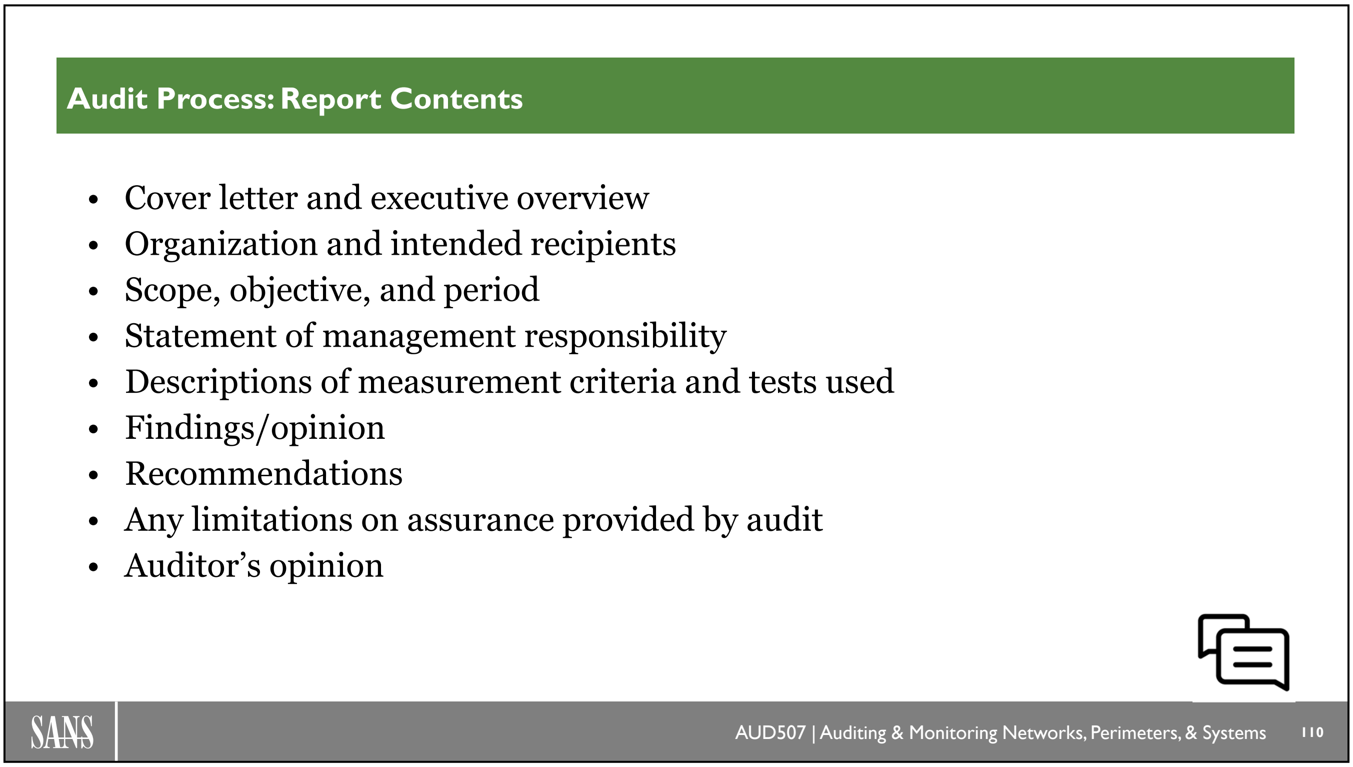

Audit Process: Report Contents

- Cover letter and executive overview
- Organization and intended recipients
- Scope, objective, and period
- Statement of management responsibility
- Descriptions of measurement criteria and tests used
- Findings/opinion
- Recommendations
- Any limitations on assurance provided by audit
- Auditor's opinion

The following slides cover the major content areas in the report. Depending on the audit scope and objectives, some of these areas may be omitted, or other content may need to be added. For example, not all reports require an "auditors opinion". If a non-CPA is performing the review, an "opinion" often cannot be stated.

The areas to be covered include:

- Cover letter and executive overview
- Organization and intended recipients
- Scope, objective, and period
- Statement of management responsibility
- Descriptions of measurement criteria and tests used
- Findings/opinion
- Recommendations
- Any limitations on assurance provided by audit
- Auditor's opinion

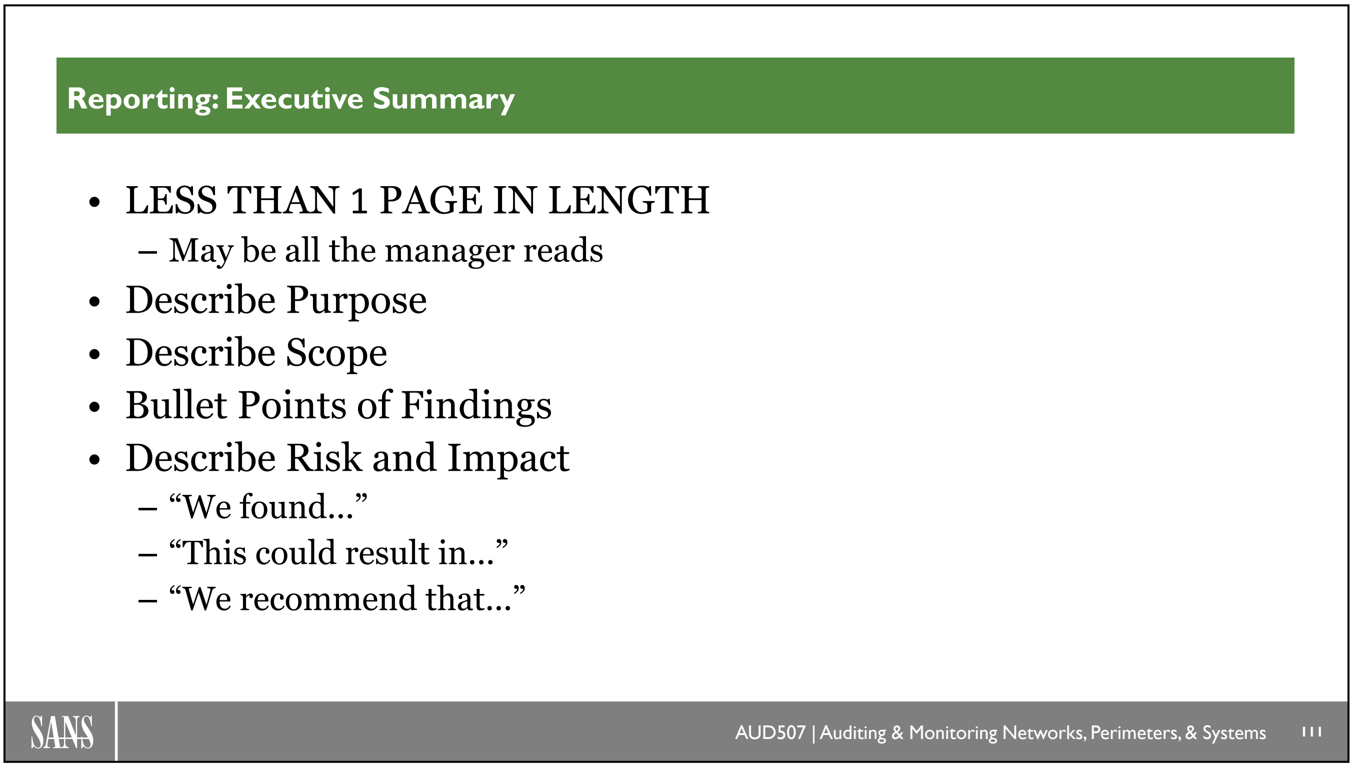

Reporting: Executive Summary

- LESS THAN 1 PAGE IN LENGTH
  - May be all the manager reads
- Describe Purpose
- Describe Scope
- Bullet Points of Findings
- Describe Risk and Impact
  - “We found…”
  - “This could result in…”
  - “We recommend that…”

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 111

Your executive summary should be the last thing you write for your audit report, even though it will be the first page or so of the report. While writing the summary, keep in mind your intended audience; this portion of the report should be specifically tailored with an executive as the target audience. Although we all enjoy pointing out how nontechnical executives tend to be, remember that their jobs don’t usually require them to be technical. Their position requires them to make good business decisions to protect assets and to further business objectives. As an auditor, you are an executive tool to measure conformance to a standard and to assess how the organization can improve.

Be sure to state the purpose of the audit clearly. Include your audit objectives or what it was you were asked to measure or assess. You should also clearly describe the scope of the audit. If possible, try to give credit to the executive-level individual you worked with during your planning stage to define your scope. After this is out of the way, you should introduce a set of bullet points. Make sure you include high points!!! If you give them only bad news, they will begin to have a negative view of the individuals responsible for the systems. This is your chance to show them where those individuals are doing a good job as well as to define what areas need improvement. Remember that the audit is not necessarily a measure of how well people do their job. If the UNIX admin, for instance, is also responsible for administering the Windows domain, user desktop support, and updating the web server, it is entirely possible that some things have slipped on the UNIX system, but not because he didn’t want to do them. You could include a summary statement as follows:

“We found that the individuals tasked with maintaining this system have too many other diverse responsibilities. This could result in degradation in the security of the systems involved because no one individual can be expected to reasonably secure all those systems simultaneously and consistently. We recommend that management review staffing levels and consider hiring new staff or cross-training existing staff to cover some of these duties.”

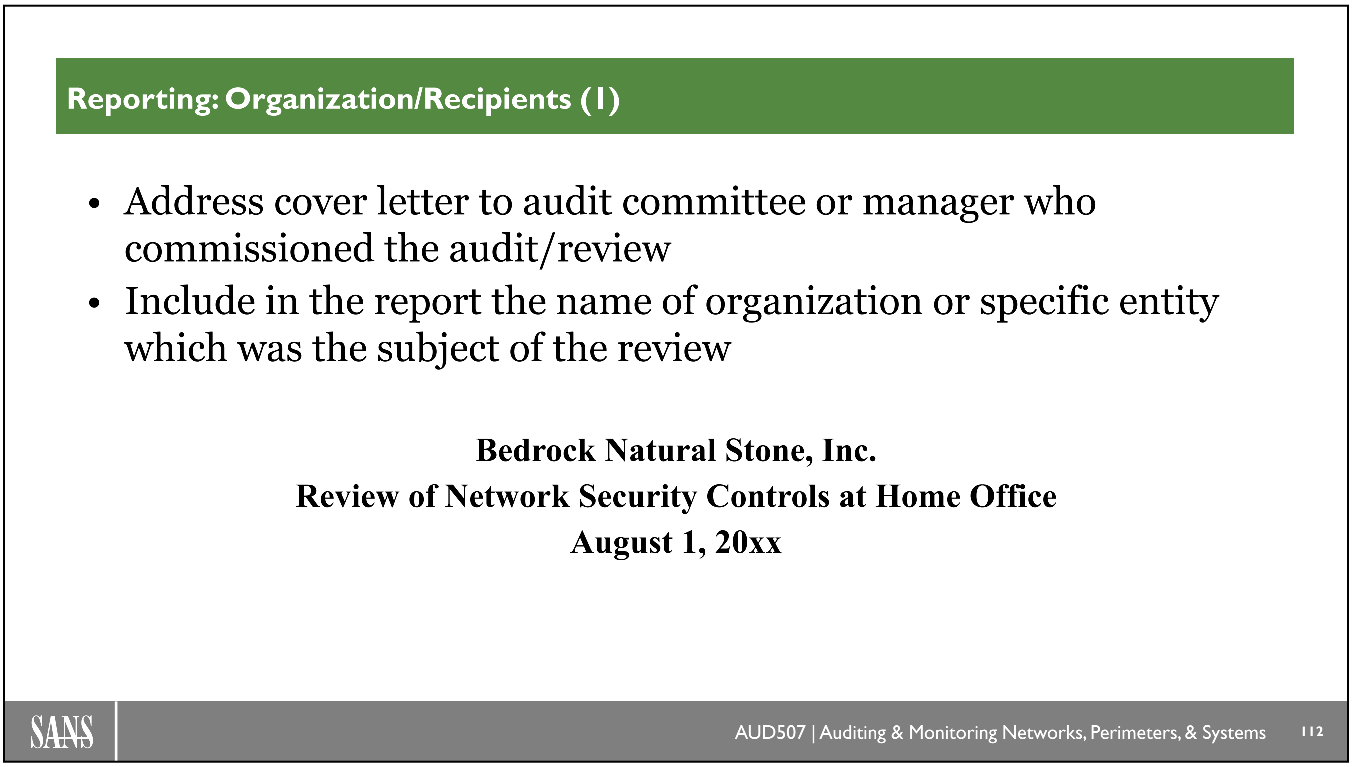

Reporting: Organization/Recipients (1)

- Address cover letter to audit committee or manager who commissioned the audit/review
- Include in the report the name of organization or specific entity which was the subject of the review

Bedrock Natural Stone, Inc.
Review of Network Security Controls at Home Office
August 1, 20xx

The audit report should include the name of the organization which was the subject of the review. It is common to place this in the title on the first page of the report.

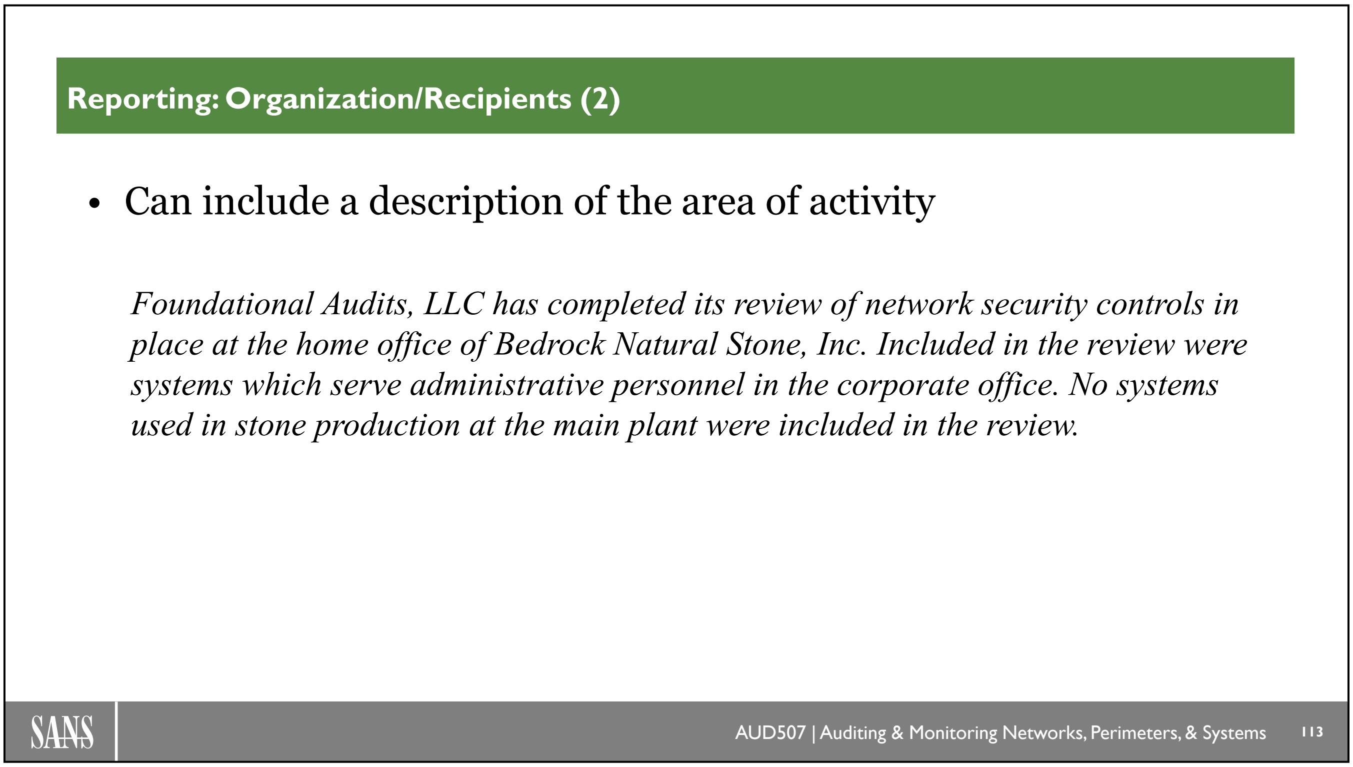

Reporting: Organization/Recipients (2)

- Can include a description of the area of activity

Foundational Audits, LLC has completed its review of network security controls in place at the home office of Bedrock Natural Stone, Inc. Included in the review were systems which serve administrative personnel in the corporate office. No systems used in stone production at the main plant were included in the review.

Particularly when the scope of the audit is limited in some way (the audit does not cover the entire organization), it is helpful to specifically describe the scope in the report. Where there are specific limitations or exclusions from the scope, you should also spell those out here.

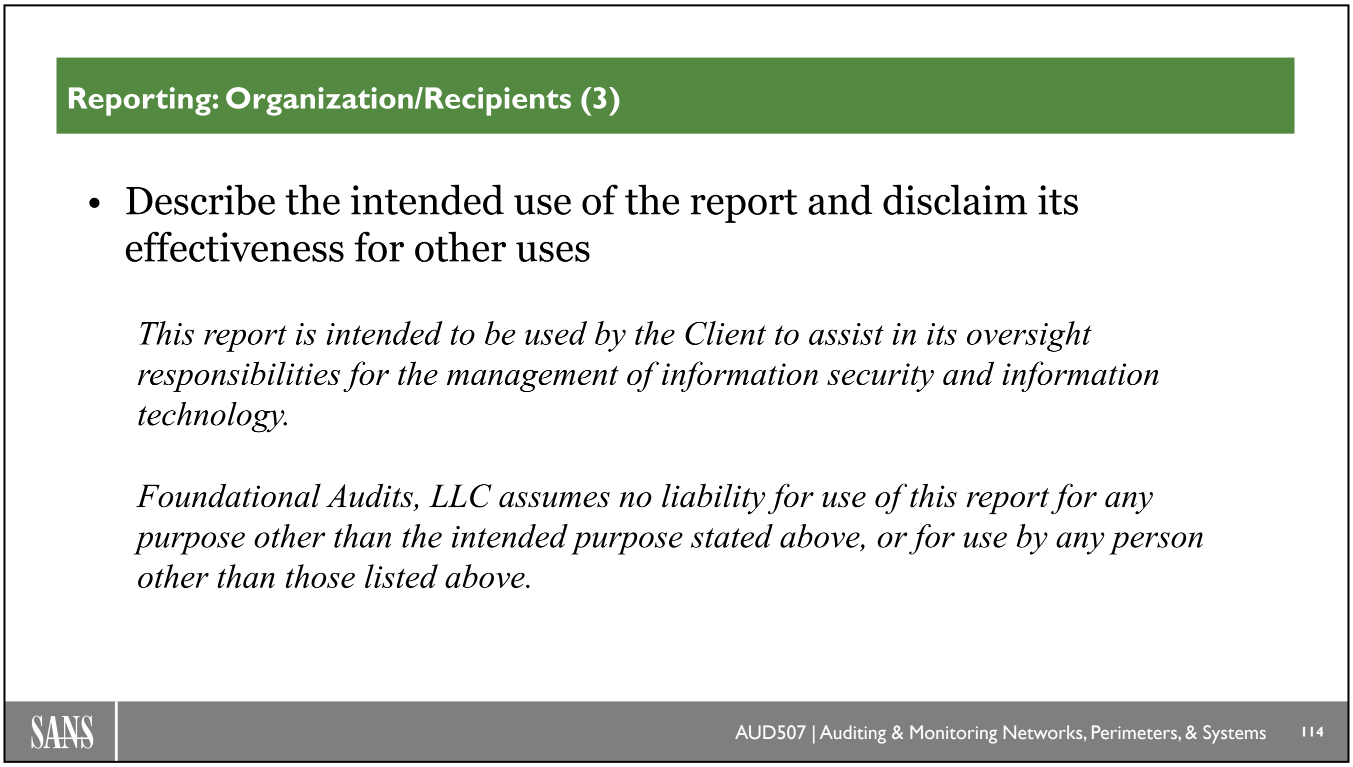

Reporting: Organization/Recipients (3)

- Describe the intended use of the report and disclaim its effectiveness for other uses

This report is intended to be used by the Client to assist in its oversight responsibilities for the management of information security and information technology.

Foundational Audits, LLC assumes no liability for use of this report for any purpose other than the intended purpose stated above, or for use by any person other than those listed above.

Audits are always commissioned to help an organization solve a very specific problem—complying with a regulatory requirement, for instance. While the final audit report may be useful to the organization or some third-party for some other use, the auditor should always disclaim those other uses in the report.

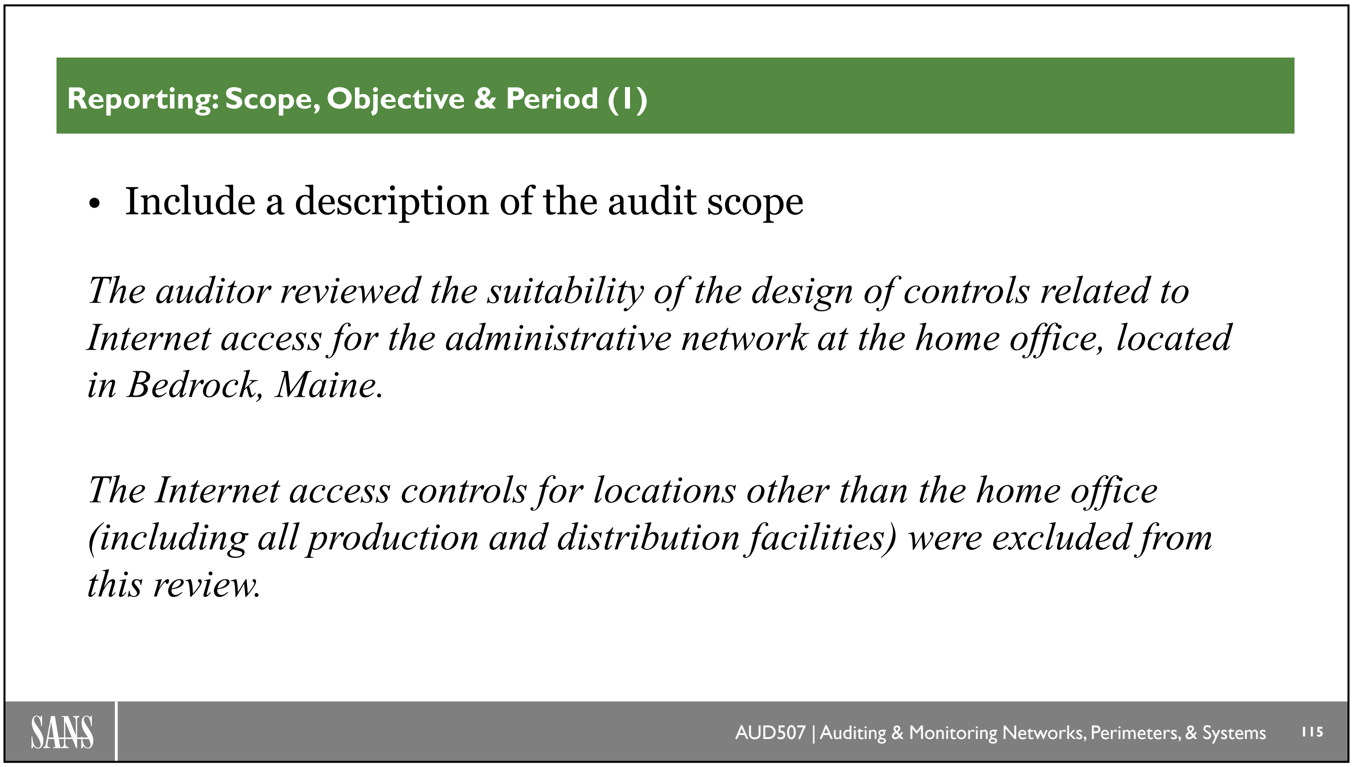

Reporting: Scope, Objective & Period (1)

- Include a description of the audit scope

The auditor reviewed the suitability of the design of controls related to Internet access for the administrative network at the home office, located in Bedrock, Maine.

The Internet access controls for locations other than the home office (including all production and distribution facilities) were excluded from this review.

The description of the audit scope should be as specific as possible. This section should include a description of the entity, business unit, or department under review, as well as a description of specific business processes, systems, or functional areas which were included.

If specific systems or areas were *excluded* from the review, this should also be noted here.

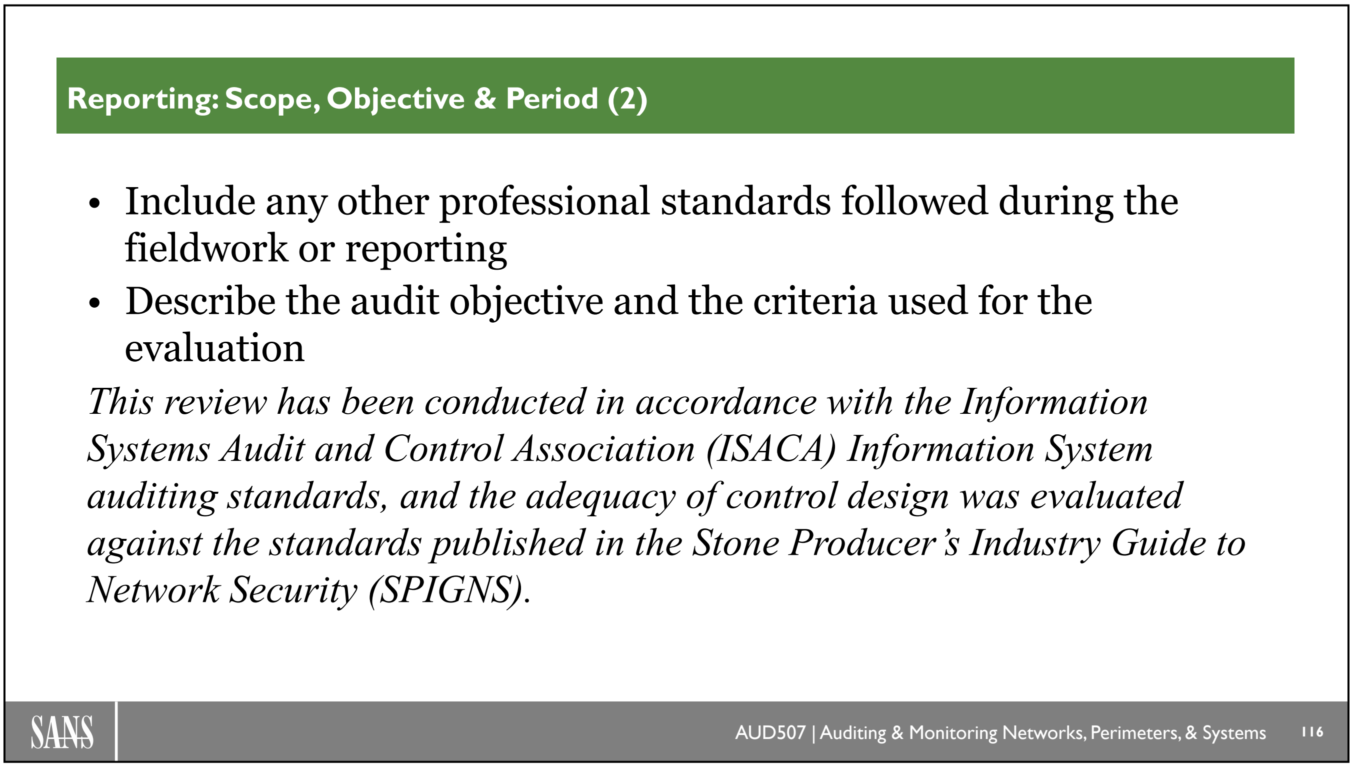

Reporting: Scope, Objective & Period (2)

- Include any other professional standards followed during the fieldwork or reporting
- Describe the audit objective and the criteria used for the evaluation

This review has been conducted in accordance with the Information Systems Audit and Control Association (ISACA) Information System auditing standards, and the adequacy of control design was evaluated against the standards published in the Stone Producer's Industry Guide to Network Security (SPIGNS).

It's not uncommon for an auditor to follow procedures and standards issued by more than one organization. In the example on this slide, the auditor used security standards from an industry organization and audit standards from ISACA. Providing this information allows future users of the report to better understand how the fieldwork was performed and what standards the auditor was using for testing.

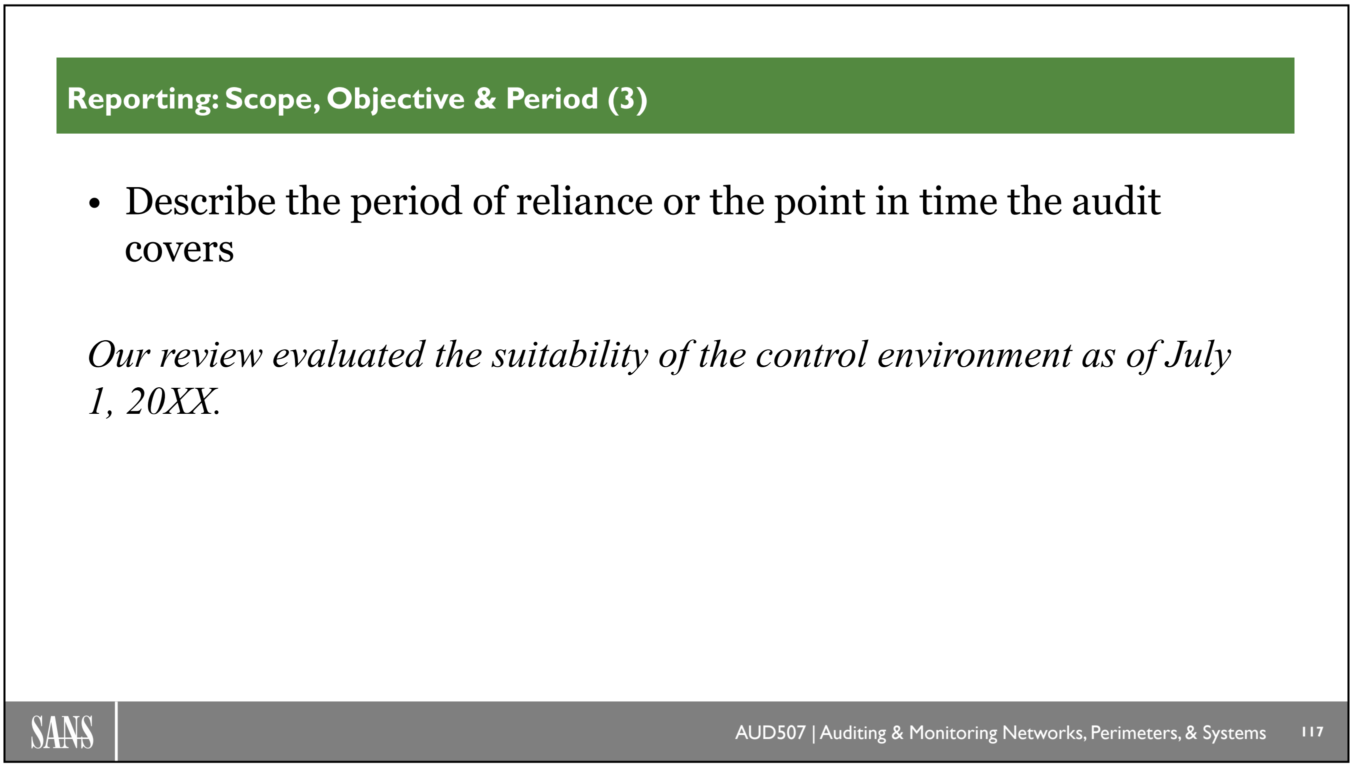

Reporting: Scope, Objective & Period (3)

- Describe the period of reliance or the point in time the audit covers

Our review evaluated the suitability of the control environment as of July 1, 20XX.

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 117

It is important for readers of the audit report to be able to easily determine whether the audit covered a point-in-time or a period of reliance. Spelling that out in the audit description will make that easier for them.

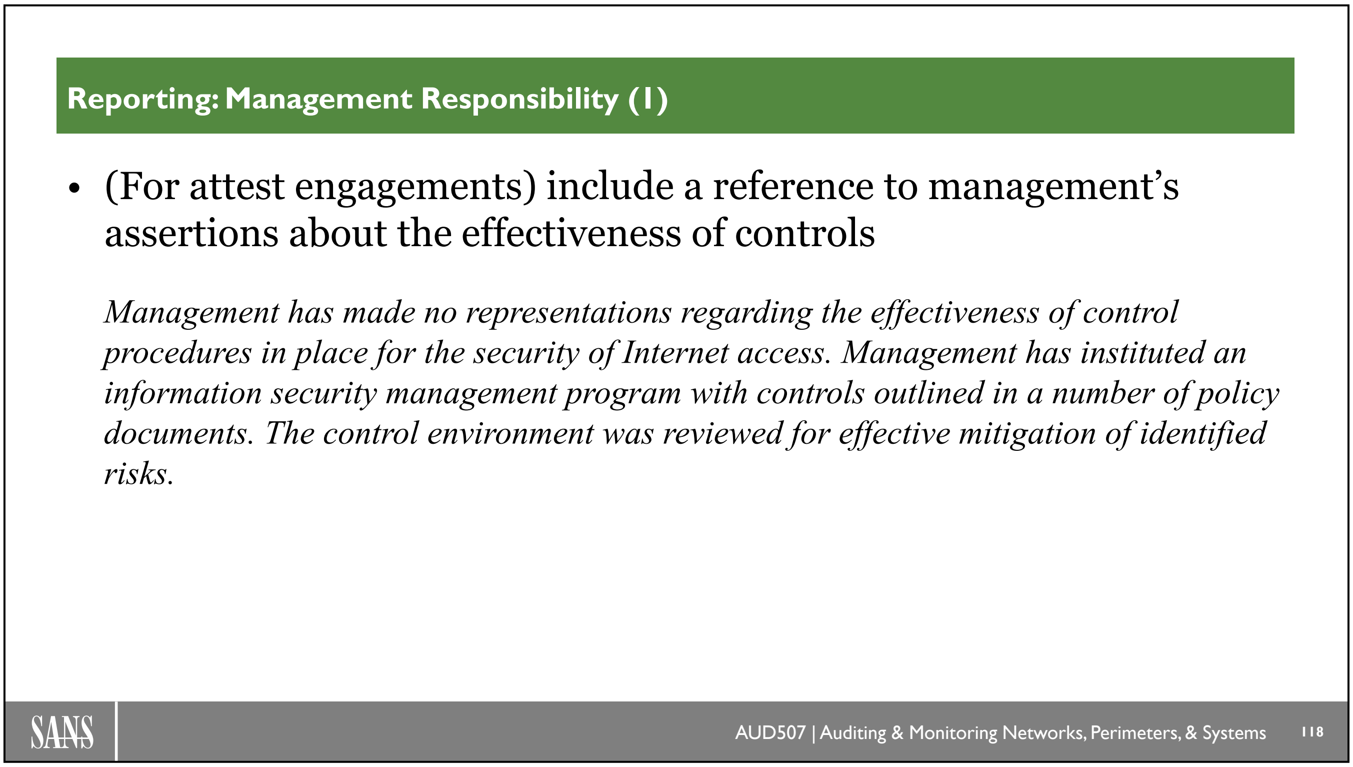

Reporting: Management Responsibility (1)

- (For attest engagements) include a reference to management's assertions about the effectiveness of controls

Management has made no representations regarding the effectiveness of control procedures in place for the security of Internet access. Management has instituted an information security management program with controls outlined in a number of policy documents. The control environment was reviewed for effective mitigation of identified risks.

In an attest engagement (like a SOC1 audit), the auditor would include a reference to management's assertions about the effectiveness of the controls under review.

Even when management has not made any assertions about the control environment, it can be a good idea to include a mention of the lack of assertions in the report. An example of that is given in this slide, taken from a fictitious review of the effectiveness of internet security controls.

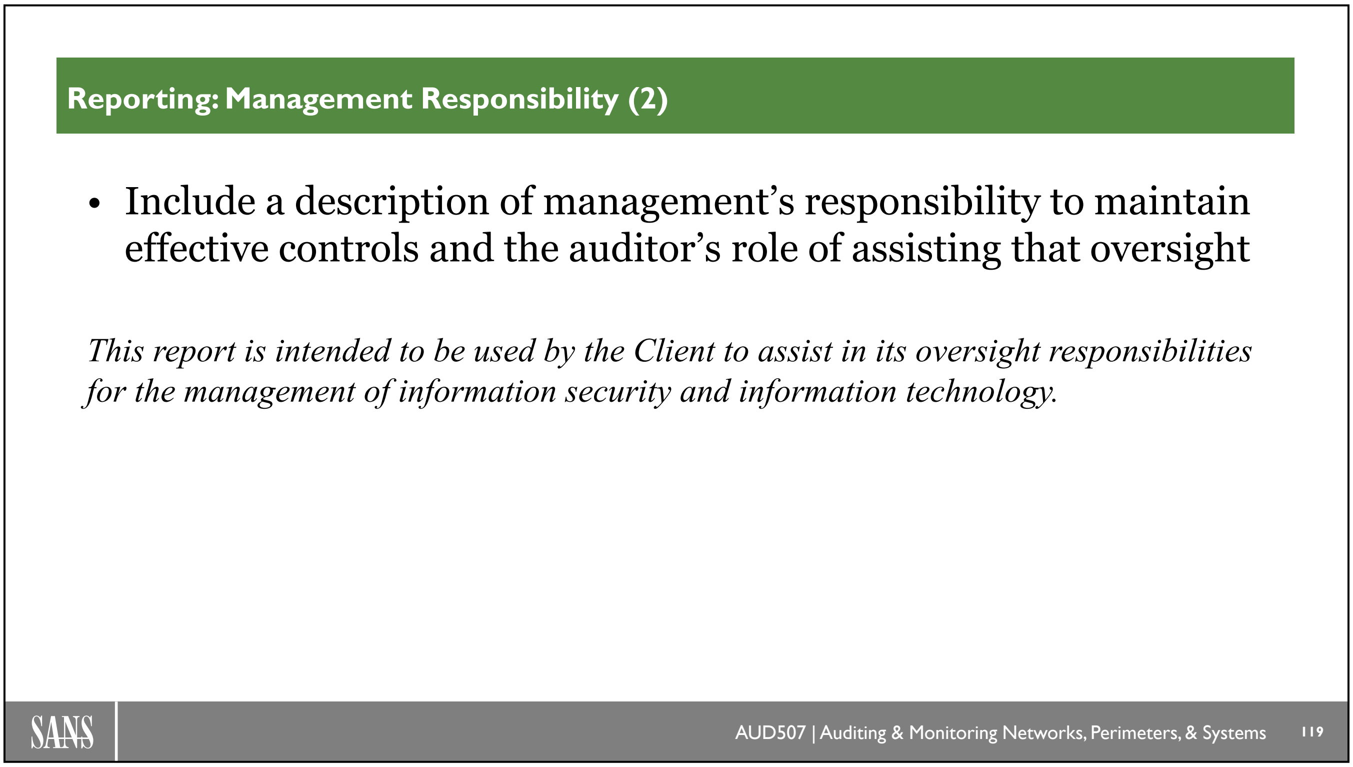

Reporting: Management Responsibility (2)

- Include a description of management's responsibility to maintain effective controls and the auditor's role of assisting that oversight

This report is intended to be used by the Client to assist in its oversight responsibilities for the management of information security and information technology.

SANS
AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems
119

The auditor is merely one tool available to management as they perform their duty to maintain controls that effectively protect the organization against identified risk.

The audit report should make it clear that the audit was conducted to help management in fulfilling its responsibilities.

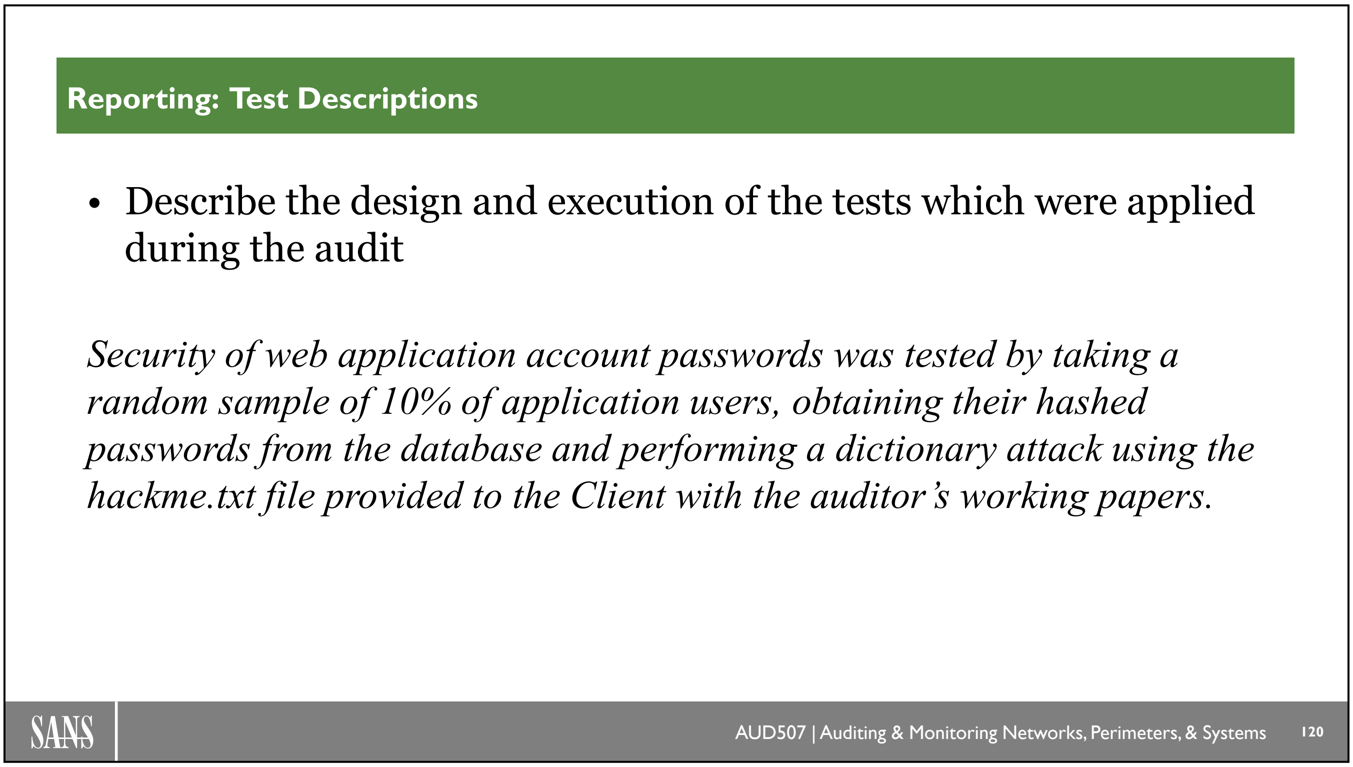

Reporting: Test Descriptions

- Describe the design and execution of the tests which were applied during the audit

Security of web application account passwords was tested by taking a random sample of 10% of application users, obtaining their hashed passwords from the database and performing a dictionary attack using the hackme.txt file provided to the Client with the auditor's working papers.

A clear description of each type of testing performed should be included somewhere in the main body of the audit report. When the auditor wishes to also include screen captures or other archival evidence of the test results, these are normally included in an appendix.

The test descriptions should be of sufficient detail that another auditor, a skilled security professional, or even a trained system administrator could reproduce the tests performed during the audit.

Where specialized tools or input files have been used, they should be provided to the auditee (when appropriate) in digital format (usually burned to disk) so that the tests may be reproduced.

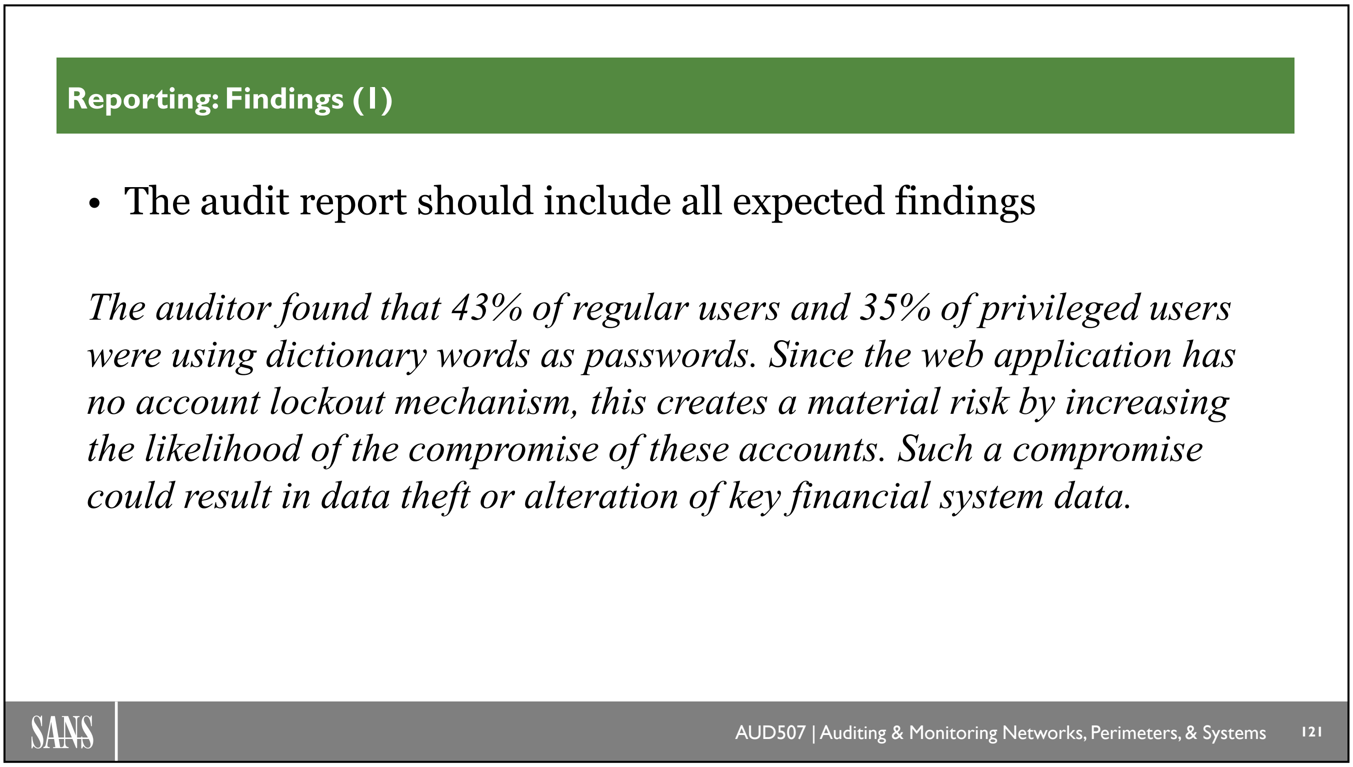

Reporting: Findings (1)

- The audit report should include all expected findings

The auditor found that 43% of regular users and 35% of privileged users were using dictionary words as passwords. Since the web application has no account lockout mechanism, this creates a material risk by increasing the likelihood of the compromise of these accounts. Such a compromise could result in data theft or alteration of key financial system data.

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 121

The audit report should include a paragraph describing each expected finding. This paragraph should include a brief explanation of why the finding is important, and what the result of a control failure or bypass could be. For some audits, all findings will be included. For other audits, only significant findings may be included. It will depend on the audit being conducted.

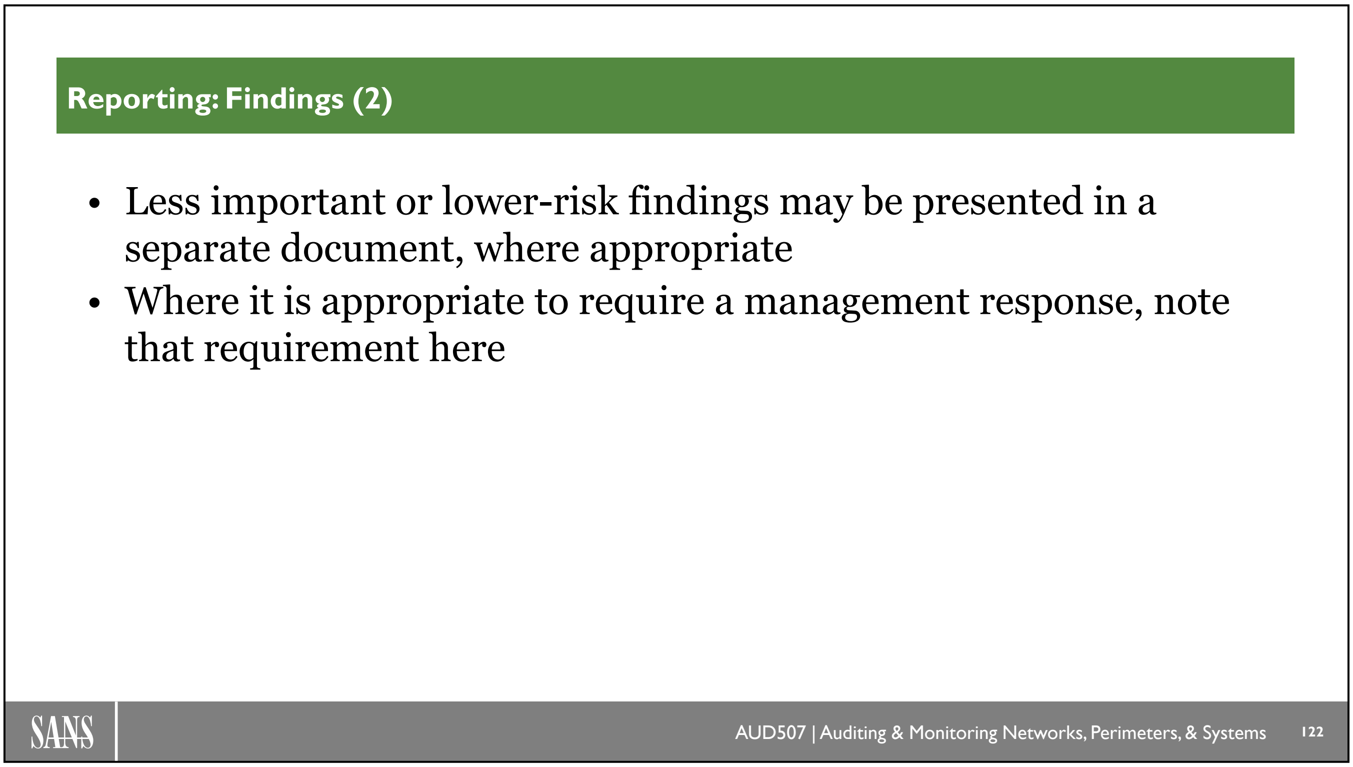

Reporting: Findings (2)

- Less important or lower-risk findings may be presented in a separate document, where appropriate
- Where it is appropriate to require a management response, note that requirement here

To keep the audit report to a manageable length, the auditor *may* decide to include lower-significance findings in an audit memorandum, report appendix, or some other document.

Depending on the type of audit engagement, it may be appropriate for the auditor to request or require an “official” response from management to certain findings. Where this is the case, the response requirement should be noted, and any timelines for follow-up should be explicitly spelled out.

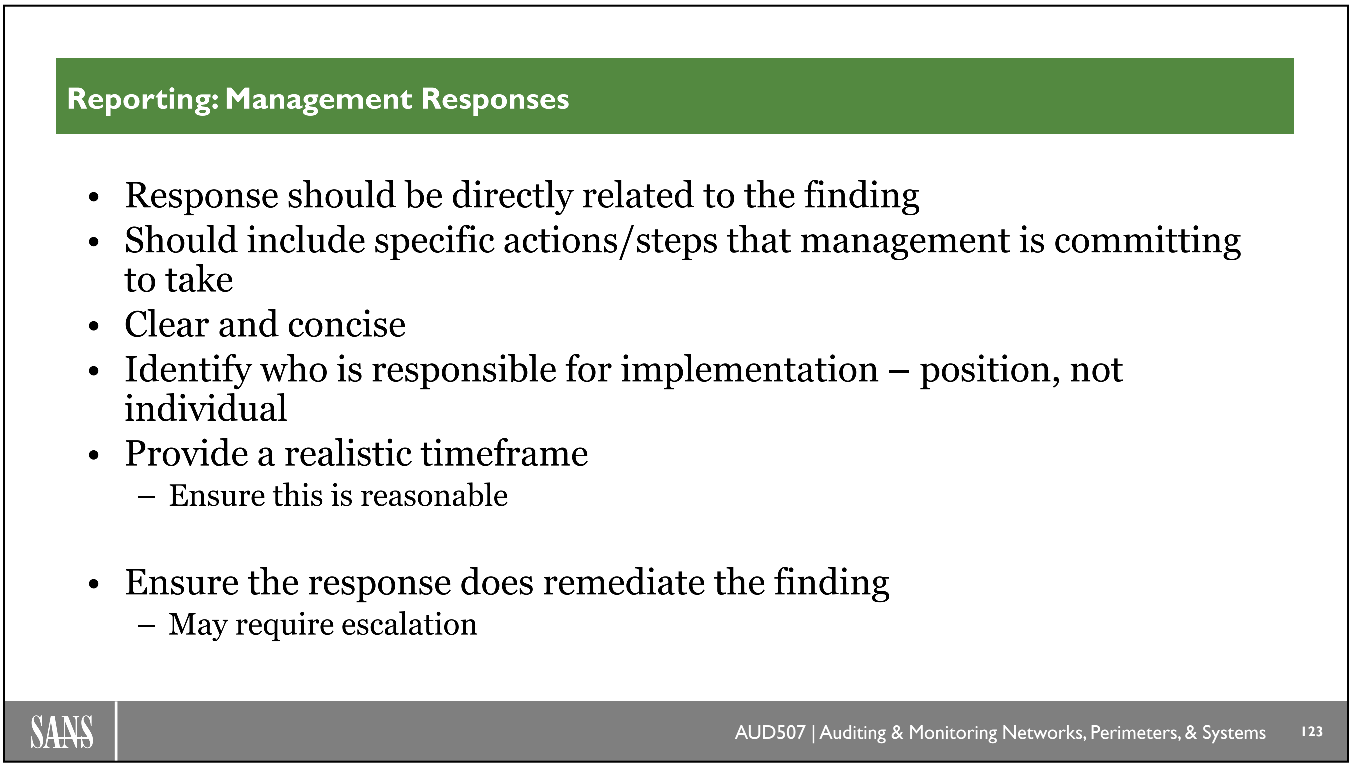

Reporting: Management Responses

- Response should be directly related to the finding
- Should include specific actions/steps that management is committing to take
- Clear and concise
- Identify who is responsible for implementation – position, not individual
- Provide a realistic timeframe
  - Ensure this is reasonable
- Ensure the response does remediate the finding
  - May require escalation

Since we're mentioning management responses, let's also look at what should be in a management response. At a minimum, the following should have the following characteristics:

- Response should be directly related to the finding
- Should include specific actions that management is committing to take
- Clear and concise
- Identify who is responsible for implementation
- Provide a realistic timeframe

Upon receipt, the auditor will want to confirm the action taken by management remediates the finding. If not, discussion should continue with management until such time. Escalation may need to take place if the auditee refuses to comply with a request for an actionable response.

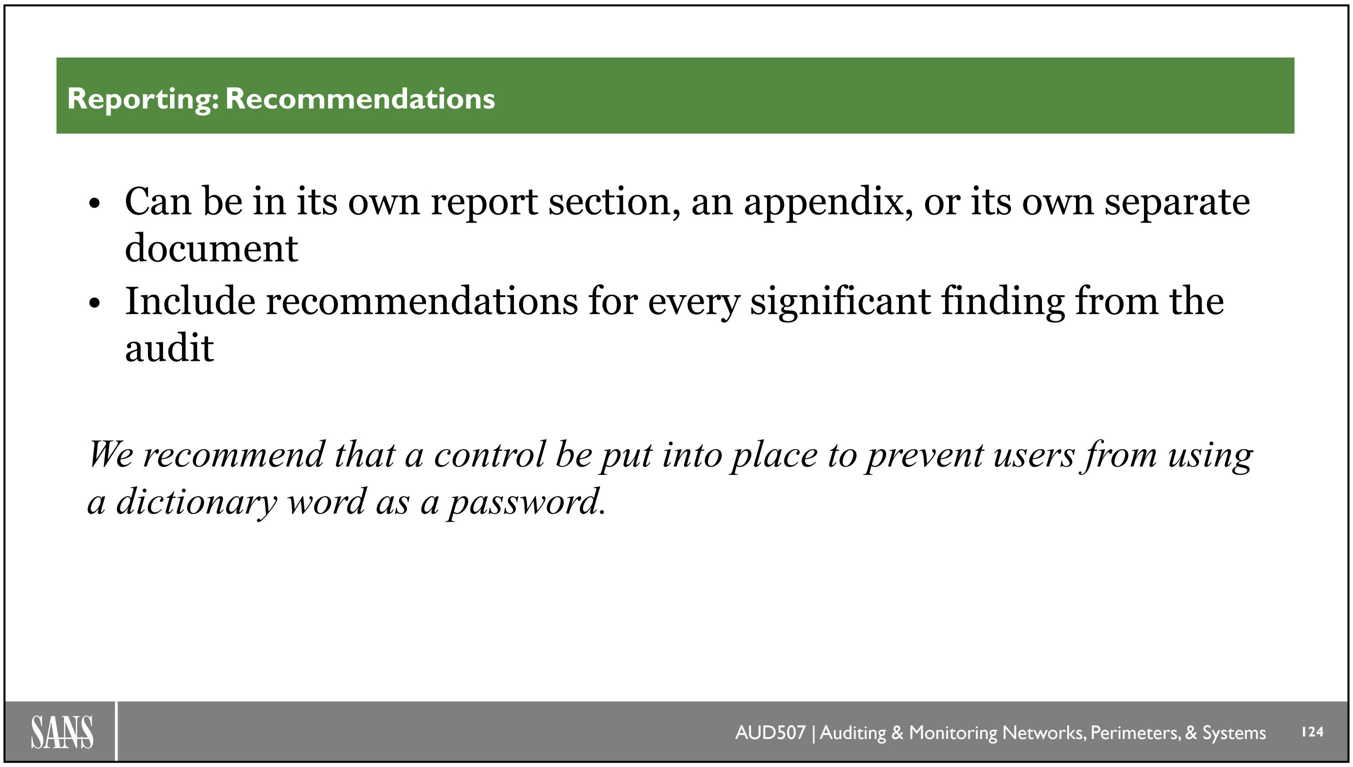

Reporting: Recommendations

- Can be in its own report section, an appendix, or its own separate document
- Include recommendations for every significant finding from the audit

We recommend that a control be put into place to prevent users from using a dictionary word as a password.

Recommendations for corrective action are an important part of the audit report and add real value to the audit. If the auditor only tells the auditee what they are doing wrong, and then leaves them on their own to determine what to do, they are missing the chance to get real problems fixed. Rather, the auditor should make recommendations about every significant finding in the report.

These recommendations can be included with each finding, in their own section of the report (which is a common approach and can make the report easier to read), or in their own document.

Please note that the implementation details of the corrective action are up to the auditee, and that the auditor is not required or expected to design and implement the fixes. Remember that the auditor has a duty to remain independent at all times.

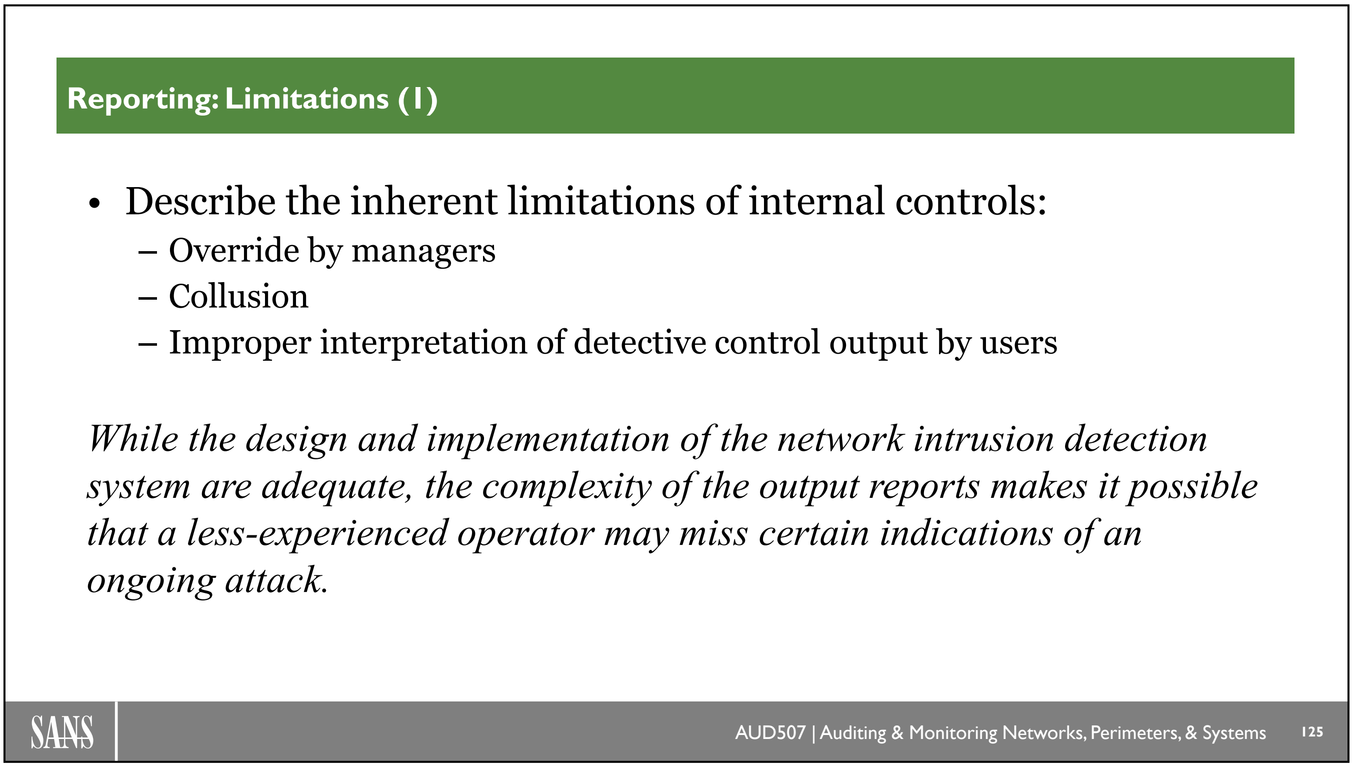

Reporting: Limitations (1)

- Describe the inherent limitations of internal controls:
  - Override by managers
  - Collusion
  - Improper interpretation of detective control output by users

While the design and implementation of the network intrusion detection system are adequate, the complexity of the output reports makes it possible that a less-experienced operator may miss certain indications of an ongoing attack.

There are two types of limitations which the auditor needs to consider when writing the report. The first, covered on this slide, is the limitations of the effectiveness of controls. The second type of limitation has to do with the performance of the audit itself, and is covered over the next two slides.

Most controls are designed to function within a specific environment and do not work outside that environment. A detective control that can be overridden by a manager only works while it is not being overridden. Likewise, a detective control that alerts a user to an unusual condition only works if the user correctly interprets its output. A control based on segregation of duties can be defeated by two people working in collusion.

The report should point out any limitations on the implementation of controls.

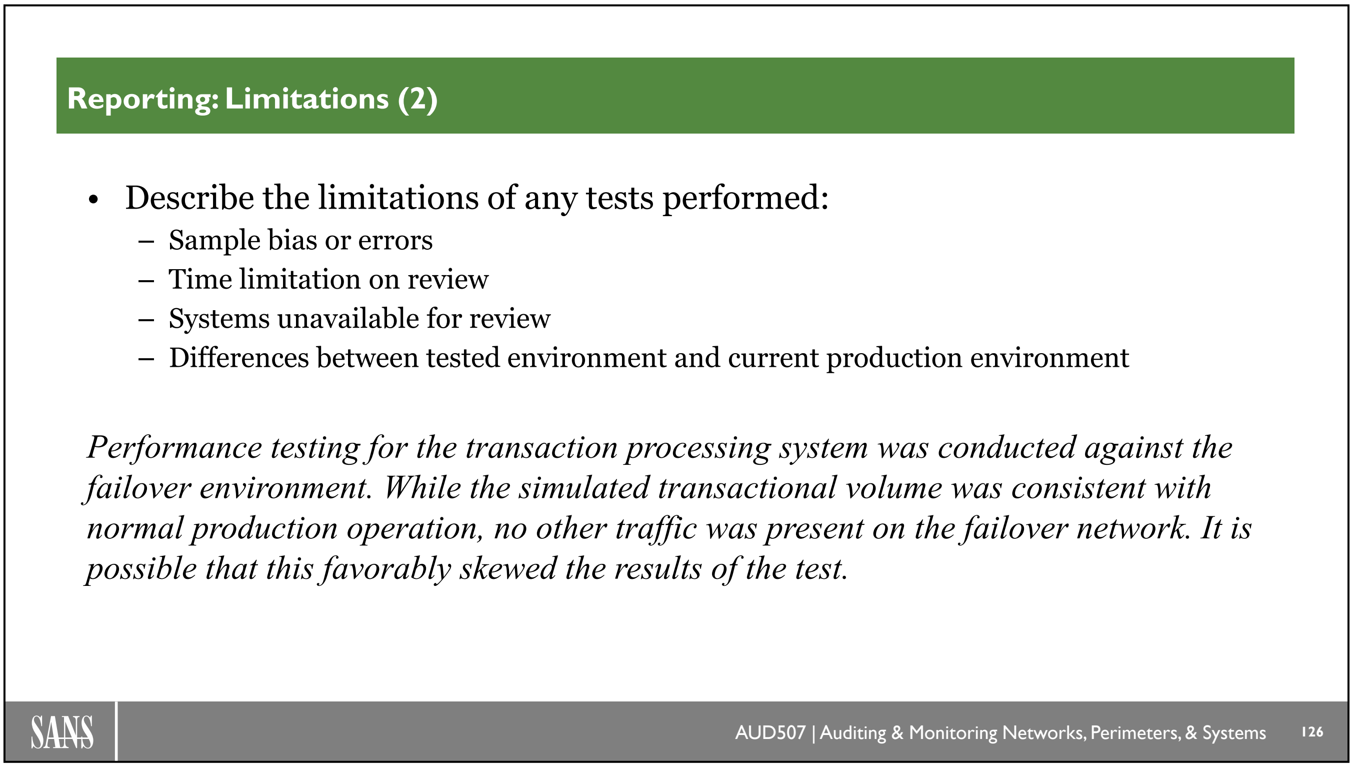

Reporting: Limitations (2)

- Describe the limitations of any tests performed:
  - Sample bias or errors
  - Time limitation on review
  - Systems unavailable for review
  - Differences between tested environment and current production environment

Performance testing for the transaction processing system was conducted against the failover environment. While the simulated transactional volume was consistent with normal production operation, no other traffic was present on the failover network. It is possible that this favorably skewed the results of the test.

The tests performed during the audit may have limitations of their own, usually owing to the fact that the auditor cannot test every transaction that ever happens. The audit has a start and end date, and there is not usually time to test all transactions during the period of reliance. These are limitations which could affect the results.

Sometimes testing is performed against failover systems or the product testing environment in order to minimize disruption to the production environment during testing. In these cases, any deviation between the environments should be noted.

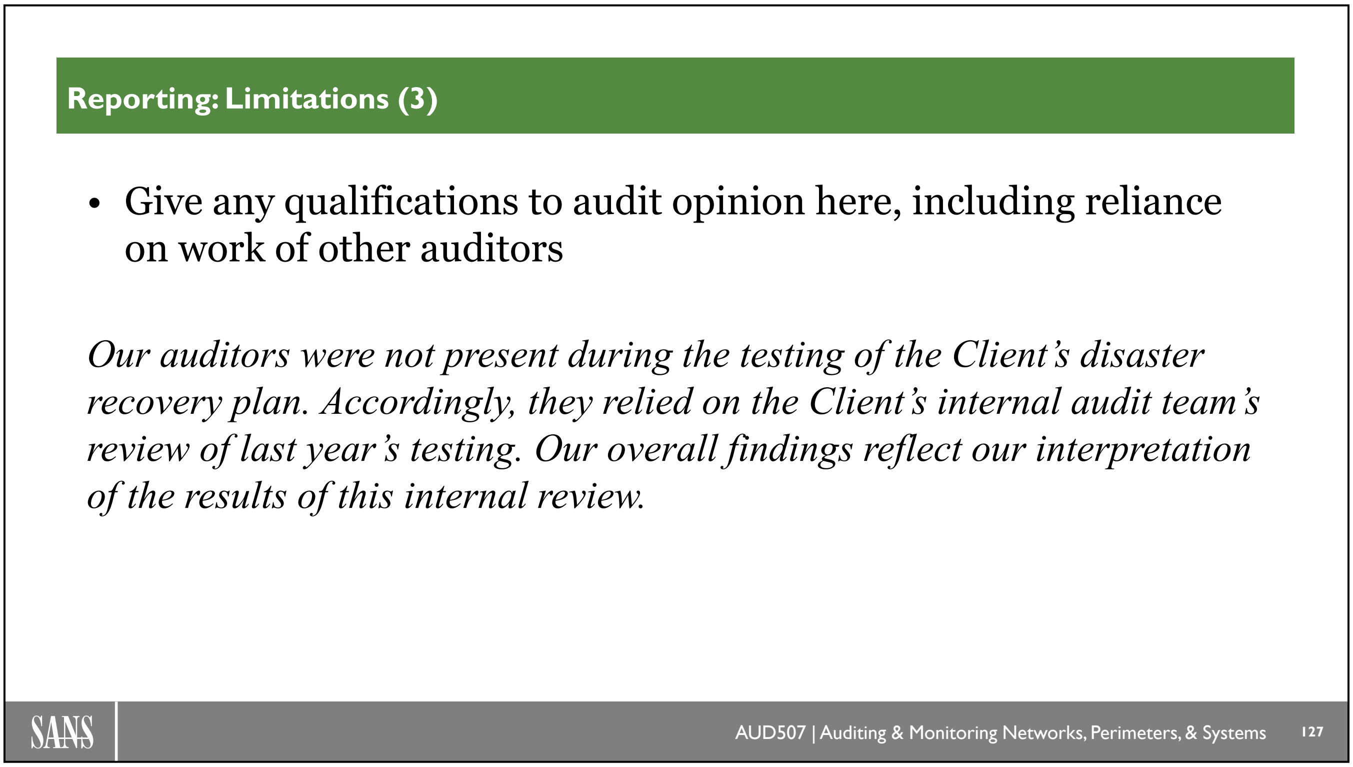

Reporting: Limitations (3)

- Give any qualifications to audit opinion here, including reliance on work of other auditors

Our auditors were not present during the testing of the Client's disaster recovery plan. Accordingly, they relied on the Client's internal audit team's review of last year's testing. Our overall findings reflect our interpretation of the results of this internal review.

Sometimes the auditor must rely on the work and opinions of other auditors who are long gone, leaving only a report behind. When the auditor must base their opinion, at least in part, on the work of another auditor, that should be disclosed.

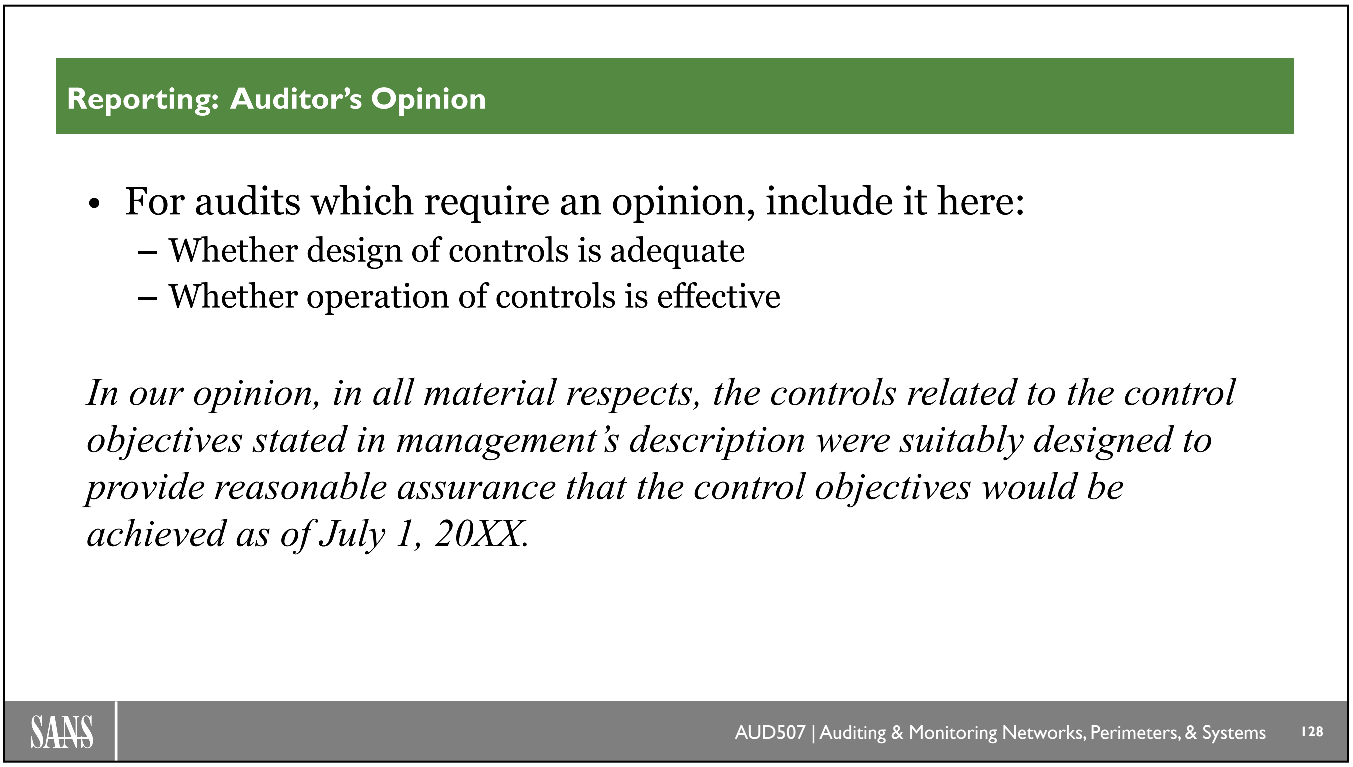

Reporting: Auditor's Opinion

- For audits which require an opinion, include it here:
  - Whether design of controls is adequate
  - Whether operation of controls is effective

In our opinion, in all material respects, the controls related to the control objectives stated in management's description were suitably designed to provide reasonable assurance that the control objectives would be achieved as of July 1, 20XX.

Finally, when the audit objective requires that the auditor issue an opinion, it should be plainly stated in the report. For SOC reporting, the auditor's opinion will often look something like the example text on this slide.

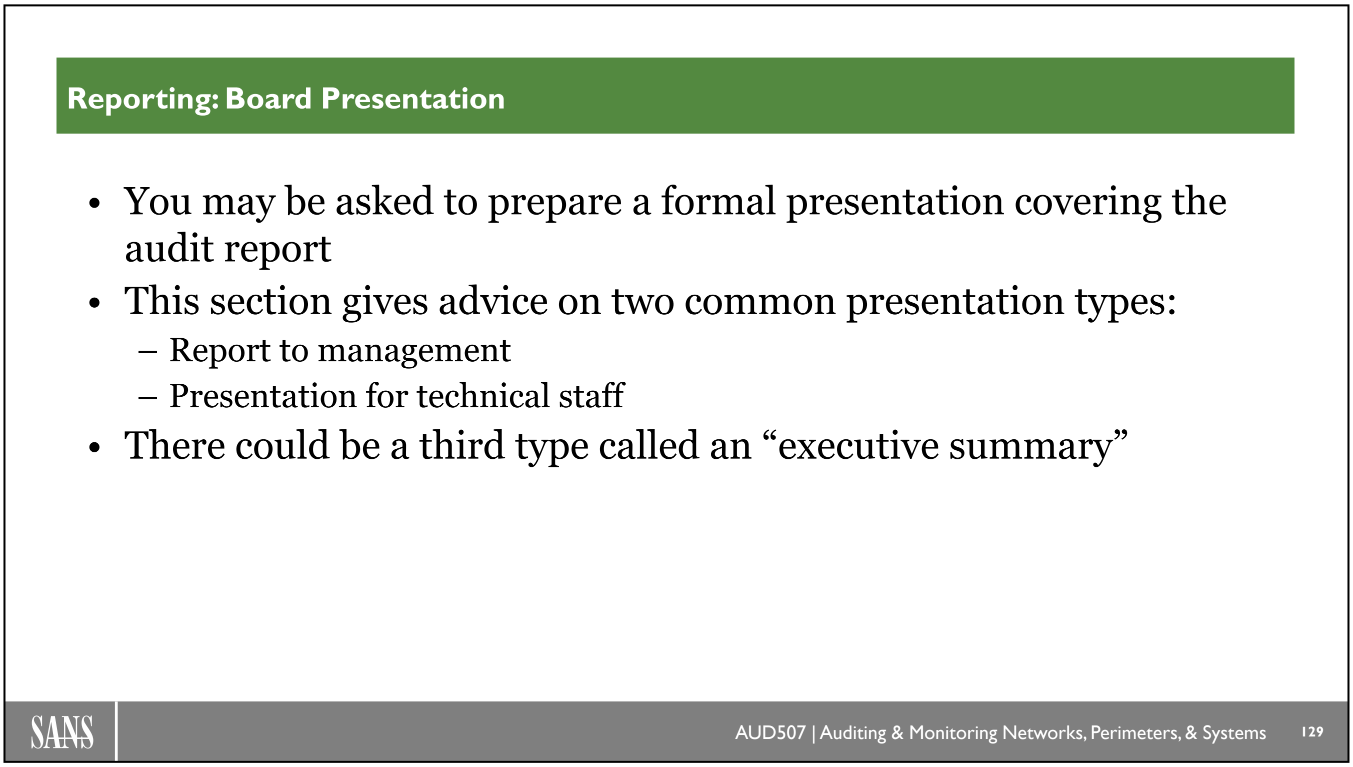

Reporting: Board Presentation

- You may be asked to prepare a formal presentation covering the audit report
- This section gives advice on two common presentation types:
  - Report to management
  - Presentation for technical staff
- There could be a third type called an "executive summary"

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 129

In addition to the written report, auditors are sometimes asked to prepare presentations to be given to the auditee after delivery of the report. These presentations most commonly fall into two areas:

- A recap of the report to management
- A more technical presentation for the auditee's technical staff

Tips for both types of presentation are given on the following slides.

As a side note, there could also be a third presentation that may be given at times called the "executive summary". This would be even more brief and to the point than a management report.

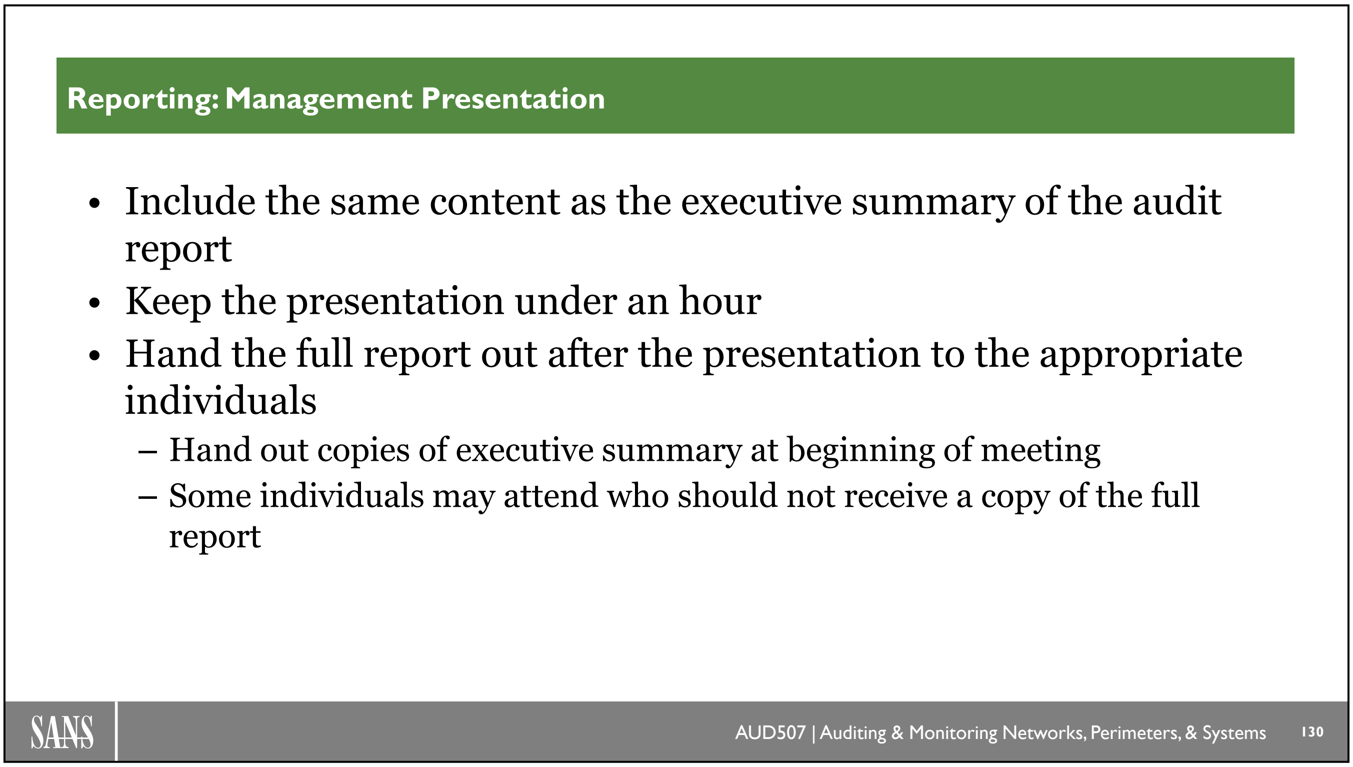

Reporting: Management Presentation

- Include the same content as the executive summary of the audit report
- Keep the presentation under an hour
- Hand the full report out after the presentation to the appropriate individuals
  - Hand out copies of executive summary at beginning of meeting
  - Some individuals may attend who should not receive a copy of the full report

The presentation given to managers should largely mirror the content in the executive summary. This presentation should be short; an hour is usually more than enough time to present the major findings and recommendations.

It's usually a good idea to hand out the full report after the presentation, when appropriate, but it might be beneficial to hand out a copy of the executive summary at the beginning of the meeting. This will give the managers something on which to write their notes or questions during the presentation. Why not hand out the full report at the beginning of the meeting? Because the managers will spend their time reading it and may lose the important points of the presentation while they are bogged down in the details.

Management presentation distribution should be discussed prior to the presentation. Individuals may attend who should not receive a copy of the full report.

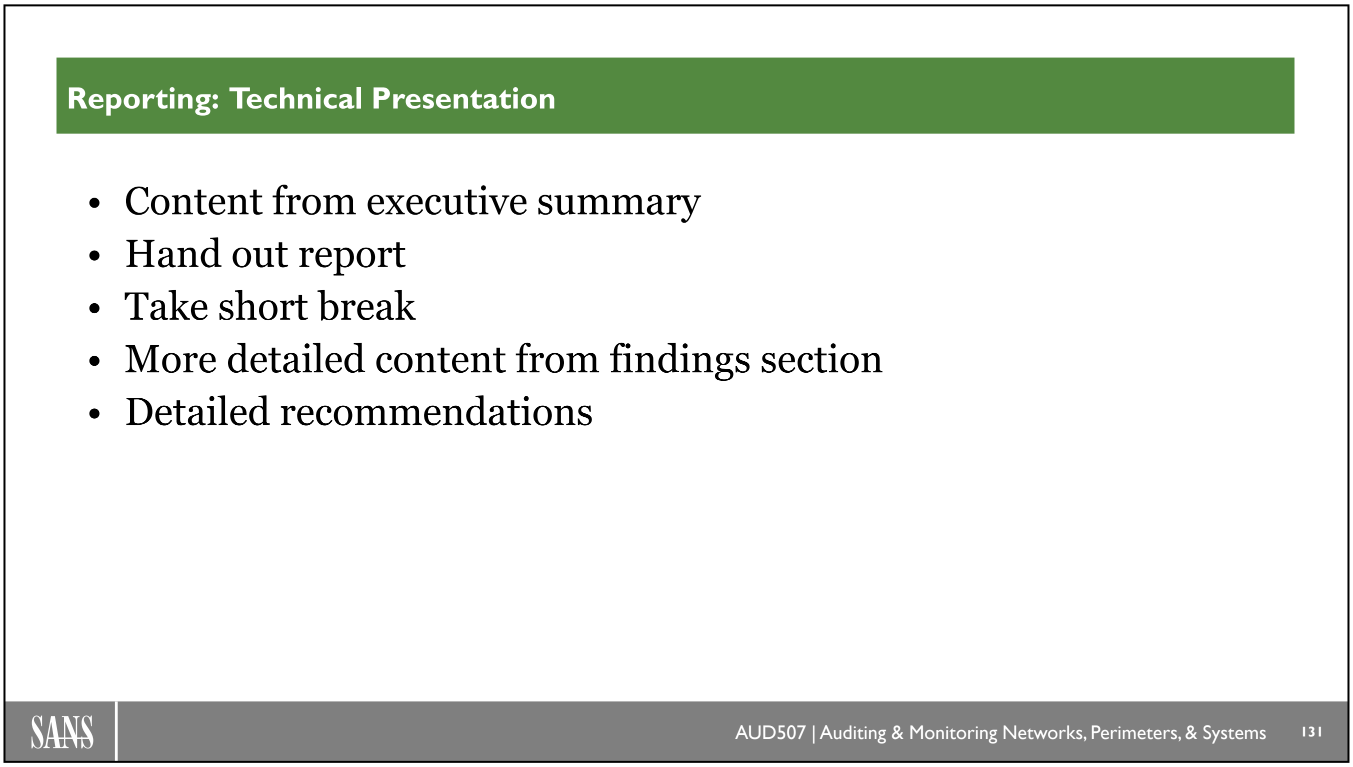

Reporting: Technical Presentation

- Content from executive summary
- Hand out report
- Take short break
- More detailed content from findings section
- Detailed recommendations

SANS | AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 131

A presentation to the technical staff should be handled differently than one made to managers. A good agenda would be to:

- Do quick presentation of items from executive summary.
- Hand out copies of the full audit report, and then take a short break to allow people to “find themselves” in the report.
- After the break, present the detailed findings one at a time, allowing staff to ask questions when appropriate.
- Finally, present recommendations for each of the findings, and again allow the staff to ask questions.

The technical staff presentation will typically take a lot longer than the presentation to management, and the auditor will need to be well-versed in the technical details of the findings.

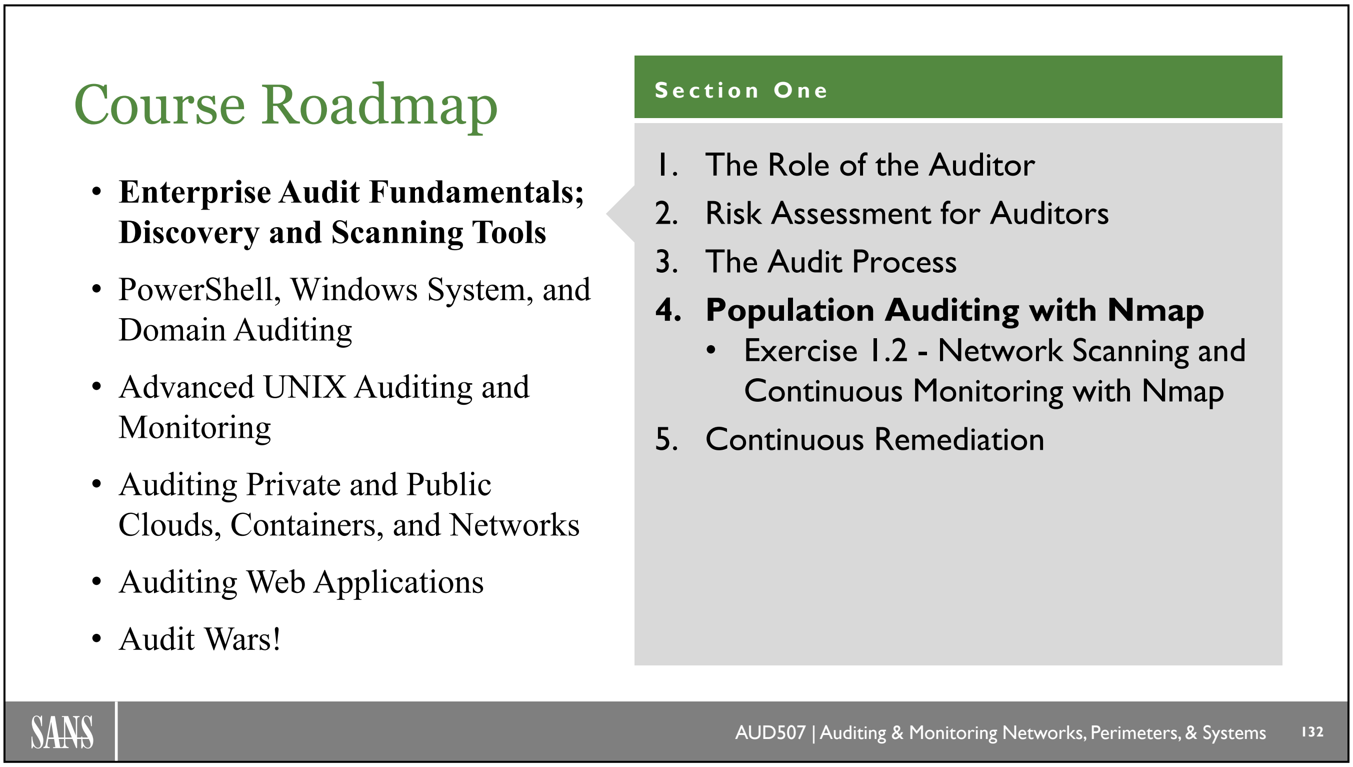

Course Roadmap

- **Enterprise Audit Fundamentals; Discovery and Scanning Tools**
- PowerShell, Windows System, and Domain Auditing
- Advanced UNIX Auditing and Monitoring
- Auditing Private and Public Clouds, Containers, and Networks
- Auditing Web Applications
- Audit Wars!

Section One

1. The Role of the Auditor
2. Risk Assessment for Auditors
3. The Audit Process
4. **Population Auditing with Nmap**
   - Exercise 1.2 - Network Scanning and Continuous Monitoring with Nmap
5. Continuous Remediation

This page intentionally left blank.

Network Maps

- Always ask for them
- Never trust them:
  - Someone's opinion of what the network might have looked like at some point in time
- Validate with scanning tools
  - Maybe before finalizing scope

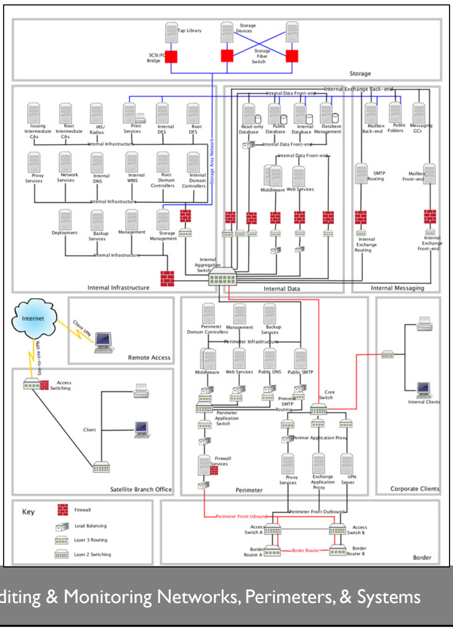

SANS

Part of the preparation work we do for our audits is to ask for a network diagram and device inventory. In our experience, these documents are often incomplete and outdated. This is usually because the diagram was created by a network engineer trying to solve a particular problem at a certain time, and the diagram represents only what was needed on that day!

This is a good opportunity for you to exercise your professional skepticism and validate that what's shown on the diagram is complete, accurate and up-to-date. In this section, we will discuss how to use scanning tools like Nmap to discover what's **REALLY** out there on the network.

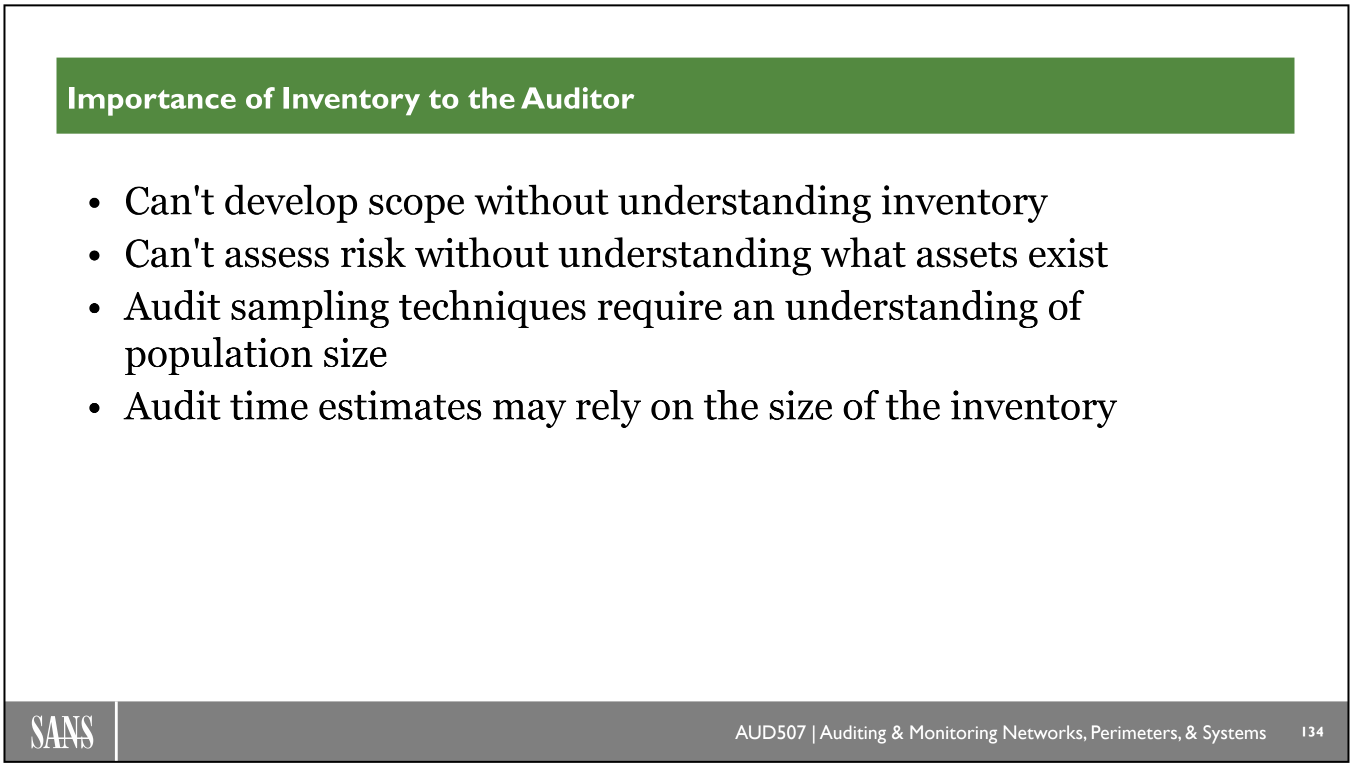

Importance of Inventory to the Auditor

- Can't develop scope without understanding inventory
- Can't assess risk without understanding what assets exist
- Audit sampling techniques require an understanding of population size
- Audit time estimates may rely on the size of the inventory

The preparation phase of auditing should include developing an understanding of the environment under review. This will often include performing discovery scans of the environment using either port scanners like Nmap or general-purpose vulnerability scanners like those from Qualys and tenable.

Just like the enterprise needs to understand its inventory before it can apply controls, the auditor needs to understand the inventory so they can assess risk and consider sampling and evidence collection techniques. It's very helpful when developing audit time and resource estimates to have an understanding of the size of the environment to be tested.

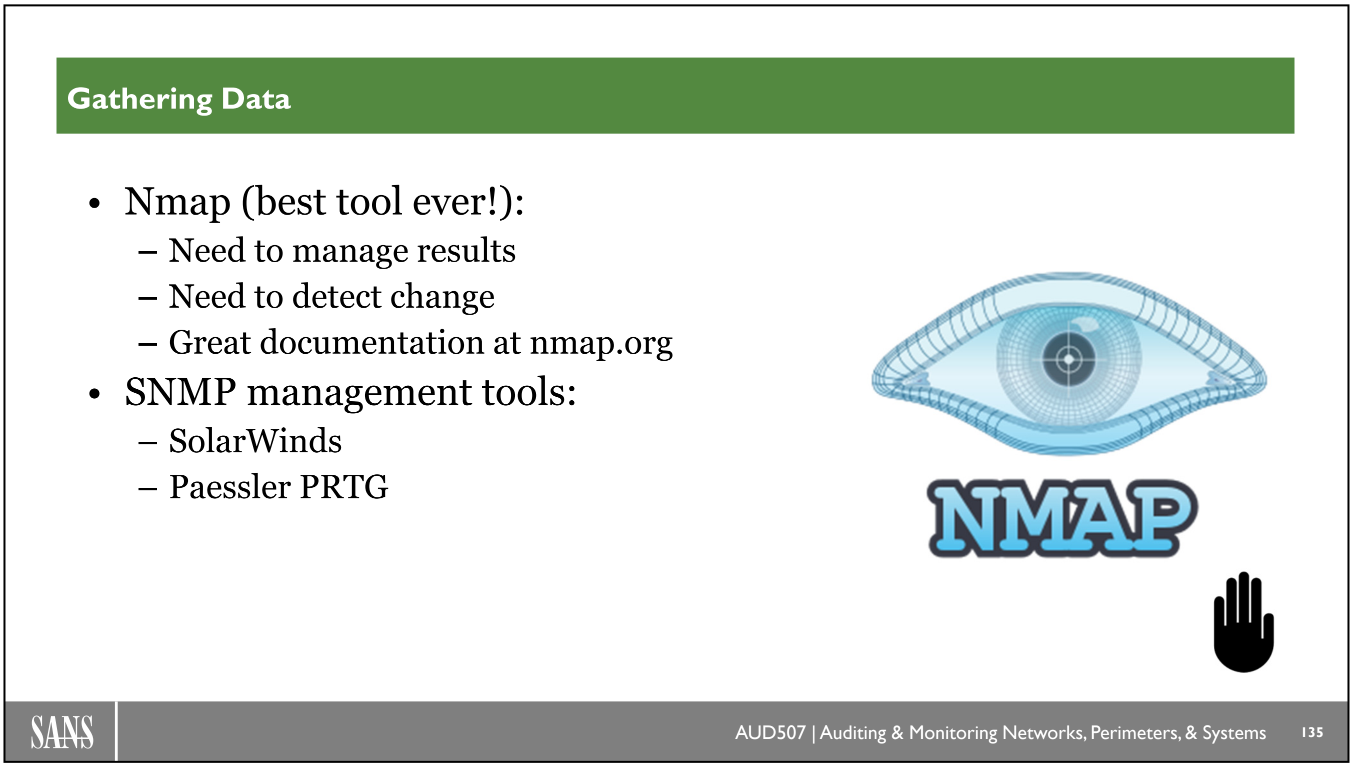

Gathering Data

- Nmap (best tool ever!):
  - Need to manage results
  - Need to detect change
  - Great documentation at nmap.org
- SNMP management tools:
  - SolarWinds
  - Paessler PRTG

How do we do the actual scanning? You can actually use any tool that you'd like to, but Nmap is a wonderful tool to get this work done. Our discussions and labs largely revolve around how to use Nmap effectively for this task.

The big challenge that you may perceive is that Nmap tends to provide verbose output and that output is fairly unwieldy because it's just text. We need to find a way to turn that into a manageable solution. We also need a way to detect changes over time.

This isn't to say that Nmap is the only way to go. For example, an extremely useful piece of information to baseline is the aggregate routing table for our internal network. There is absolutely no way for Nmap to get you this data. This isn't something you "scan"; it's something that you have to ask for from the routers. We may also be interested in tracking hosts moving around the subnets.

Many IT shops use Solar Winds commercial tools for monitoring network devices and servers. In those enterprises, you may be able to leverage that technology to gather your population data. Paessler's PRTG product is another monitoring tool, and it has a version that's free for up to 100 devices.

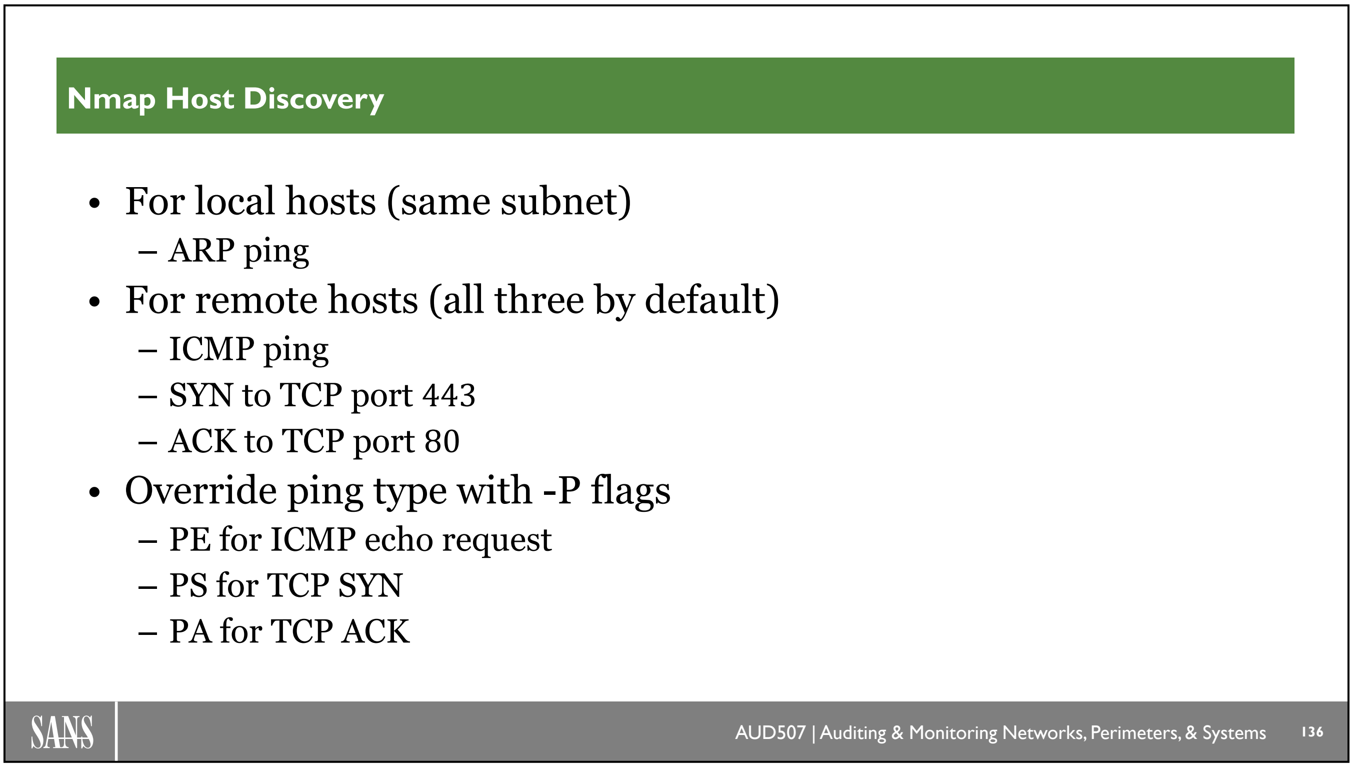

Nmap Host Discovery

- For local hosts (same subnet)
  - ARP ping
- For remote hosts (all three by default)
  - ICMP ping
  - SYN to TCP port 443
  - ACK to TCP port 80
- Override ping type with -P flags
  - PE for ICMP echo request
  - PS for TCP SYN
  - PA for TCP ACK

Let’s take a minute to explain exactly how Nmap works. Even though we have already used this tool today, the only thing that we’ve used it for is scanning the firewall. Scanning one host with Nmap is easy. Scanning a lot of hosts can be extremely time-consuming if we don’t understand how the tool actually works.

In its default configuration, Nmap begins any scan by checking to see if each target host is online. Nmap has several ways that it can determine this information, but the default is to perform an ICMP ping (echo request) and an ACK ping (to port 80). The only exception to this would be if Nmap detects that the target address is on the local subnet. If this is the case, it instead uses an ARP who has (mentioned earlier) to determine if the host is up.

You can get help understanding all of Nmap’s functions by running the command: nmap --help.

The host discovery flags listed in Nmap’s help file are given below:

```
HOST DISCOVERY:
  -sL: List Scan - simply list targets to scan
  -sn: Ping Scan - disable port scan
  -Pn: Treat all hosts as online -- skip host discovery
  -PS/PA/PU/PY[portlist]: TCP SYN/ACK, UDP or SCTP discovery to given ports
  -PE/PP/PM: ICMP echo, timestamp, and netmask request discovery probes
  -PO[protocol list]: IP Protocol Ping
  -n/-R: Never do DNS resolution/Always resolve [default: sometimes]
  --dns-servers <serv1[,serv2],...>: Specify custom DNS servers
  --system-dns: Use OS's DNS resolver
  --traceroute: Trace hop path to each host
```

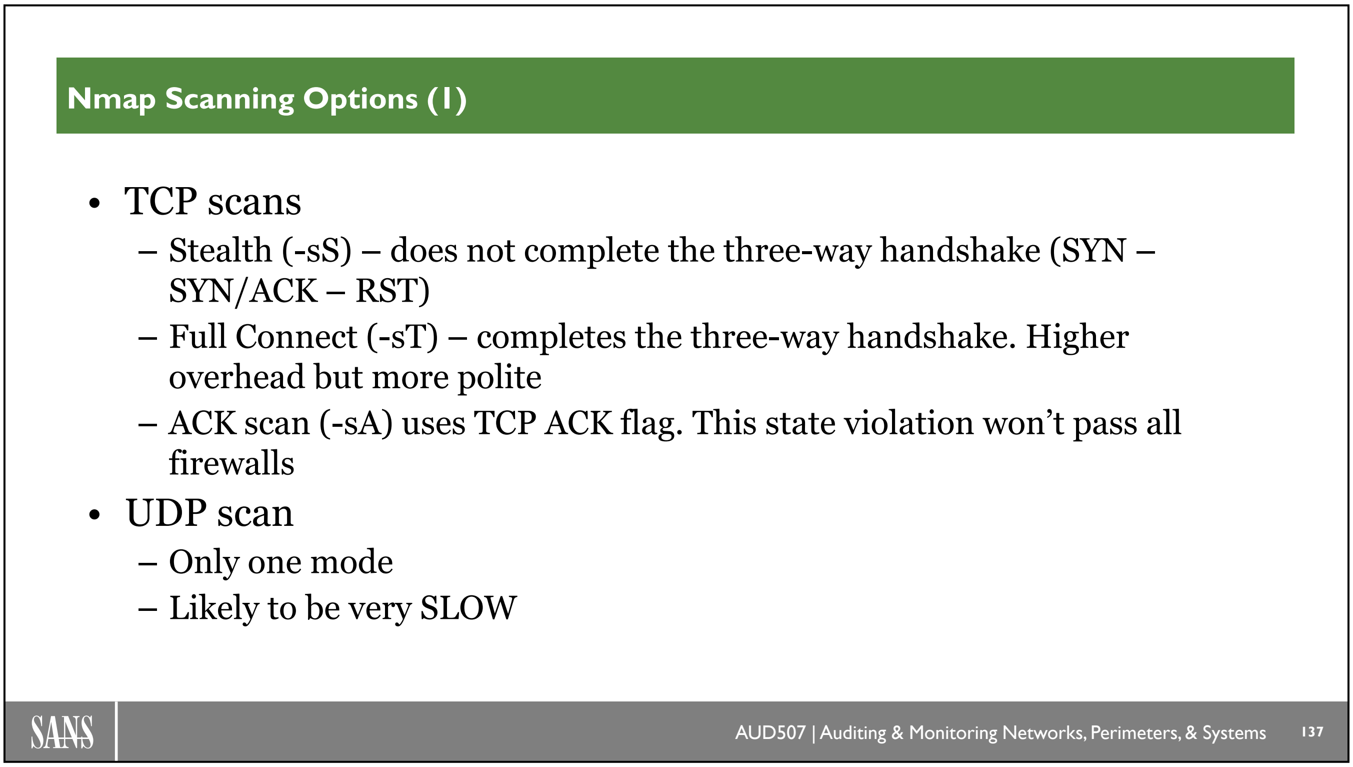

Nmap Scanning Options (1)

- TCP scans
  - Stealth (-sS) – does not complete the three-way handshake (SYN – SYN/ACK – RST)
  - Full Connect (-sT) – completes the three-way handshake. Higher overhead but more polite
  - ACK scan (-sA) uses TCP ACK flag. This state violation won't pass all firewalls
- UDP scan
  - Only one mode
  - Likely to be very SLOW

If Nmap determines that the host is online, it then begins scanning for TCP ports by default. You can, of course, specify which ports to scan for. You can also control how the scan is performed. For our purposes, the most interesting scan type to use is a full connection scan. This means that Nmap complies with the typical three-way handshake approach to connections.

This scan type is important to use, especially with legacy systems, because it is much friendlier for the target network devices. This gives us the most reliable way to scan hosts while causing the least damage. In addition, we may want to look at the speed with which the scan runs because this, too, can create issues on a fragile network.

Nmap can also perform UDP scans. For the next lab that we do, we will not perform any UDP scans. The biggest reason is that they are incredibly slow. Remember that if a UDP service is listening, it will not answer a probe. This means that we have to wait for a *lack* of response. As you can imagine, this takes a long time.

Full scanning options from the help file are below:

```
SCAN TECHNIQUES:
  -sS/sT/sA/sW/sM: TCP SYN/Connect()/ACK/Window/Maimon scans
  -sU: UDP Scan
  -sN/sF/sX: TCP Null, FIN, and Xmas scans
  --scanflags <flags>: Customize TCP scan flags
  -sI <zombie host[:probeport]>: Idle scan
  -sY/sZ: SCTP INIT/COOKIE-ECHO scans
  -sO: IP protocol scan
  -b <FTP relay host>: FTP bounce scan
```

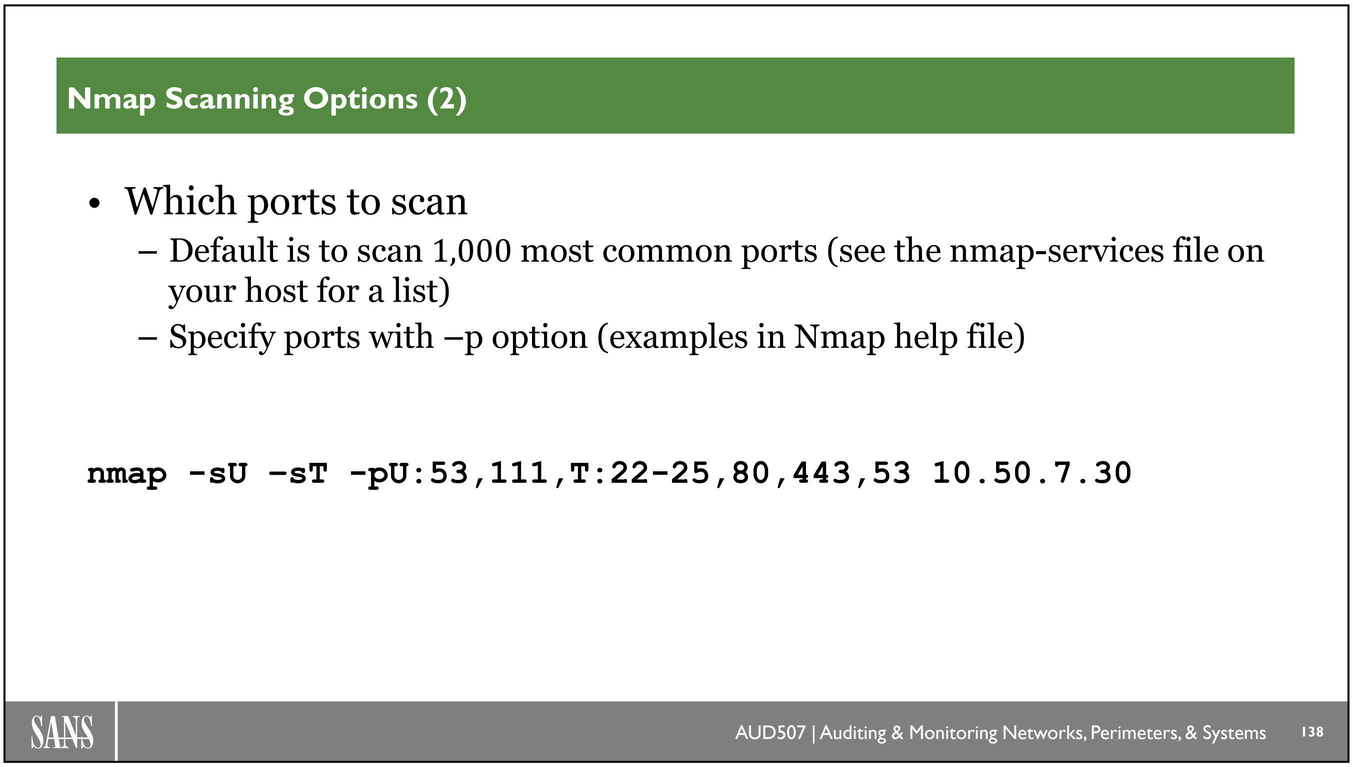

Nmap Scanning Options (2)

- Which ports to scan
  - Default is to scan 1,000 most common ports (see the nmap-services file on your host for a list)
  - Specify ports with –p option (examples in Nmap help file)

```
nmap -sU –sT -pU:53,111,T:22-25,80,443,53 10.50.7.30
```

Nmap ships with a list of the 1,000 most commonly open ports, based on years of research by the author. The full list, with some frequency information is saved in the nmap-services file. On the Debian virtual machine you'll use in most of your labs, that file is located at /usr/local/share/nmap/nmap-services.

Specifying ports in Nmap requires the use of the –p flag. Ports can be listed as ranges with a hyphen, or individually, separated by commas. To combine both TCP and UDP ports into one scan, you must use the –sU flag plus one of the TCP scan flags, like –sT.

The help section on port specification is below:

```
PORT SPECIFICATION AND SCAN ORDER:
  -p <port ranges>: Only scan specified ports
    Ex: -p22; -p1-65535; -p U:53,111,137,T:21-25,80,139,8080,S:9
  --exclude-ports <port ranges>: Exclude the specified ports from scanning
  -F: Fast mode - Scan fewer ports than the default scan
  -r: Scan ports consecutively - don't randomize
  --top-ports <number>: Scan <number> most common ports
  --port-ratio <ratio>: Scan ports more common than <ratio>
```

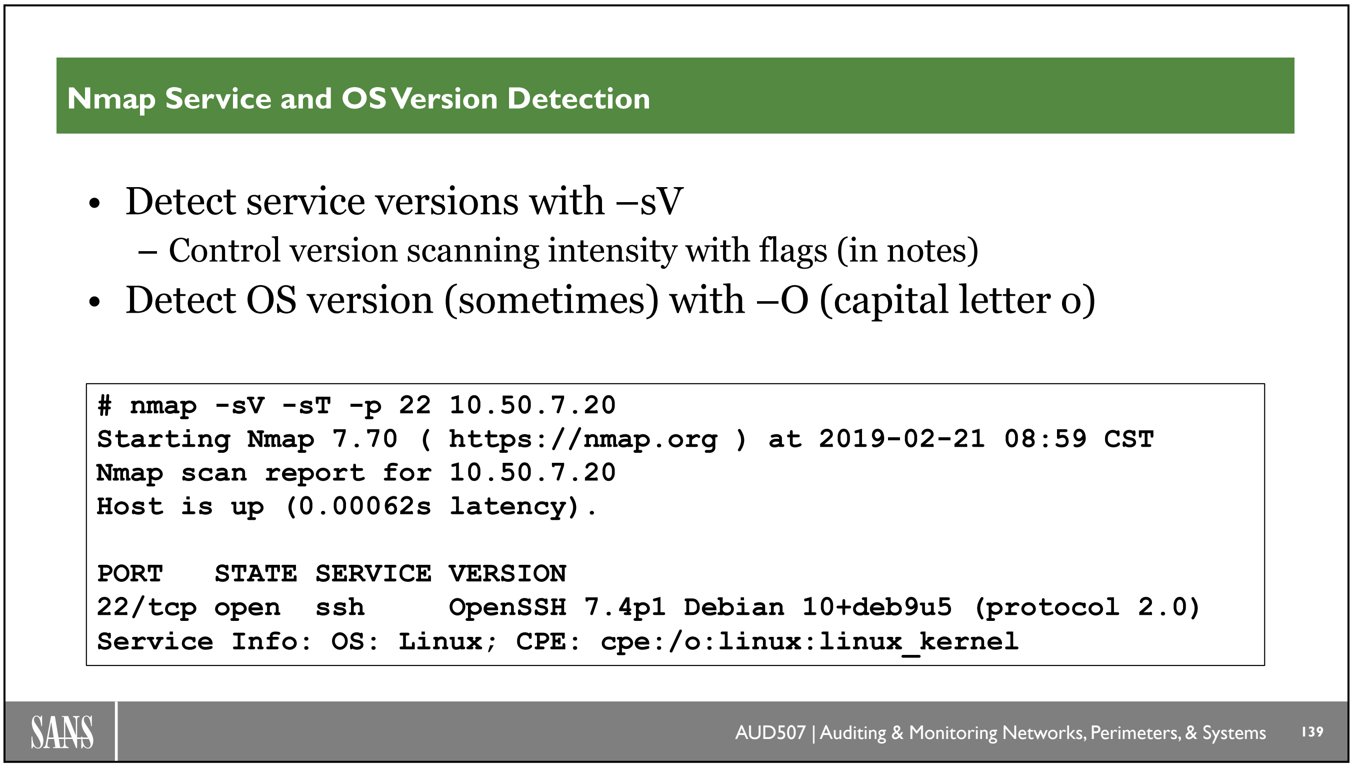

Nmap Service and OS Version Detection

- Detect service versions with –sV
  - Control version scanning intensity with flags (in notes)
- Detect OS version (sometimes) with –O (capital letter o)

```
# nmap -sV -sT -p 22 10.50.7.20
Starting Nmap 7.70 ( https://nmap.org ) at 2019-02-21 08:59 CST
Nmap scan report for 10.50.7.20
Host is up (0.00062s latency).

PORT   STATE SERVICE VERSION
22/tcp open  ssh     OpenSSH 7.4p1 Debian 10+deb9u5 (protocol 2.0)
Service Info: OS: Linux; CPE: cpe:/o:linux:linux_kernel
```

Version detection is very useful for auditors because it provides a starting point for researching how up-to-date the services on a host are. The operating system detection capabilities of Nmap are largely based on stimulating the host with very specific traffic and seeing how it responds. OS detection works best when there are a combination of open and closed ports on the host being scanned.

The relevant section of the help file is below:

```
SERVICE/VERSION DETECTION:
  -sV: Probe open ports to determine service/version info
  --version-intensity <level>: Set from 0 (light) to 9 (try all probes)
  --version-light: Limit to most likely probes (intensity 2)
  --version-all: Try every single probe (intensity 9)
  --version-trace: Show detailed version scan activity (for debugging)
```

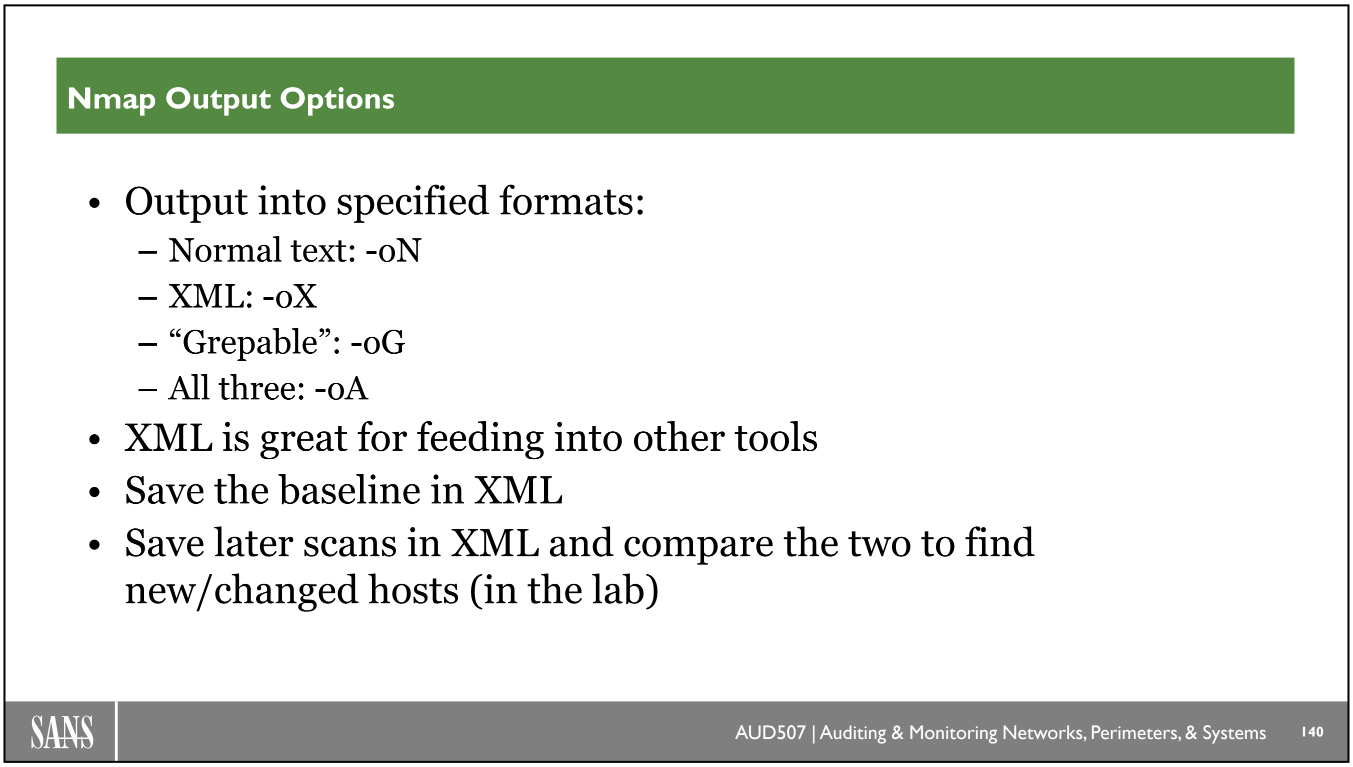

Nmap Output Options

- Output into specified formats:
  - Normal text: -oN
  - XML: -oX
  - “Grepable”: -oG
  - All three: -oA
- XML is great for feeding into other tools
- Save the baseline in XML
- Save later scans in XML and compare the two to find new/changed hosts (in the lab)

Nmap allows the user lots of control over the output from the scanner. One of my favorites as an auditor is the ability to save the results to files for later analysis. The XML output format is especially useful when I need to import the scan results into another tool. I often use Nmap to perform the initial scan of the network, and then import its results into my vulnerability scanning tool for later testing.

Nmap’s output control options are listed below:

```
OUTPUT:
  -oN/-oX/-oS/-oG <file>: Output scan in normal, XML, s|<rIpt kIddi3,
     and Grepable format, respectively, to the given filename.
  -oA <basename>: Output in the three major formats at once
  -v: Increase verbosity level (use -vv or more for greater effect)
  -d: Increase debugging level (use -dd or more for greater effect)
  --reason: Display the reason a port is in a particular state
  --open: Only show open (or possibly open) ports
  --packet-trace: Show all packets sent and received
  --iflist: Print host interfaces and routes (for debugging)
  --append-output: Append to rather than clobber specified output files
  --resume <filename>: Resume an aborted scan
  --stylesheet <path/URL>: XSL stylesheet to transform XML output to HTML
  --webxml: Reference stylesheet from Nmap.Org for more portable XML
  --no-stylesheet: Prevent associating of XSL stylesheet w/XML output
```

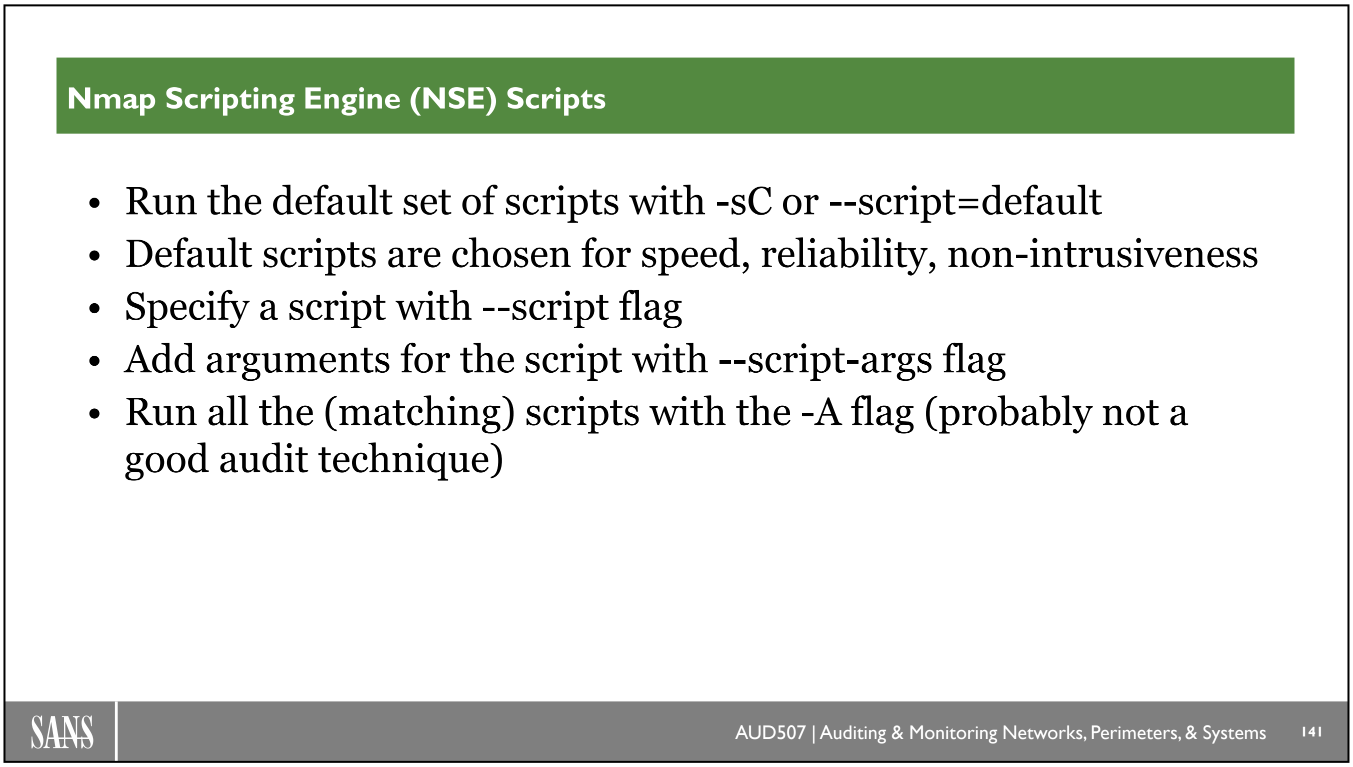

Nmap Scripting Engine (NSE) Scripts

- Run the default set of scripts with -sC or --script=default
- Default scripts are chosen for speed, reliability, non-intrusiveness
- Specify a script with --script flag
- Add arguments for the script with --script-args flag
- Run all the (matching) scripts with the -A flag (probably not a good audit technique)

Nmap may have started as a simple port scanner, but with the advent of scripts, it has become a multi-purpose audit chainsaw with very wide functionality. Nmap scripts exist to do everything from inventory the TLS ciphers on a web server (we'll do that in the web application section) to brute forcing credentials on an SMTP server. Nmap currently ships with well over 500 scripts. On the Debian VM you'll use in today's labs, the scripts are stored in the /usr/local/share/nmap/scripts/ directory. The scripts usually include usage instructions in the source code, so it can be worth reading the code for scripts which look like they could help with an audit task.

The script section of the help file is below:

```
SCRIPT SCAN:
  -sC: equivalent to --script=default
  --script=<Lua scripts>: <Lua scripts> is a comma separated list of
           directories, script-files or script-categories
  --script-args=<n1=v1,[n2=v2,...]>: provide arguments to scripts
  --script-args-file=filename: provide NSE script args in a file
  --script-trace: Show all data sent and received
  --script-updatedb: Update the script database.
  --script-help=<Lua scripts>: Show help about scripts.
           <Lua scripts> is a comma-separated list of script-files or
           script-categories.
```

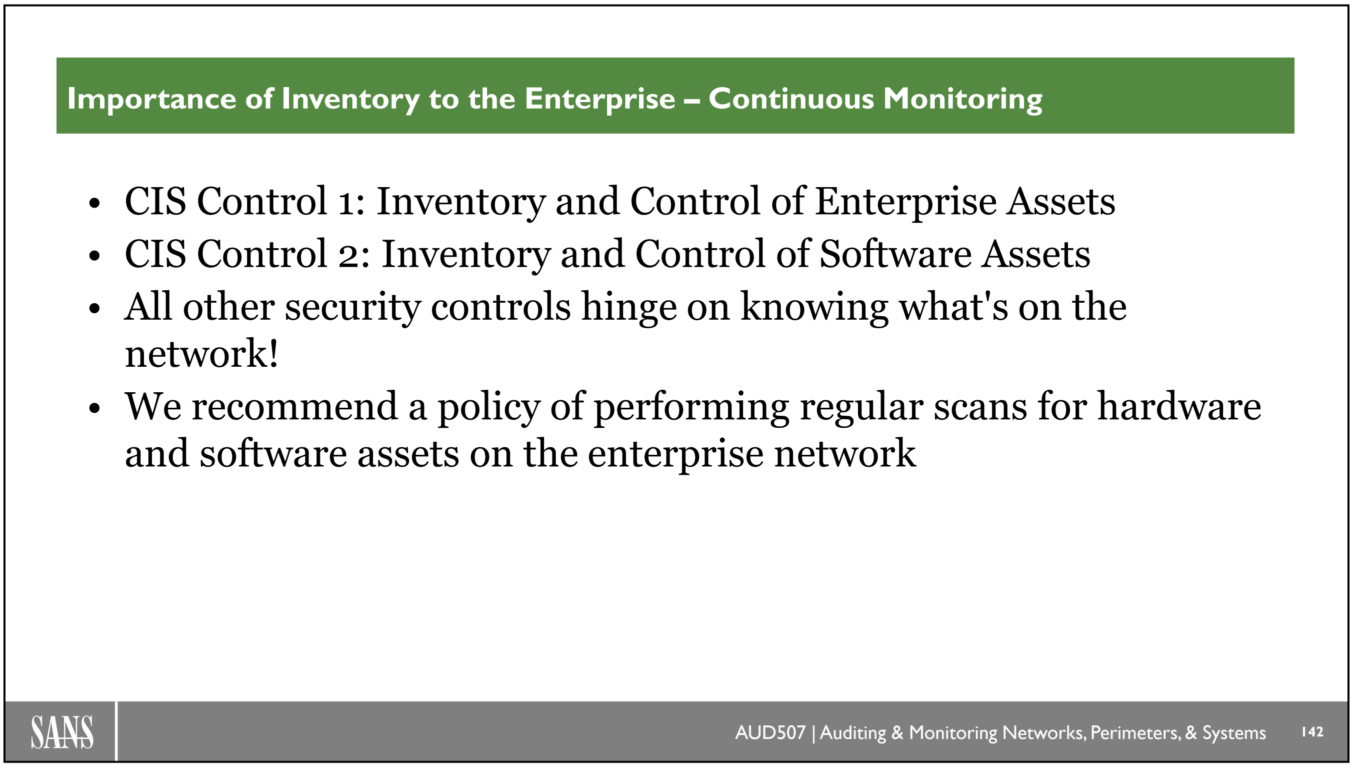

Importance of Inventory to the Enterprise – Continuous Monitoring

- CIS Control 1: Inventory and Control of Enterprise Assets
- CIS Control 2: Inventory and Control of Software Assets
- All other security controls hinge on knowing what's on the network!
- We recommend a policy of performing regular scans for hardware and software assets on the enterprise network

The first two CIS Critical Security Controls (which should be implemented in order) cover inventory and control of enterprise hardware and software assets. Scanning tools can be an excellent source of information for creating and validating these inventories. The reason that these controls occur first is that it's impossible to secure an enterprise or enclave if you don't understand what assets exist in the environment.

Developing and maintaining this inventory is really not the job of the auditor, but we can encourage the practice in the enterprises which we audit!

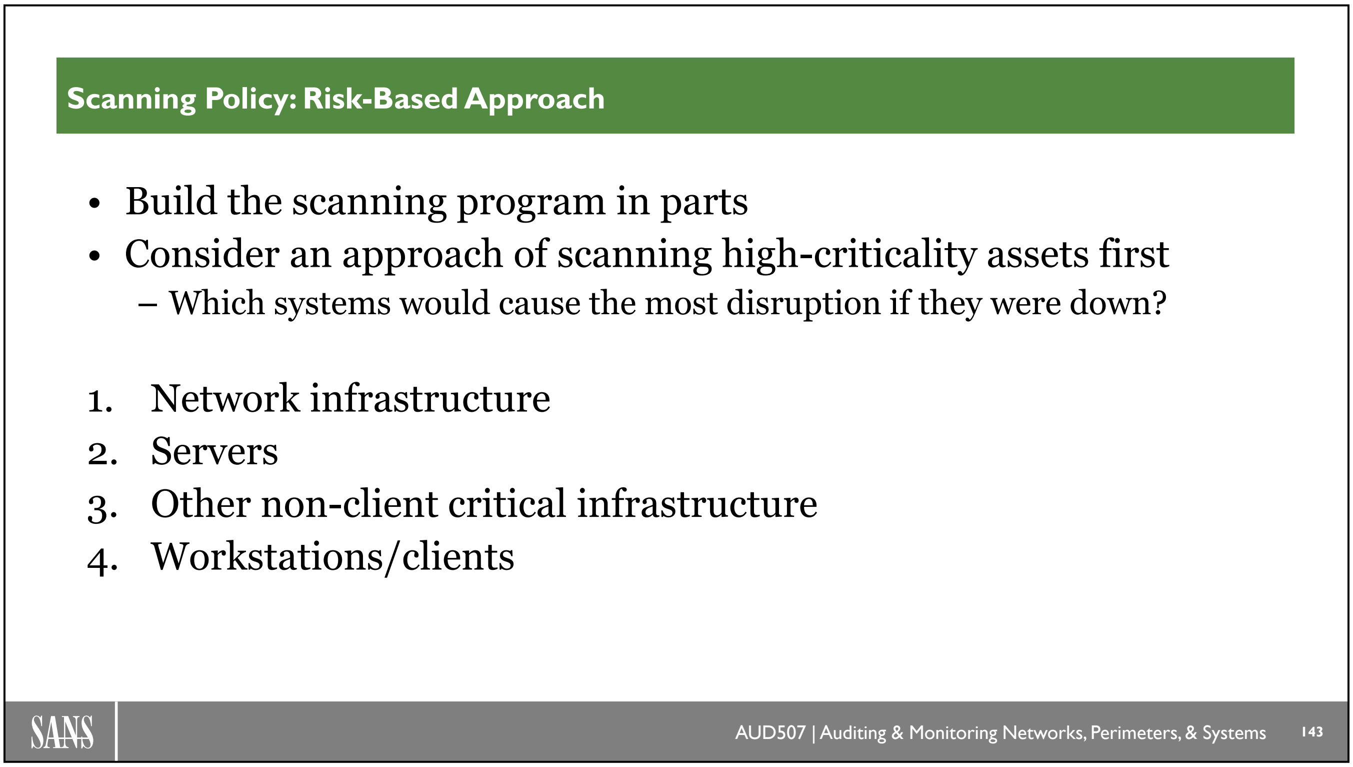

Scanning Policy: Risk-Based Approach

- Build the scanning program in parts
- Consider an approach of scanning high-criticality assets first
  - Which systems would cause the most disruption if they were down?

1. Network infrastructure
2. Servers
3. Other non-client critical infrastructure
4. Workstations/clients

When approaching this problem, whether we do so as an auditor or as an administrator of some kind, we'd advise that you take a risk-based approach to the problem. When we think about which systems have the most critical impact on our ability to operate, we usually see that the network infrastructure and servers are the first items on our list.

This may be counterintuitive. If you're thinking of starting to do network scanning, you may be tempted to start with client systems. The reasoning is that this will likely cause the least harm.

Although this is certainly a valid point of view, it is likely also true that these systems tend to provide you the least value. In addition, there are far more systems when we think of client systems. It will likely take us a long time to get to a point where we can reliably scan our client systems without any side effects. This, in turn, means that it will take a long time for us to get around to scanning our critical infrastructure.

Instead, starting with the servers, routers, and switches allows us to start with the most critical systems. It also means that we will be dealing with a much smaller number of systems. This allows us to move to the later phases quicker.

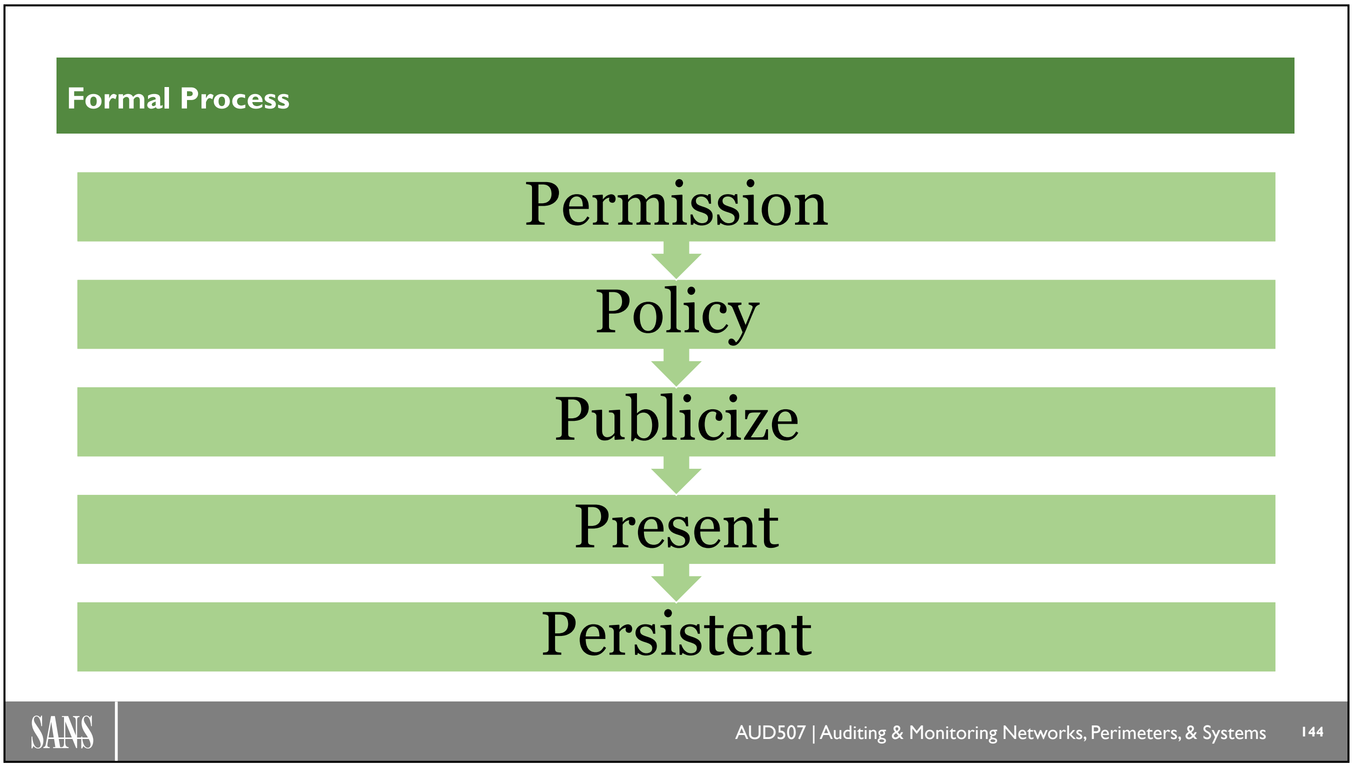

Formal Process

Permission

Policy

Publicize

Present

Persistent

It is also important to have an overall process to follow. We recommend that you follow these five steps.

First, obtain proper permission. Of course, you want the administrators to perform these tasks, so it is these administrators who would need to obtain permission. Whether it is you or them, always make sure that you obtain permission from individuals who actually have the authority to grant that permission. Also, it's best to get that permission in writing.

One way to get permission in writing is to formalize a process. The process would include who can scan, what can be scanned, how frequently, and with which tools; or it may possibly detail the intensity and type of scan.

Next, we would like to publicize the scan. You may at first think that publicizing is a bad idea. After all, if you tell people you're going to scan, they'll turn the stuff off that they're not supposed to have! Just hold your horses. We'll deal with that in a later step. We want people to know the scan is happening so that if the scan starts interfering with operations, people will realize that the culprit might be the scan.

When people realize that the scan is causing issues, they need a way to get in touch with us. We should remain available the entire time that the scan is running, and perhaps even several hours afterward. We should include contact information including our landline, cell phone, and email address. It would be terrible for the mail server to go down during the scan, preventing people from contacting us if that is the only mechanism that we provided!

Finally, we want to be persistent. This means that we will not be satisfied with one scan. Instead, when a subnet can be scanned without any harm, we will stop announcing the scan. The scan will become a periodic and automatic process. This allows us to discover anything that may have been hiding from us during step 3.

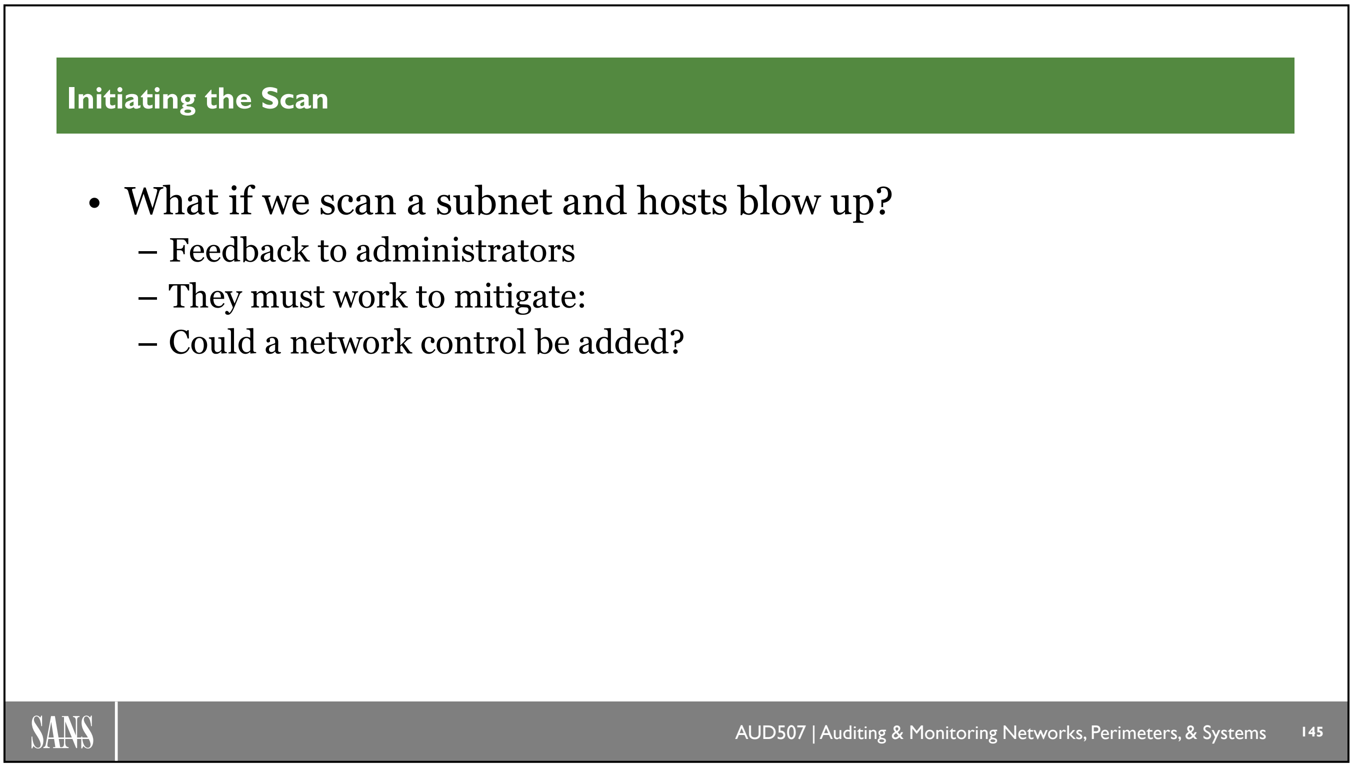

Initiating the Scan

- What if we scan a subnet and hosts blow up?
  - Feedback to administrators
  - They must work to mitigate:
  - Could a network control be added?

What do you do if you run a scan and it crashes hosts or services? We absolutely need to make the scan results available to the administrators who are responsible for the systems that were scanned. Those administrators must take the data and verify first that all the enabled services are actually required. Next, if a service is required and fails for some reason during the scan, the administrator is responsible for patching the host or otherwise fixing the issue.

What if the system in question cannot be patched? Perhaps you have a legacy system, and no patches are available, yet the system is critical to your operations.

Consider whether it's possible to introduce an additional network control to mitigate the problem. Let's go with an extremely simple solution. Could it be as easy as installing a $30 Linksys router onto the network between the legacy host and the rest of the network? This Linksys device can be configured to perform port forwarding for any required service, but everything else on that legacy host is now completely insulated from our network and, as a result, from our scan.

The whole point here is that we *must* scan the network. It is simply not acceptable for an administrator to request a long-term exception to scanning his system if it is connected to the network.

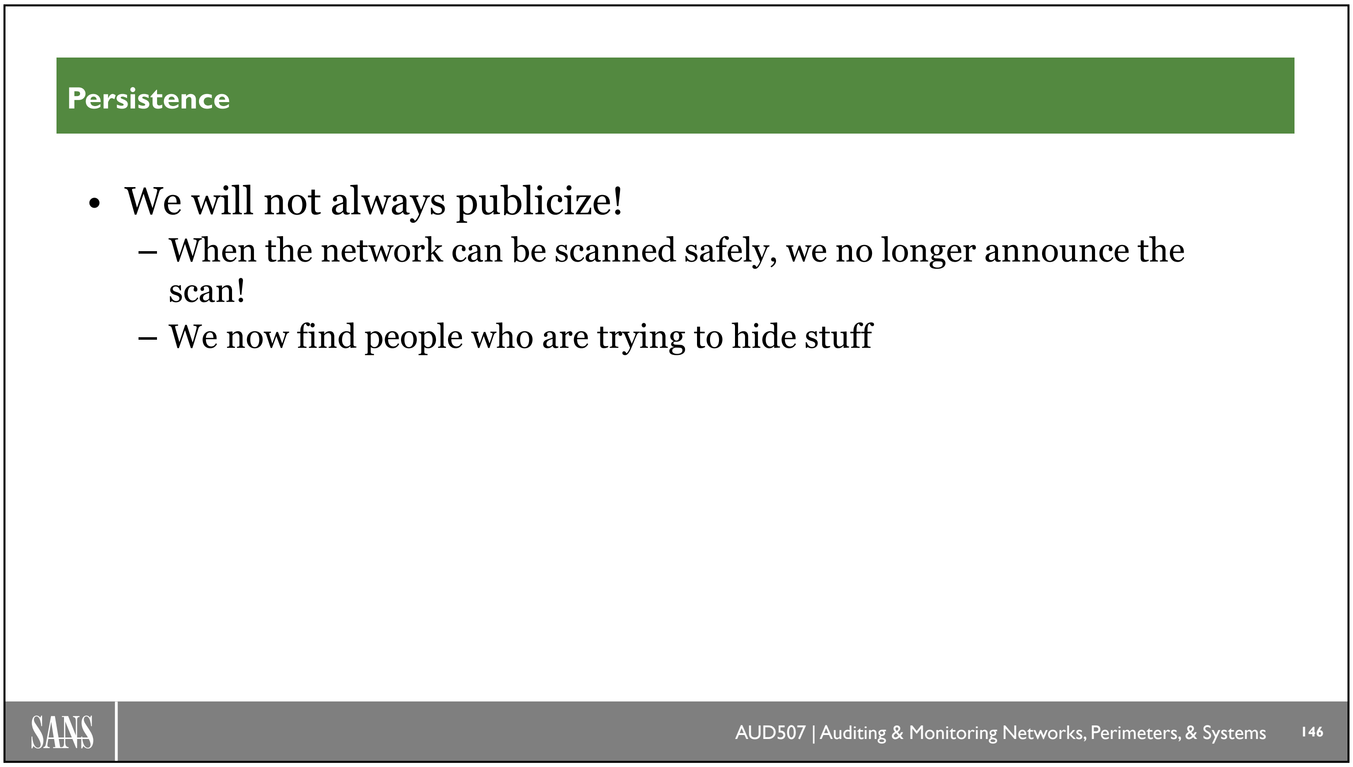

Persistence

- We will not always publicize!
  - When the network can be scanned safely, we no longer announce the scan!
  - We now find people who are trying to hide stuff

SANS
AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems
146

After we get to the point where we can scan each network with impunity, we are now ready to worry about that persistence element.

This raises the question of how often the network should be scanned. The answer depends on the size of your network and what you scan for. Under ideal circumstances, you need the scan to run every night and scan your entire address space. If your network is only a few thousand nodes, you can likely scan through this entire space every night with no problems.

If your network is significantly larger than a few thousand nodes, you need to break the scan into several pieces. Even so, think carefully about how you do this. I would still prefer to find that the critical systems and routing/switching infrastructure are scanned as frequently as possible.

The other thing that can stretch this out is UDP. If we choose to include UDP in our baseline, each scan can take an extremely long time. In most enterprises in which we have chosen to include UDP in our scanning system, we typically let the scanner run continuously, allowing it to work nonstop. When we take this approach, we also tune the scanner using the timing options so that it is not trying to scan too quickly, leading to other network issues.

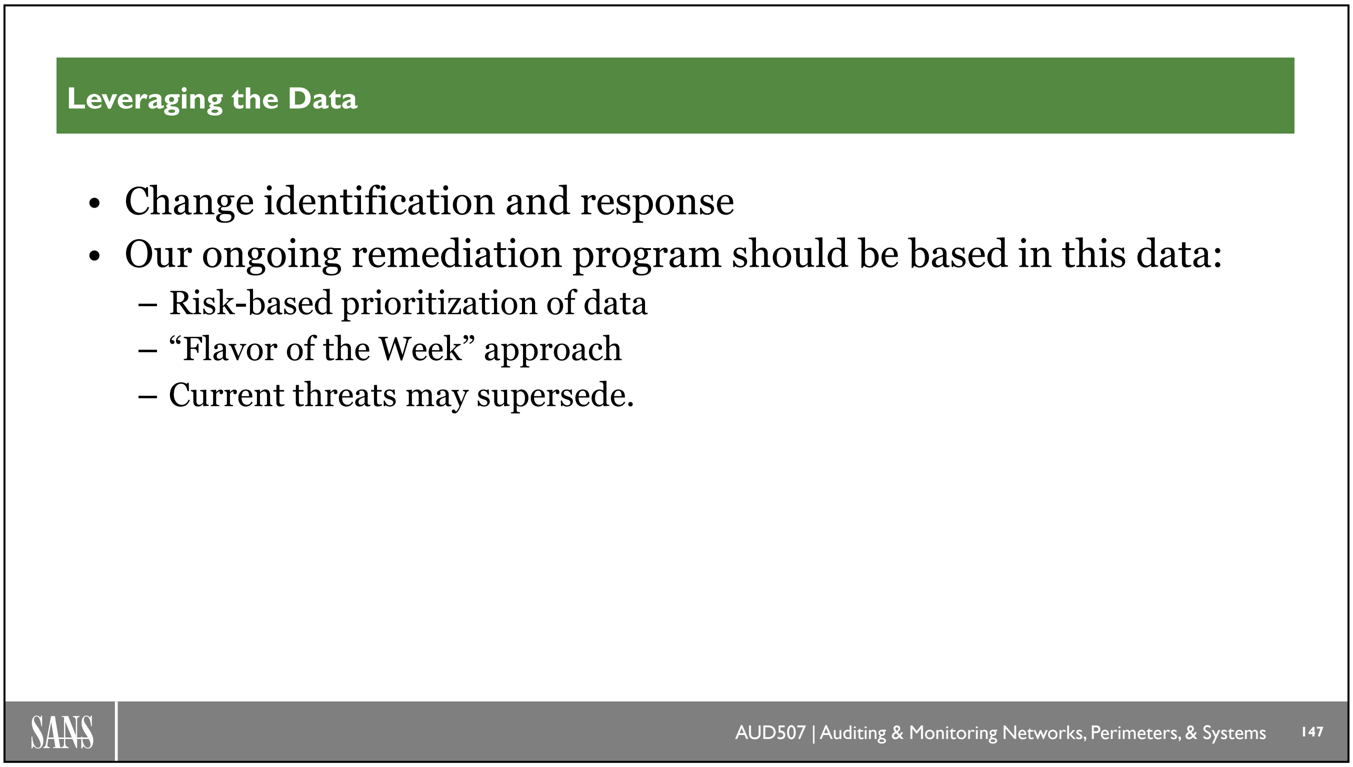

Leveraging the Data

- Change identification and response
- Our ongoing remediation program should be based in this data:
  - Risk-based prioritization of data
  - “Flavor of the Week” approach
  - Current threats may supersede.

Now that we have all this data, what do we do with it? Well, network engineers and security analysts would likely want to monitor this data daily for change. We’ll look at an extremely easy-to-understand reporting tool for this, yandiff, in our lab.

We also want to see how the security and administration teams handle continuous remediation and monitoring. I expect to find that the security officer is overseeing and driving this process using a “Flavor of the Week” approach: Each Monday morning when the security officer comes to work, he begins his day by reading about all the threats discovered over the weekend or last week. If there hasn’t been anything particularly important for his enterprise, he consults a prioritized list of services within his organization that he is working through. Where does this list come from? It is an artifact of the network scanning! He doesn’t have to wonder; he *knows* what’s running.

After one of these ports is selected, he then leverages his vulnerability scanner to look for issues with that service on all the hosts that, according to the network mapping system, have that service running. This report is then sent to the relevant administrators, asking them to determine if the service is necessary, and to either disable it or patch it.

After checking in once or twice during the week, if he has not heard back from the administrator and the issue has not been resolved by Friday, the next step is to terminate service to the offending node. Of course, if the administrator is out on leave, the security officer would hold off acting until the administrator has an opportunity to resolve the issue. As an auditor, if I found a security team following this process, I would be extremely satisfied with its continuous monitoring and remediation program!

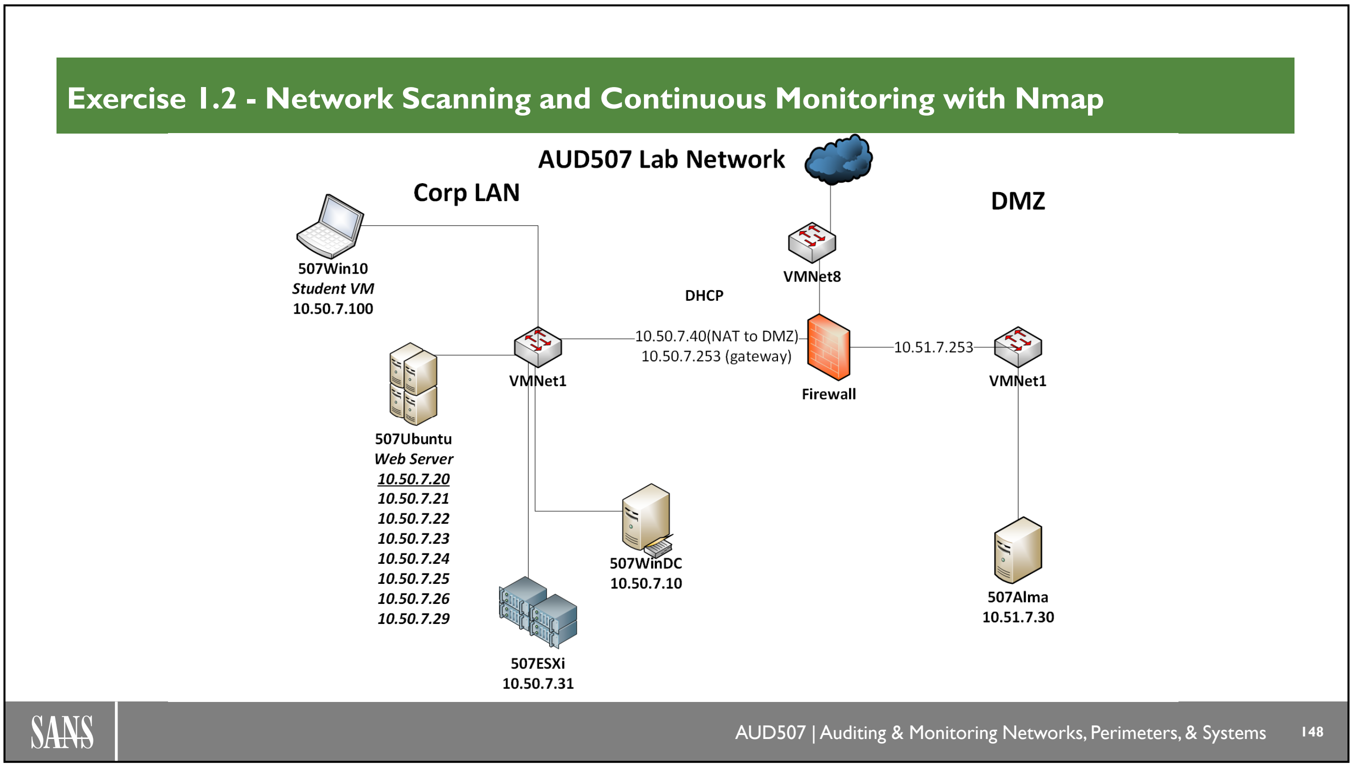

Exercise 1.2 - Network Scanning and Continuous Monitoring with Nmap

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 148

This page intentionally left blank.

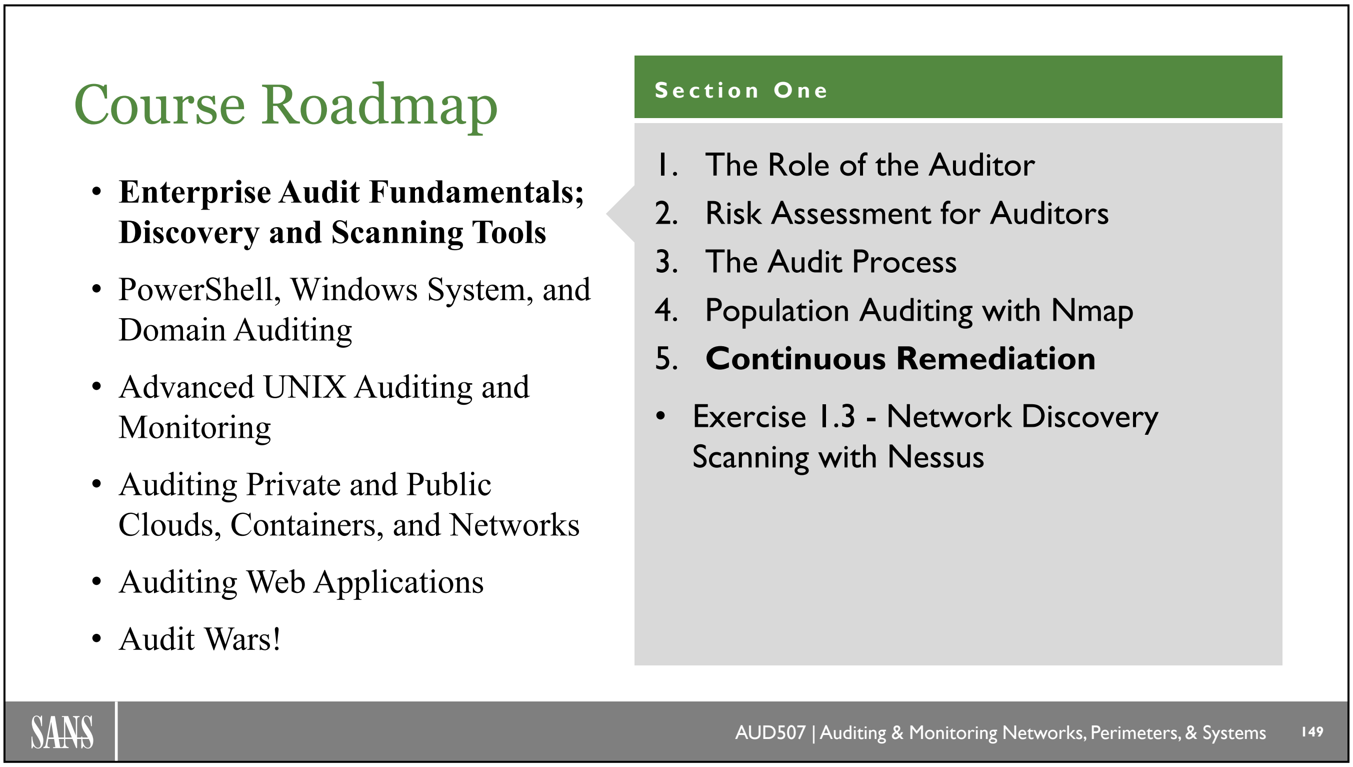

Course Roadmap

- **Enterprise Audit Fundamentals; Discovery and Scanning Tools**
- PowerShell, Windows System, and Domain Auditing
- Advanced UNIX Auditing and Monitoring
- Auditing Private and Public Clouds, Containers, and Networks
- Auditing Web Applications
- Audit Wars!

Section One

1. The Role of the Auditor
2. Risk Assessment for Auditors
3. The Audit Process
4. Population Auditing with Nmap
5. **Continuous Remediation**

- Exercise 1.3 - Network Discovery Scanning with Nessus

This page intentionally left blank.

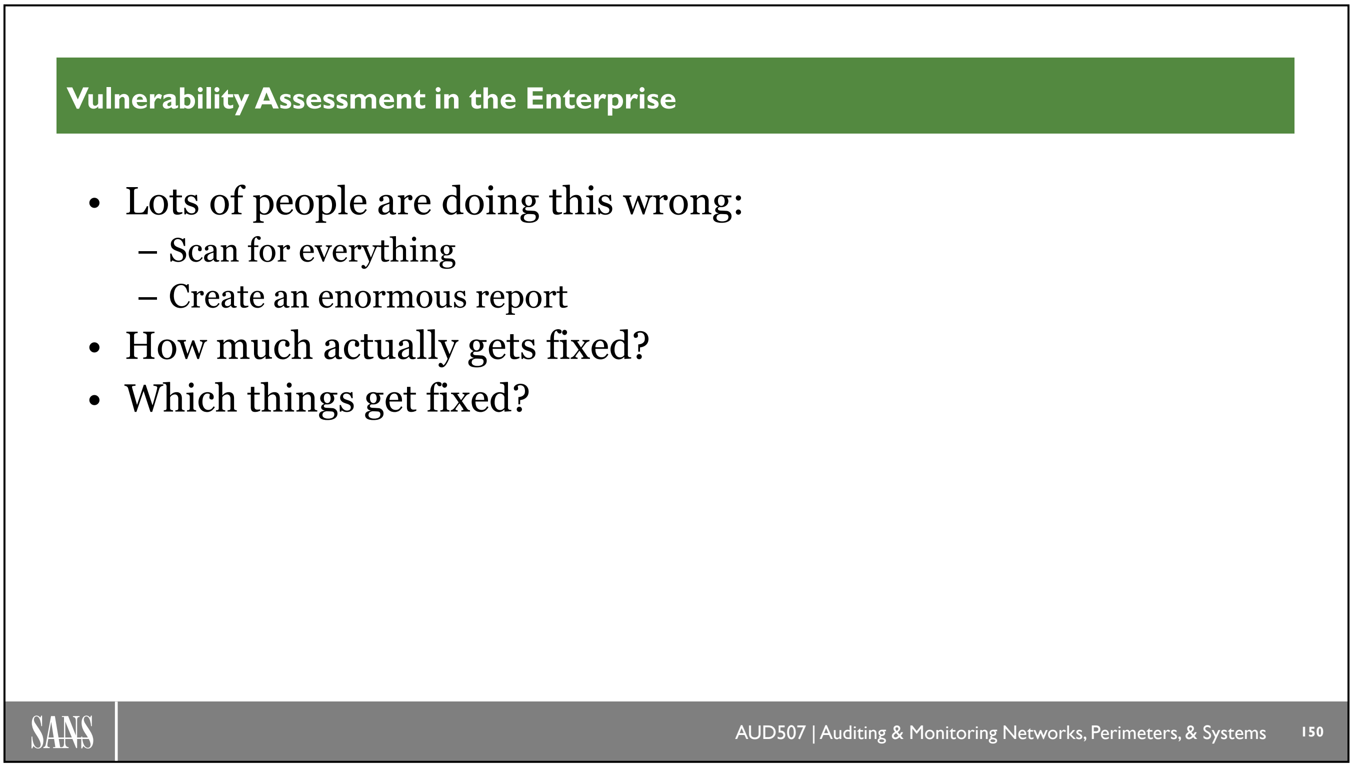

Vulnerability Assessment in the Enterprise

- Lots of people are doing this wrong:
  - Scan for everything
  - Create an enormous report
- How much actually gets fixed?
- Which things get fixed?

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 150

Vulnerability assessment is an important activity, but we find that most organizations are just plain doing it wrong. This may sound puzzling because the tools are not terribly hard to use or configure. Think about this question, though: If you have ever run a scan to find everything that's wrong, be honest… how much of that actually got fixed?

When management asks to have a vulnerability scan performed, there is almost always some kind of initiating event that sensitized them to the need for a scan. However, the more time that passes between the initiating event and the ongoing scans, the less investment management is willing to make because they are far less sensitive to the issues.

You can understand this, but how can we communicate this effectively to management? How can we help them get the kind of scan that will actually give them what they want and need?

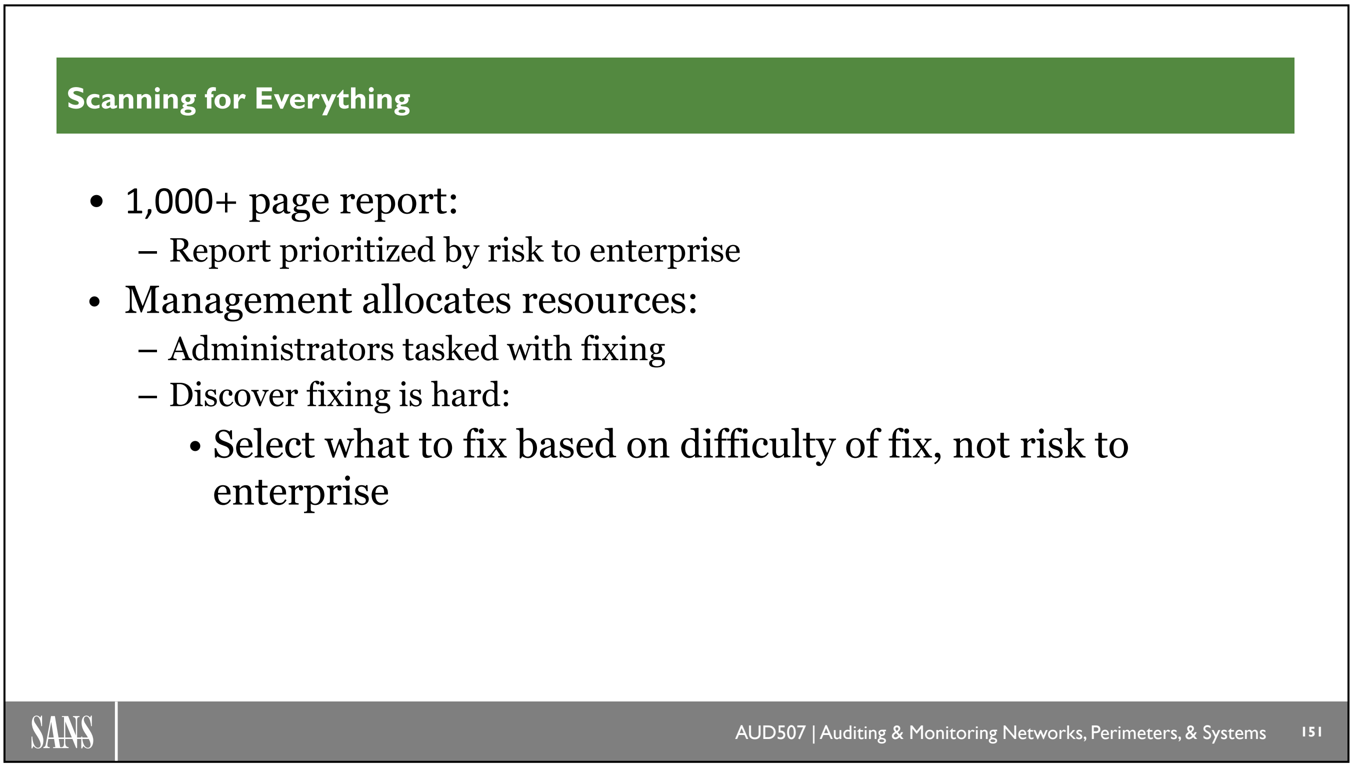

Scanning for Everything

- 1,000+ page report:
  - Report prioritized by risk to enterprise
- Management allocates resources:
  - Administrators tasked with fixing
  - Discover fixing is hard:
    - Select what to fix based on difficulty of fix, not risk to enterprise

I usually approach this by discussing it frankly with management. When management asks me, “Can you do a vulnerability scan for us?” I typically respond by saying, “Sure. What is it that you’d like me to find?”

Because of my question, the response I usually get is, “Everything…?” My question has made management unsure, but that’s okay! I now follow this up by explaining what will happen if I scan for everything.

After the scan is complete, I’ll generate a report. That report will, for argument’s sake, be 1,000 pages. You may think that’s big, but if you’ve done this before and if your network is of any reasonable size, you can attest to the accuracy of my claim. In any event, I explain to management that I will organize that report based on risk and then I will come in and scare them. I’ll scare management so badly that it will tell the administrators that their one and only job is to fix everything in the report.

The administrators don’t mind this because they recognize that there are a lot of things that they’d love to fix, they just haven’t had the time. The administrators now begin fixing issues. They discover quickly that fixing things is actually a lot of work. In fact, after fixing one or two issues, they begin leafing through the report. Rather than looking to fix things based on how much risk they represent, they are now choosing what to fix based on how hard it appears to be to fix.

As time goes on, their energy wanes. In addition, management’s attention, too, becomes unfocused. Within a few weeks, everything is back to business as usual.

Time Passes

- Within a month or so, no more fixes:
  - Remediation follow-up report is 1,100 pages:
    - New vulnerabilities
    - Fixing things uncovered other flaws
  - Nothing bad has happened
  - No more resources

Now I come back to the enterprise a few months later and scan the network again. This time the report isn't 1,000 pages long, it's 1,100 pages long. "How could it be bigger?" they wonder. "We fixed things, didn't we?"

Although it is true that they have fixed things, it is also a fact that security is not a stable system. Not only have new vulnerabilities been discovered but in the process of fixing things, previously concealed flaws have now been exposed to the light of day.

When I try to talk to management about my findings, they may pay lip service to how serious this is, but they will not actually devote any resources to it. Why? Because they see no return on investment, and nothing bad seems to have happened.

Clearly this process doesn't work.

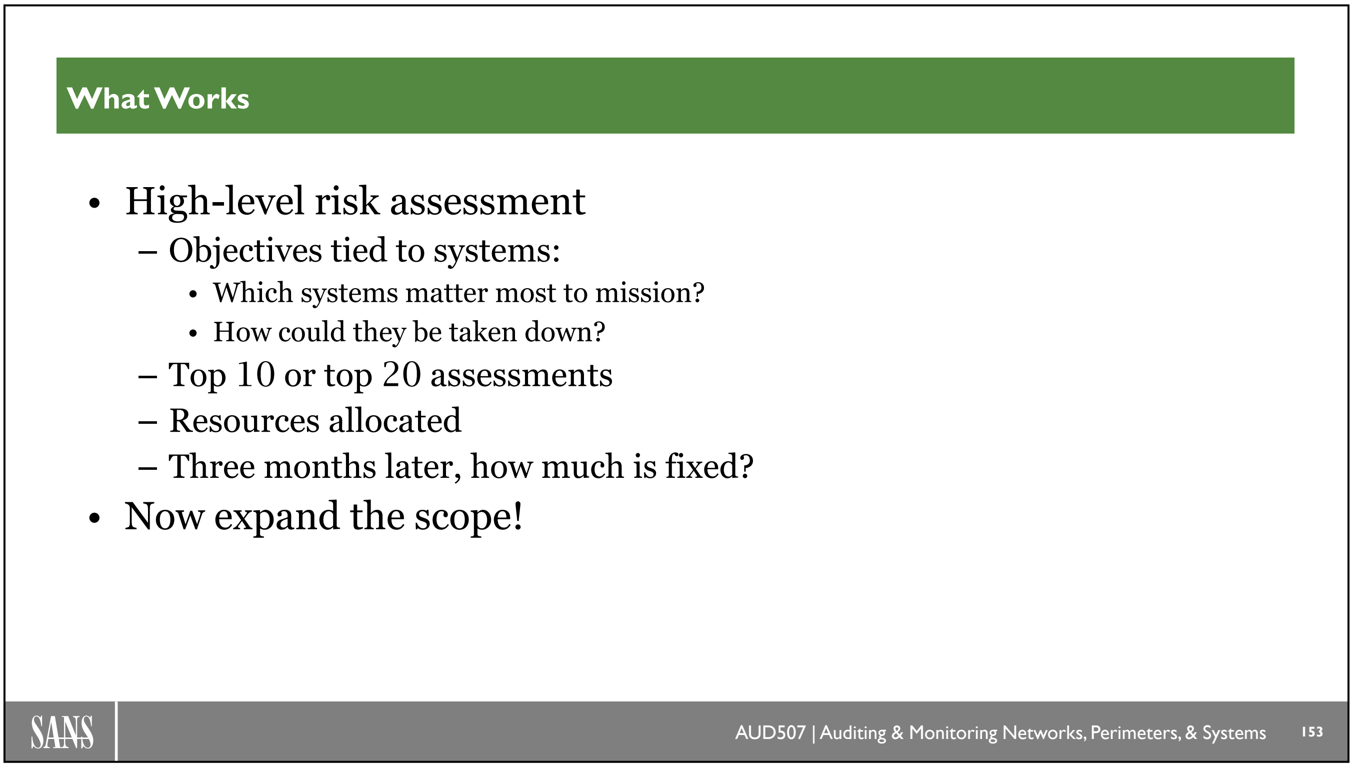

What Works

- High-level risk assessment
  - Objectives tied to systems:
    - Which systems matter most to mission?
    - How could they be taken down?
  - Top 10 or top 20 assessments
  - Resources allocated
  - Three months later, how much is fixed?
- Now expand the scope!

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 153

What does work is to approach the problem in an entirely different way. We begin by performing a high-level risk assessment, working to identify the organizational mission and identify which operational components most directly influence the ability to achieve that mission. With that knowledge, we now work to identify which systems support those operational components.

Now that we are at a system level, we take the time to identify the top 10 or 20 risks that might exist in the context of those systems. This is a great time, by the way, to have a look at vulnerability lists and resources such as the 20 critical controls.

With these top 10 or top 20 items identified, we now perform a vulnerability scan for just these things. The report is now only approximately 100 pages, and it’s actually important that all these items are fixed. The administrators now have the same three months to work through remediation.

When we return in three months to check on remediation, we still have findings, but the report is only 10 or 20 pages. When we report to management now, they can see clear improvement and real return on investment. They actually feel more secure, and they have good reason for feeling that way!

Given that we’ve made progress, funding will likely continue. We can now expand out to the top 25 or the top 50 issues, progressively working through issues based on actual risk rather than arbitrary vulnerability.

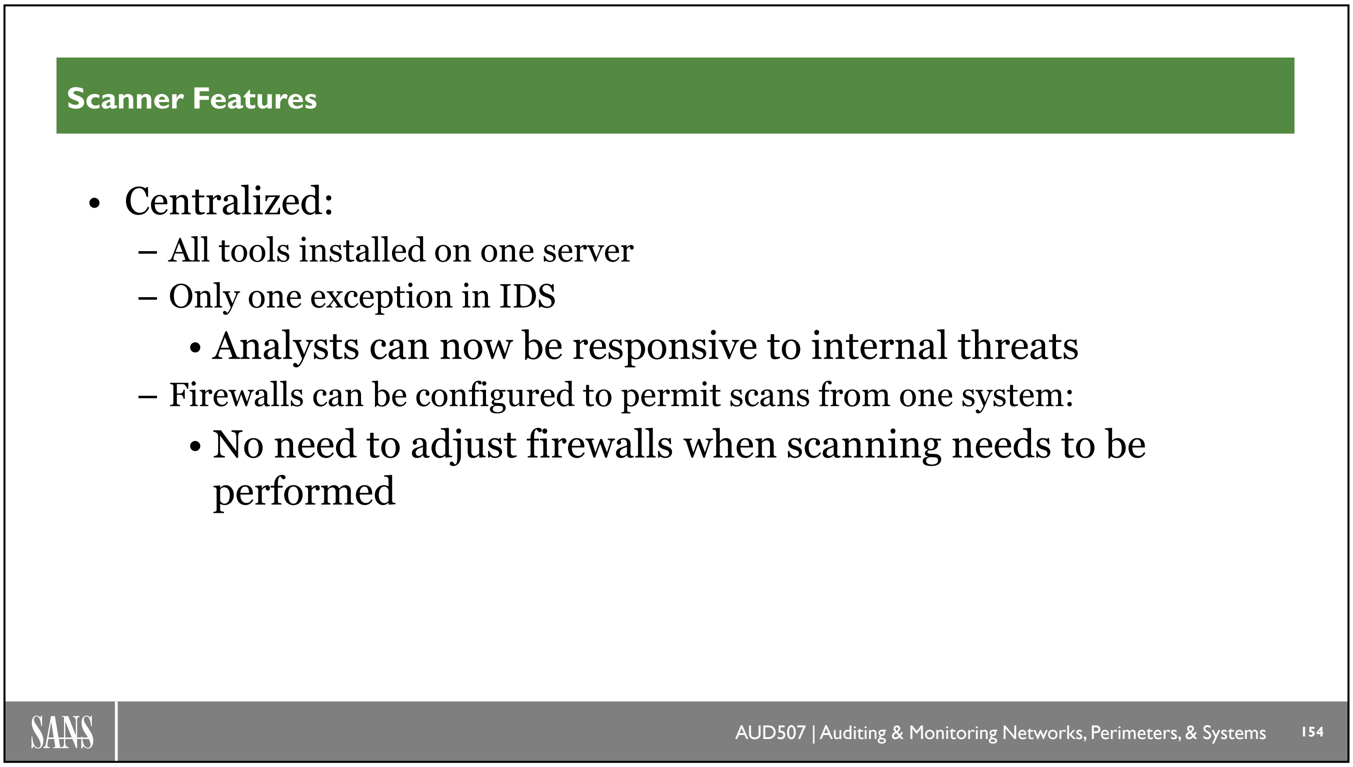

Scanner Features

- Centralized:
  - All tools installed on one server
  - Only one exception in IDS
    - Analysts can now be responsive to internal threats
  - Firewalls can be configured to permit scans from one system:
    - No need to adjust firewalls when scanning needs to be performed

Now that we understand a process that works, let's talk about features that you'd like to find in a vulnerability scanner. There are many scanners that can fit the bill. There's no single vulnerability scanner that we would recommend over all the others, and trust me, if there were one that I preferred, I'd let you know!

One of the more important features I'd like, from the perspective of a security officer, is a centralized scanner. I don't want something that I have to install onto endpoints. I'd greatly prefer something with a web-based interface that I can deploy centrally in my network. This means that I no longer need to allow administrators to install hacker and other security tools onto their systems. They also don't need to be local administrators. Instead, the location of the vulnerability scanner becomes, essentially, the approved security scanning host. If there's some tool that someone needs, it can be installed there, and that user given the access needed to use it.

With this centralized, I can also make better use of my IDS/IPS. If you don't have a centralized scanning device, IDS analysts eventually become desensitized to internal scans because administrators run them from all over the place. With the centralized server, any scan coming from anywhere else is unauthorized, and we can deal with it as a security threat.

Another advantage is that it will not be necessary for administrators to turn firewalls off and on to conduct scans. Instead, systems can be configured to allow the authorized scanner to conduct scans, but no other system is permitted. This all becomes manageable.

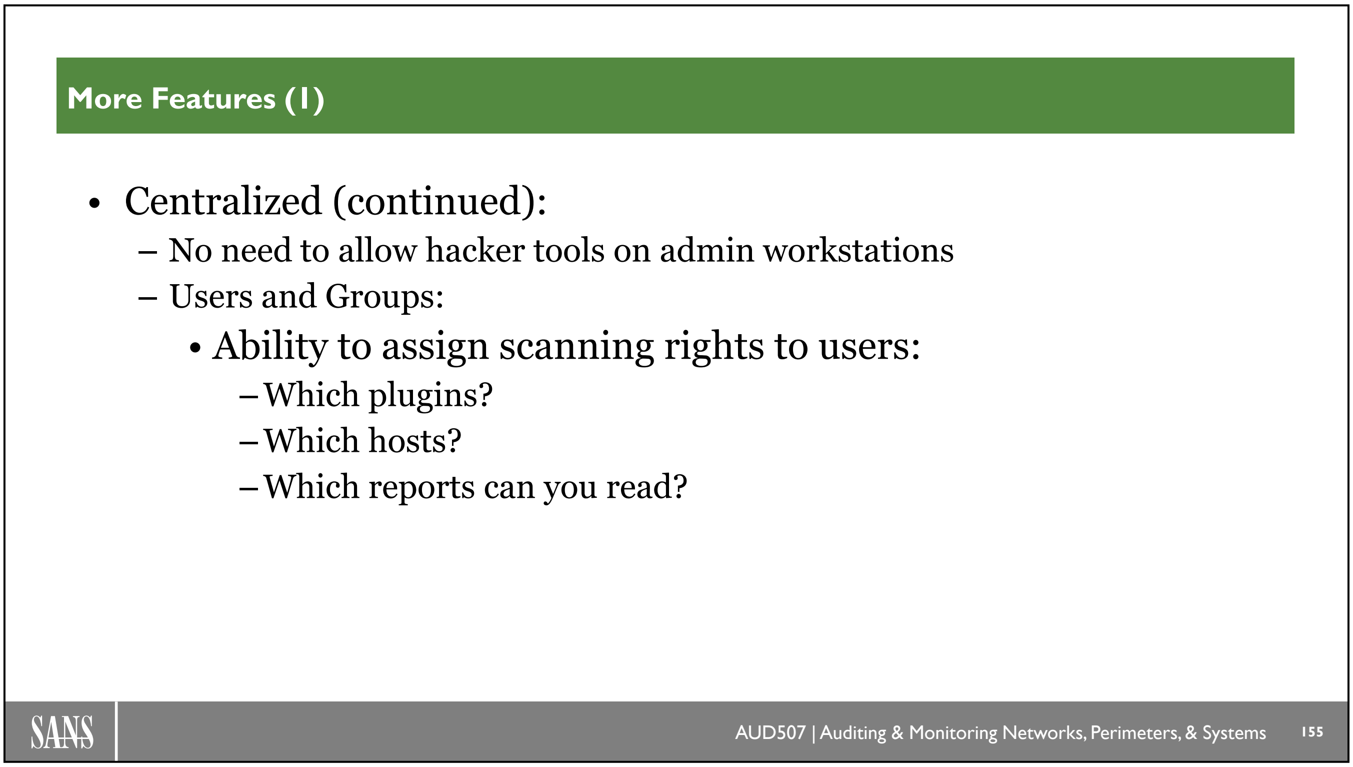

More Features (1)

- Centralized (continued):
  - No need to allow hacker tools on admin workstations
  - Users and Groups:
    - Ability to assign scanning rights to users:
      - Which plugins?
      - Which hosts?
      - Which reports can you read?

Another feature that I'd like is the ability to distinguish users and groups within the scanning system. In particular, I'd like the ability to assign rights and restrict features from various user groups.

Imagine that you're a Windows administrator. If this is the case, you should certainly have the ability to run scans against the servers that you administer. However, you should *not* scan anything else. Furthermore, although you are allowed to scan Window hosts, I may want to restrict you from using certain "dangerous" plugins that are known to cause issues. All this should be configurable by the administrator.

Don't overlook the ability to restrict which reports any particular user can view. Certainly, the security officer and auditor can view any report. As a Windows administrator, however, you should view only the aspects of reports that actually apply to the systems that you manage, and nothing else.

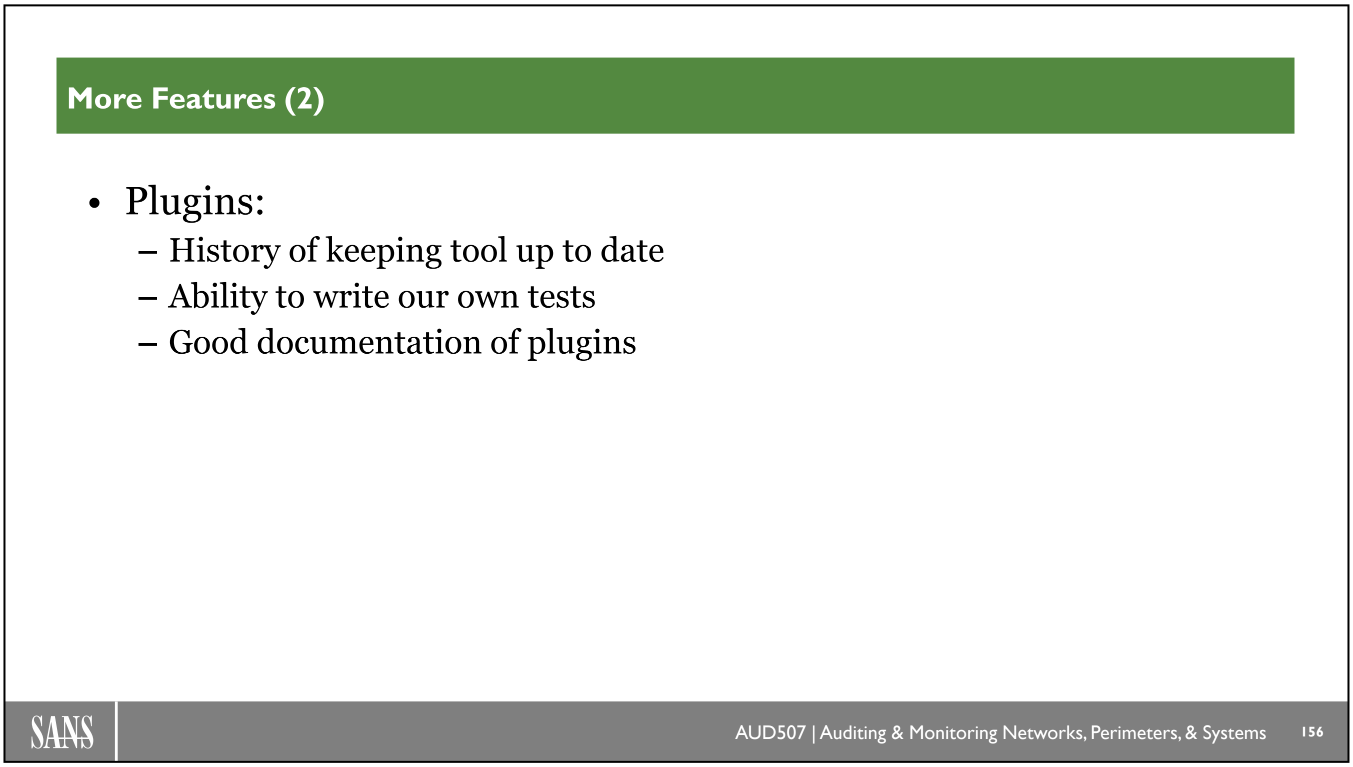

More Features (2)

- Plugins:
  - History of keeping tool up to date
  - Ability to write our own tests
  - Good documentation of plugins

This brings us to plugins. We want to make sure that whichever scanning tool we use, it has a good history of updates from the vendor. Many vendors begin quite well intentioned, but some just can't seem to keep up with the vulnerabilities that are discovered every day.

We'd like to find, too, that the vendor provides good documentation with the plugins. It is frustrating to find something in your report as an important finding, but when you try to dig into it and figure out what it means, the tool has only the most rudimentary data that is essentially meaningless.

Another important feature is the ability for us to create our own plugins or signatures. This is especially important if we develop software or hardware within our organization. Now, when we test our own internal stuff and find issues, we can build plugins to test for these things automatically. This allows us to leverage our existing vulnerability-testing infrastructure rather than having to build some other tool for every test that we need to automate.

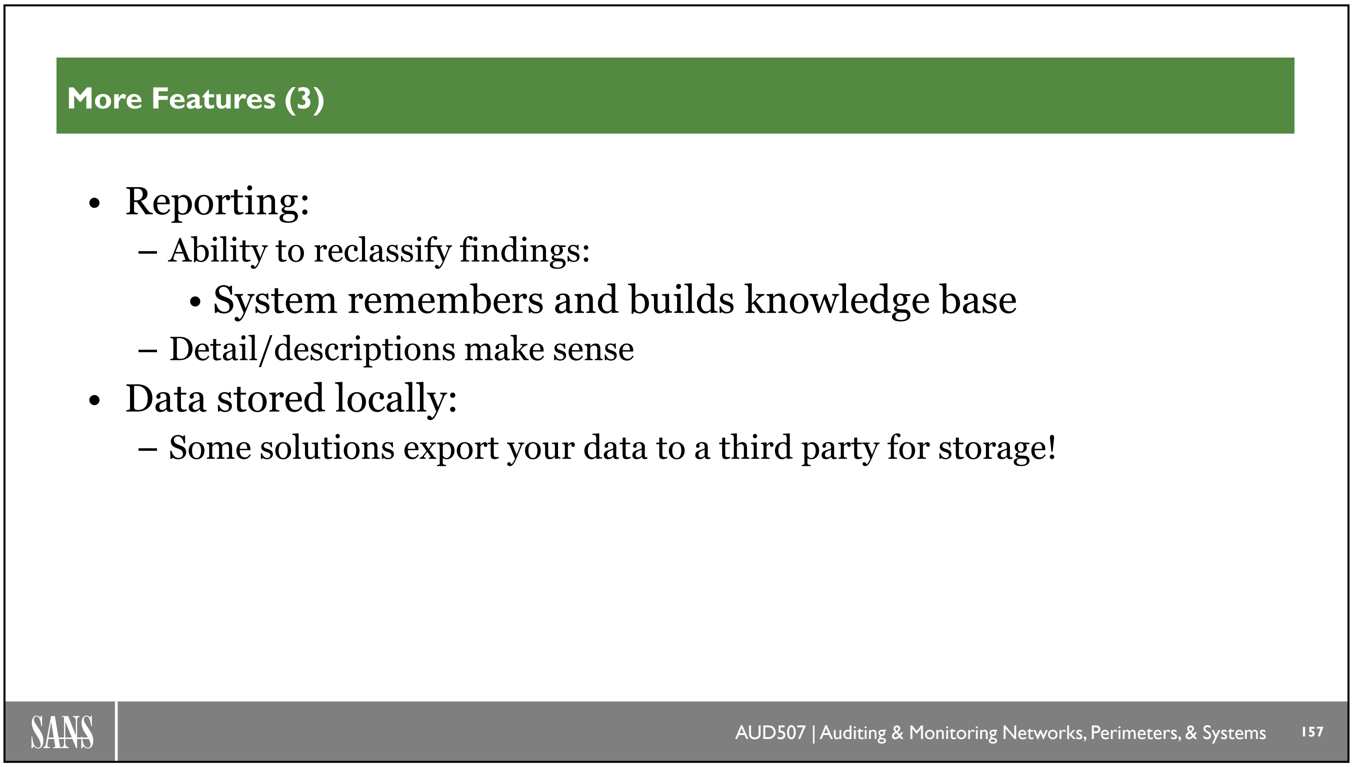

More Features (3)

- Reporting:
  - Ability to reclassify findings:
    - System remembers and builds knowledge base
  - Detail/descriptions make sense
- Data stored locally:
  - Some solutions export your data to a third party for storage!

SANS

AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 157

For reporting, we should look for a couple important features. Obviously, the reports should be easy to read. More than this, though, we'd like the ability to reclassify findings. We'd also prefer that the system remembers how we have classified a particular finding for a particular host so that all future scans will take this into account.

Understand that we're not just talking about the ability to mark and exclude false positives. We also want the ability to simply reclassify something as unimportant or as important regardless of how the tool feels about that issue.

We also want to be sure that all the vulnerability data and reports are stored locally. We point this feature out because there is an otherwise wonderful vulnerability-scanning tool that provides what some feel is the best reporting in the industry, but all your report data is sent out to its site for management. When you want to view a report about your system, you must log in to its site to retrieve the report.

In all fairness, this vendor does claim that your data is all encrypted and it doesn't have a key that can read it. Even if that is the case, I am extremely uncomfortable with my vulnerability information being stored anywhere other than inside of my network.

I'm sure you're wondering who this vendor is. Don't worry, I'll tell you in a couple slides.

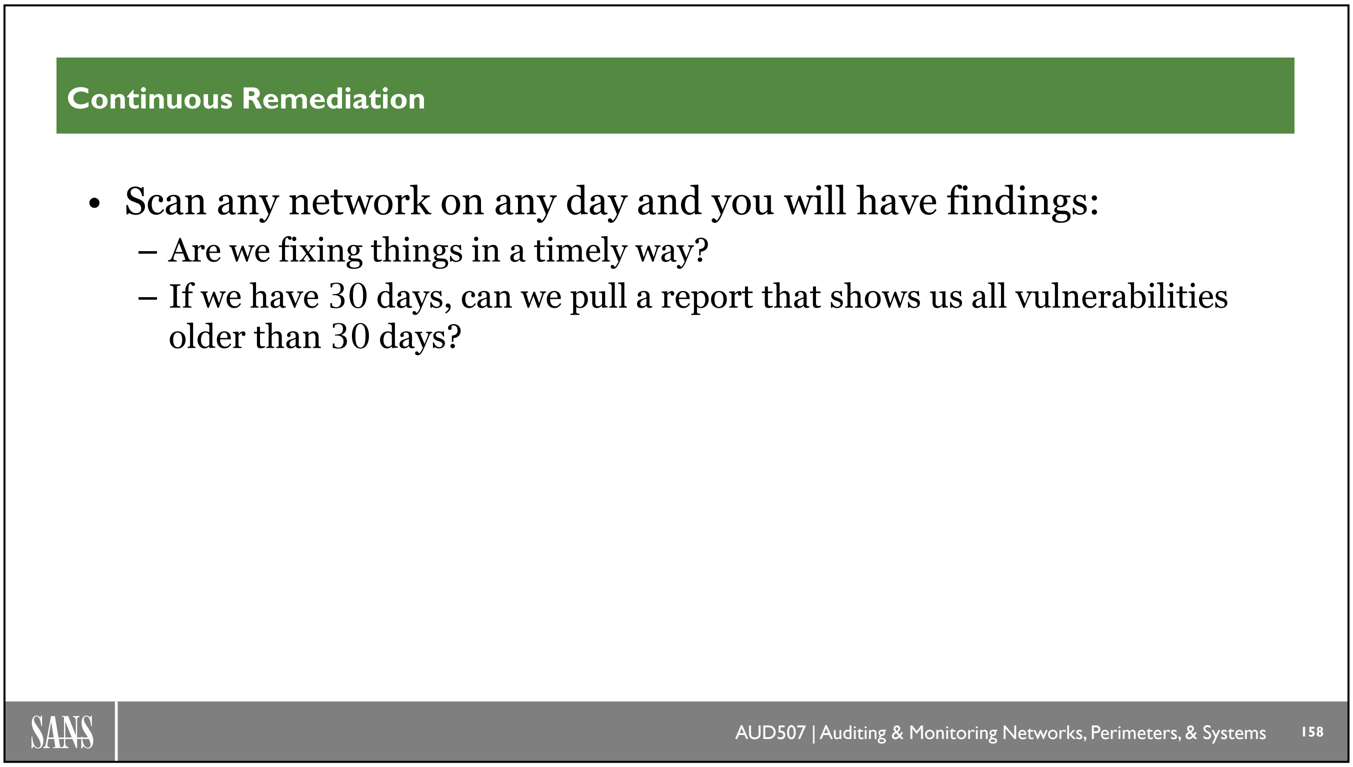

Continuous Remediation

- Scan any network on any day and you will have findings:
  - Are we fixing things in a timely way?
  - If we have 30 days, can we pull a report that shows us all vulnerabilities older than 30 days?

When working to validate the security of networks, keep in mind that if you scan any network in the world on any given day, you will absolutely find vulnerabilities in it. I'm not saying that an attacker can just walk right in, but I am saying that there are patches missing or configurations that are not ideal.

This problem is not a result of a poor security budget or a lack of focus on security issues. This is because networks are constantly evolving, vulnerabilities are constantly being discovered, and we must go through testing and patching processes. This takes time. Add to this the sheer number of issues that must be fixed every month, and it's a wonder that we are keeping up as well as we are!

What we would like to do is look at the patch requirements that the organization states. If, organizationally, patches are supposed to be applied within 30 days of release, then we should be able to look at a vulnerability scan that includes everything older than 30 days and it should be pretty clean.

This gives us another feature that we'd like to find. Can the system's reports show us how quickly we are fixing things? This is sometimes called the Mean Time To Repair (MTTR). If we find that our MTTR is within the threshold specified, we wouldn't get too excited if there are one or two issues that took 40 or 50 days to resolve because of difficulties in applying the patches.

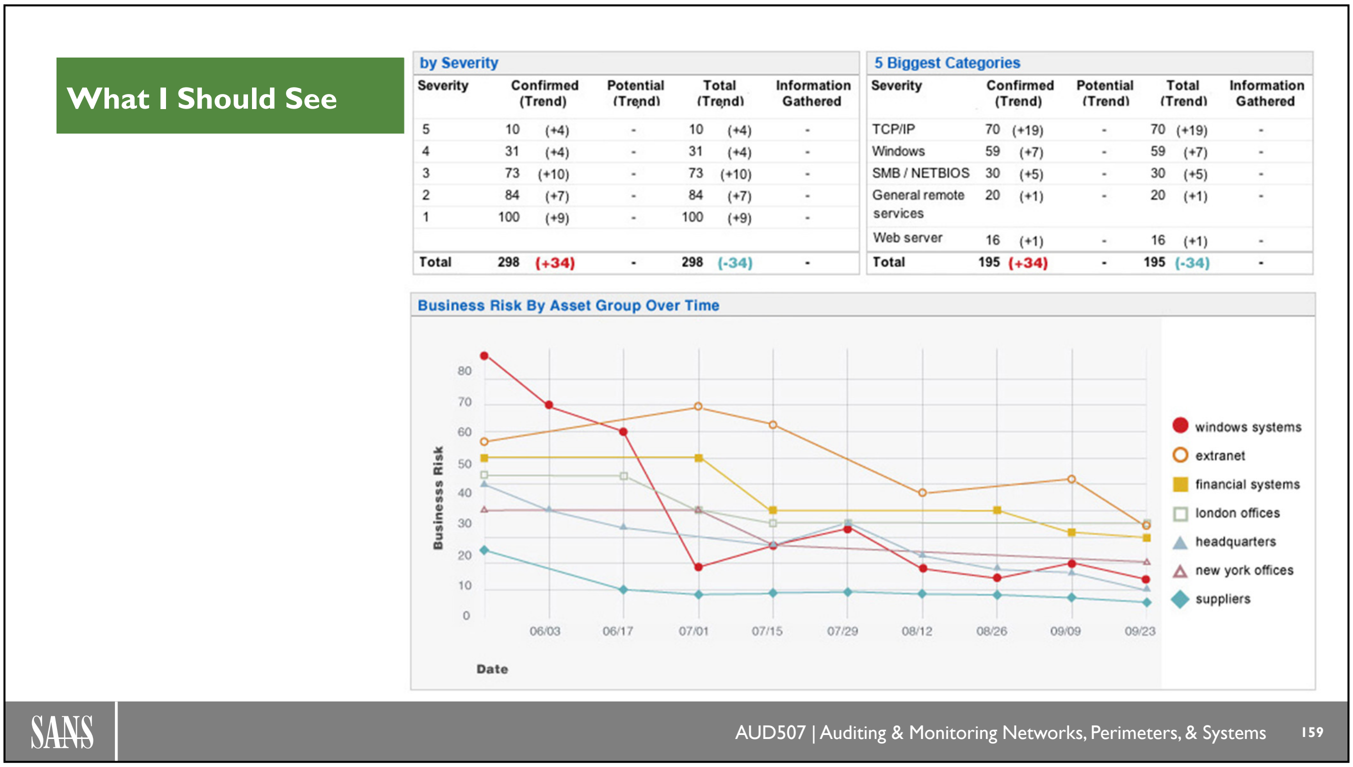

What I Should See

by Severity

Severity	Confirmed (Trend)	Potential (Trend)	Total (Trend)	Information Gathered
5	10 (+4)	-	10 (+4)	-
4	31 (+4)	-	31 (+4)	-
3	73 (+10)	-	73 (+10)	-
2	84 (+7)	-	84 (+7)	-
1	100 (+9)	-	100 (+9)	-
Total	298 (+34)	-	298 (-34)	-

5 Biggest Categories

Severity	Confirmed (Trend)	Potential (Trend)	Total (Trend)	Information Gathered
TCP/IP	70 (+19)	-	70 (+19)	-
Windows	59 (+7)	-	59 (+7)	-
SMB / NETBIOS	30 (+5)	-	30 (+5)	-
General remote services	20 (+1)	-	20 (+1)	-
Web server	16 (+1)	-	16 (+1)	-
Total	195 (+34)	-	195 (-34)	-

Business Risk By Asset Group Over Time

This slide shows the type of reporting you'd like to see. This report comes from Qualys. The QualysGuard tool is, in the opinion of many people in the industry, one of the best for reporting. It doesn't just tell the techies what they need to know, but it also provides useful dashboards that executives can leverage in managing an overall risk-response strategy. The only big downside to Qualys is that these are the folks that we mentioned two slides ago. All the information on your vulnerabilities is actually stored out in the Qualys servers.

Notice the bottom chart in the slide. Here we can see overall trends for vulnerability over time. In the slide, it is broken out by type of system, but this is an arbitrary classification. You can actually configure it to produce the report in just about any way that makes sense to you.

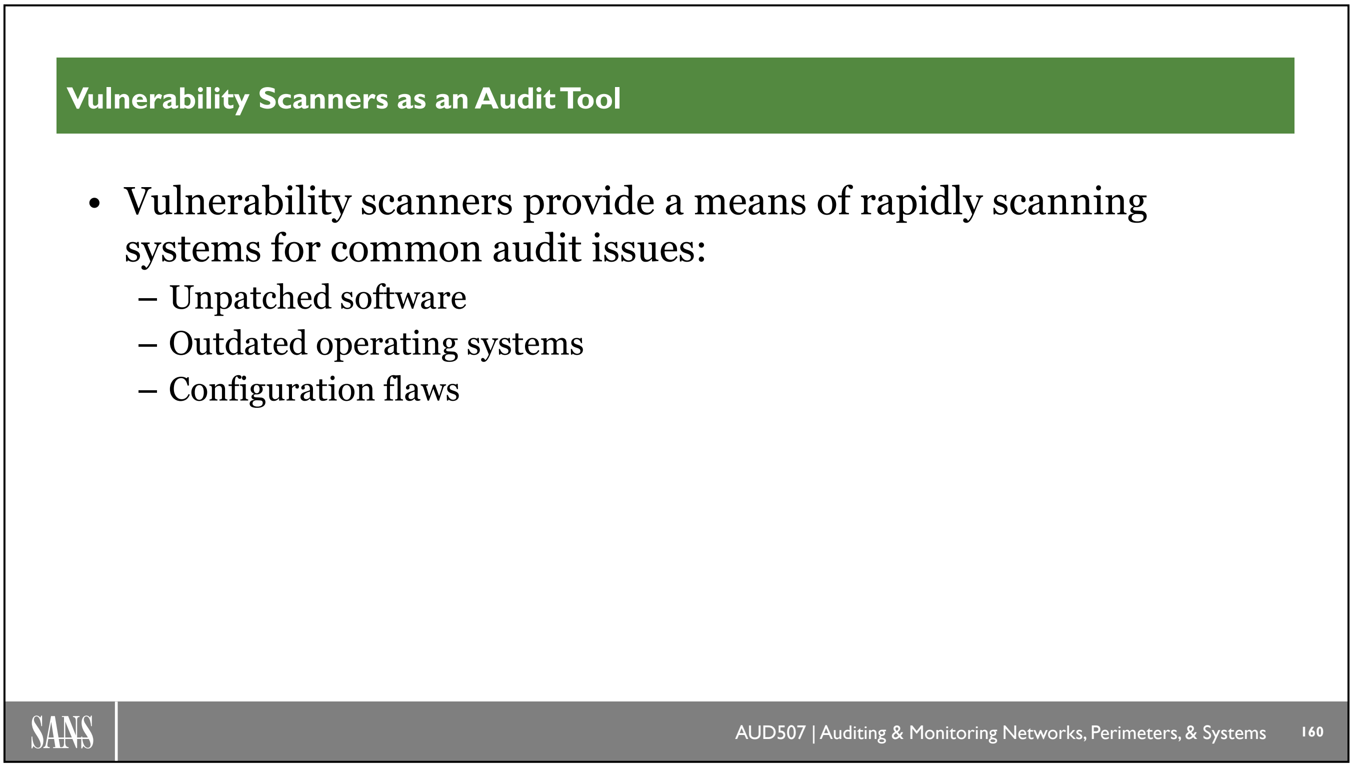

Vulnerability Scanners as an Audit Tool

- Vulnerability scanners provide a means of rapidly scanning systems for common audit issues:
  - Unpatched software
  - Outdated operating systems
  - Configuration flaws

I've always told my students that there are two tools which will make you "look smart" as an auditor – Nmap and a vulnerability scanner! Nmap is awesome for collecting information about inventory, gathering version information, and even testing for some vulnerabilities or misconfigurations, but vulnerability scanners are designed specifically for this kind of work, and updated with new plugins frequently.

Most vuln scanners give auditors a way of testing for outdated or unsupported operating systems or application software. They also allow for checking for common OS and software misconfigurations like incorrect windows registry settings, improper Linux system configuration items, or default content left on a webserver.

While much of the audit work we do will require a finer level of detail than the scanner can provide, it is nice to let your scanner do the "heavy lifting" while you focus on only the details which require your attention.

Compliance Scanning

- Scores systems against published standards
  - CIS Benchmarks
  - DISA STIGs
- Can include offline network device configuration audits

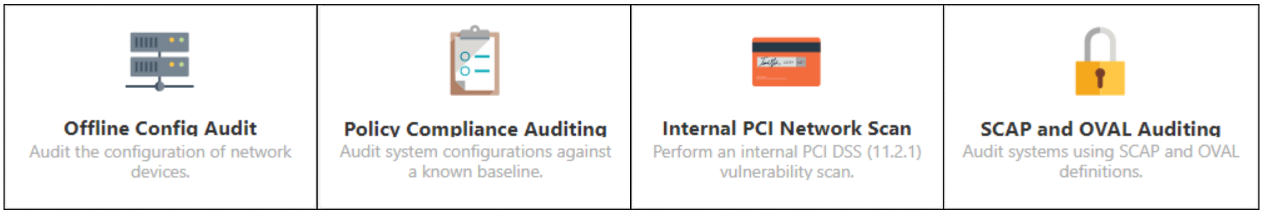

Offline Config Audit
Audit the configuration of network devices.

Policy Compliance Auditing
Audit system configurations against a known baseline.

Internal PCI Network Scan
Perform an internal PCI DSS (11.2.1) vulnerability scan.

SCAP and OVAL Auditing
Audit systems using SCAP and OVAL definitions.

Using the compliance scanning feature of most vulnerability scanners is usually a simple process consisting of choosing the standard, answering questions about enterprise-specific settings like logon banners and subnet addresses, then running the scan against all selected hosts.

The upgraded Professional version of the Nessus Essentials scanner we use in the lab comes with the CIS Benchmarks and DISA STIGs pre-loaded as audit policies, making scanning for compliance with these standards quite easy.

Most scanners can also do compliance scanning against text configuration files from network devices (like switches, routers, and firewalls) for compliance with supported standards.

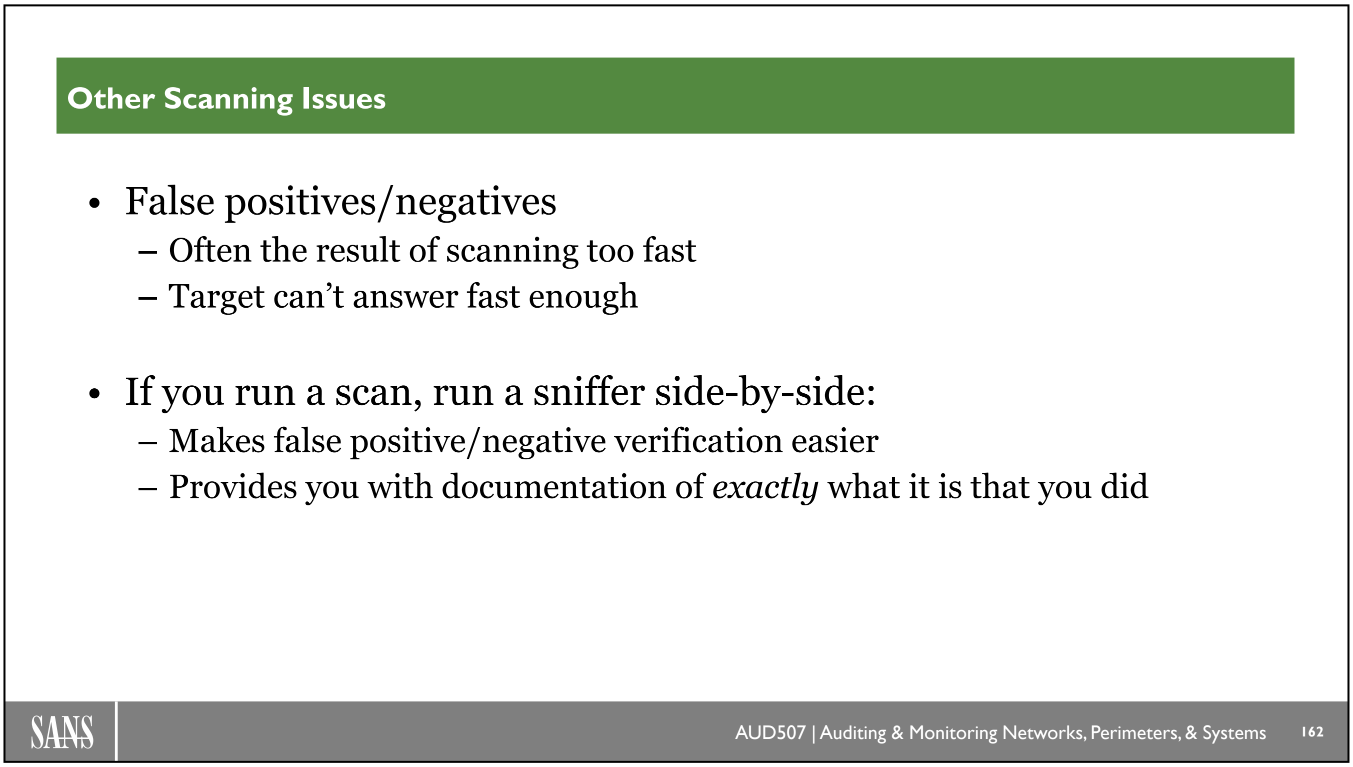

Other Scanning Issues

- False positives/negatives
  - Often the result of scanning too fast
  - Target can't answer fast enough
- If you run a scan, run a sniffer side-by-side:
  - Makes false positive/negative verification easier
  - Provides you with documentation of *exactly* what it is that you did

SANS | AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems 162

False positives and false negatives are a fact of life in any vulnerability scanning system. A false positive occurs when the tool tells you that there is an issue when, in reality, there is nothing actually wrong. A false negative is when the system fails to identify an issue even though you do, in fact, have a problem.

There are a few common causes for both false positives and false negatives. Probably the most common cause is misconfiguration of the tool. Everyone is always concerned about running scans as fast as they possibly can. Although we can understand this desire for immediate gratification, the faster the scan is run, the more likely that the tool will miss things. This can be because the system being scanned is getting loaded down and just can't answer fast enough. It could also be because as the system is becoming overwhelmed and services are failing, the scanner begins to assume that a compromise has actually occurred when it hasn't.

Another major cause, especially of false positives, is poorly written signatures. This is why manual validation of any important findings is valuable.

One last suggestion before we wrap up today's material: If you are ever in a position to run a vulnerability scan to perform some measure of validation, we strongly recommend that you run a sniffer side-by-side with it.

Running a sniffer to monitor the vulnerability scanner provides us with some extremely important data. Not only do we have a record of absolutely everything that was actually done, but we now have a means to research poorly documented plugins and signatures! If we need to figure out how a particular issue was assessed, we simply need to open up our sniffer, find that probe, and extract the content!

This allows us to perform manual validation in addition to better research what the alert is about.

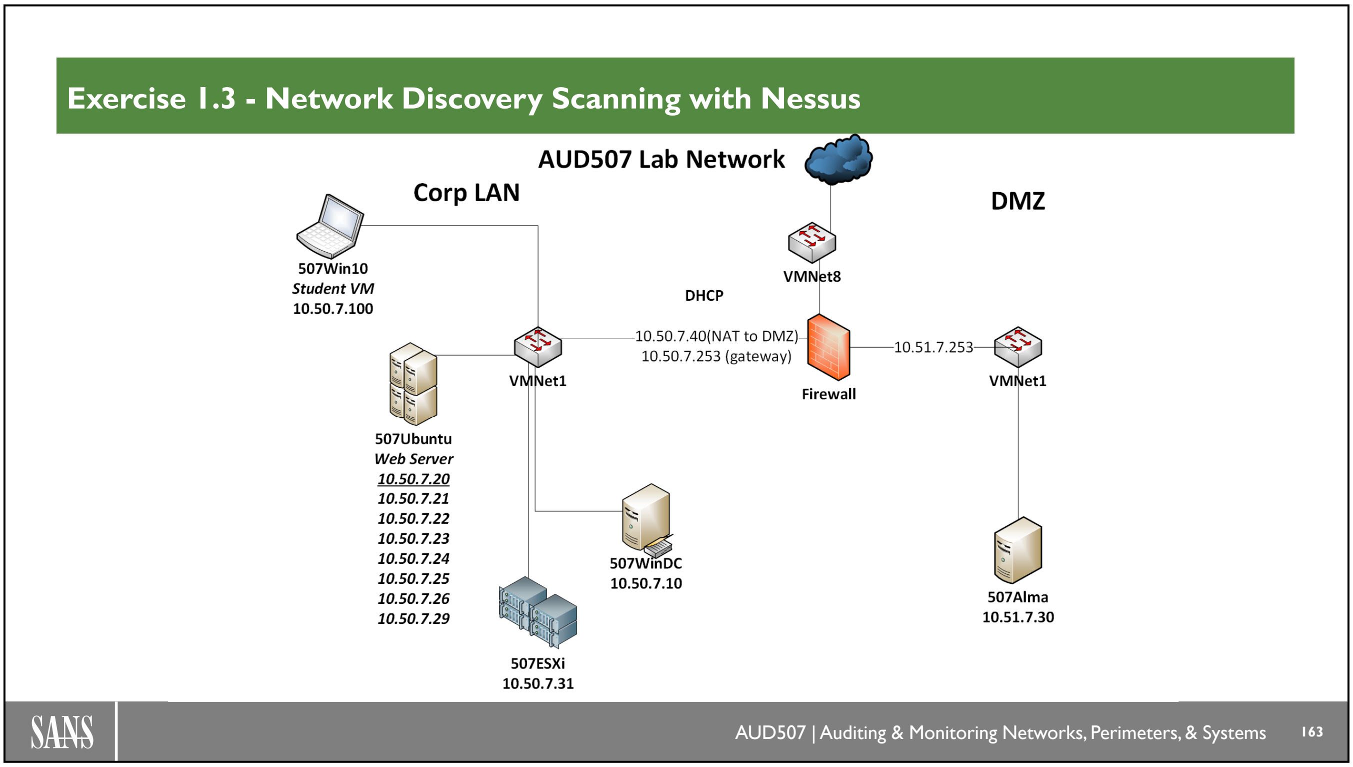

Exercise 1.3 - Network Discovery Scanning with Nessus

AUD507 Lab Network

Corp LAN

DMZ

507Win10
Student VM
10.50.7.100

VMNet8

DHCP

10.50.7.40(NAT to DMZ)
10.50.7.253 (gateway)

10.51.7.253

VMNet1

Firewall

VMNet1

507Ubuntu
Web Server
10.50.7.20
10.50.7.21
10.50.7.22
10.50.7.23
10.50.7.24
10.50.7.25
10.50.7.26
10.50.7.29

507WinDC
10.50.7.10

507Alma
10.51.7.30

507ESXi
10.50.7.31

This page intentionally left blank.

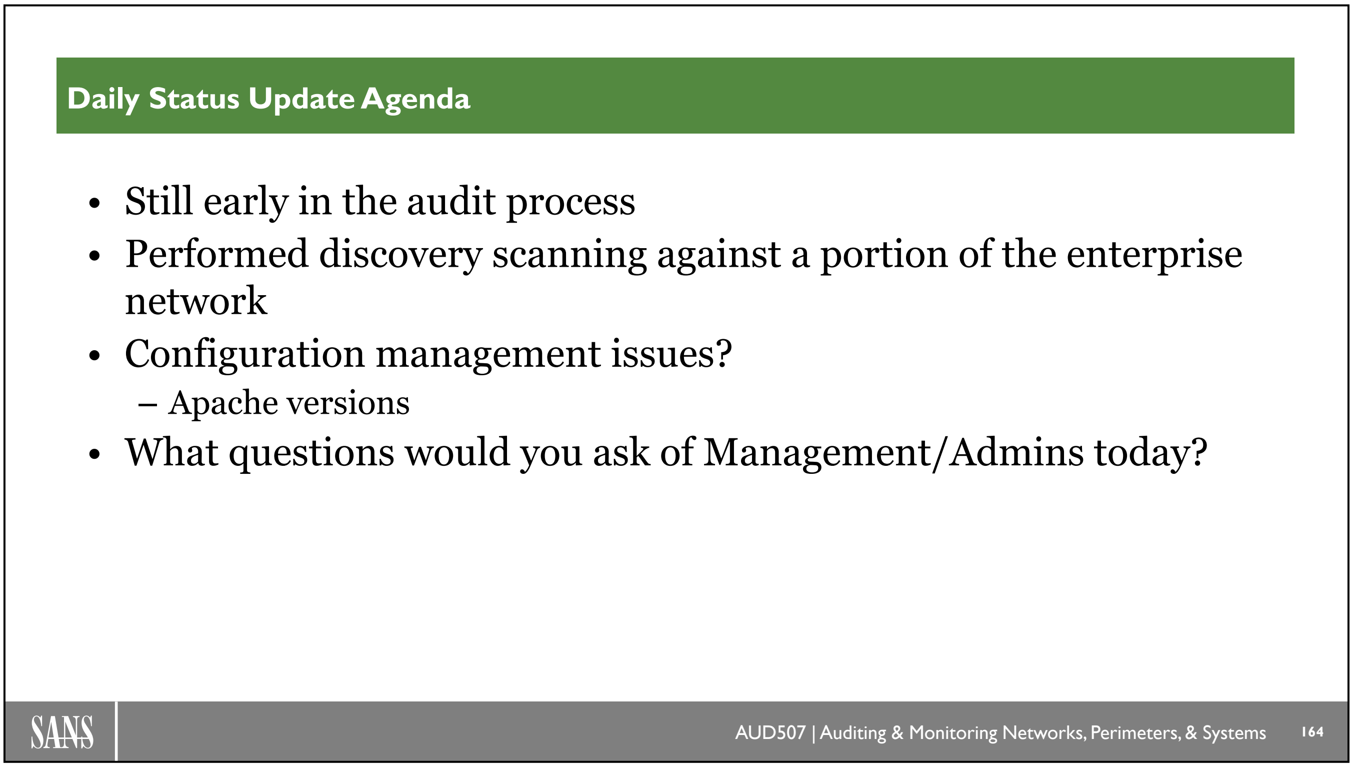

Daily Status Update Agenda

- Still early in the audit process
- Performed discovery scanning against a portion of the enterprise network
- Configuration management issues?
  - Apache versions
- What questions would you ask of Management/Admins today?

SANS
AUD507 | Auditing & Monitoring Networks, Perimeters, & Systems
164

For the daily client update meeting, we can use this agenda.

While it’s still early in the audit, we have begun to gather information that might be useful in understanding some of the systemic issues the enterprise might be facing. Think about the variance you noticed in Apache versions on the web servers in the network.

- Could this indicate problems with the change management and system update processes?
- What questions would you want to ask management about change and configuration management?

Thank You!

This brings us to the end of Section 1. If you are taking the class at a conference, please take a moment to complete an evaluation form. You will be given a different evaluation every section of the class.