501.6
Data Loss Prevention

SANS

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

Data Loss Prevention (DLP)

Version A12_02

SEC 501 Advanced Security Essentials

This page intentionally left blank.

Course Outline

- Risk Management
  - Calculating and understanding risk across an organization
  - Applying proactive risk-management processes
  - Incorporating risk management into all business processes
- BCP/DRP
- Insider Threat
- Data Classification
  - Building a data classification program
  - Key aspects on deploying and implementing classification of critical information
  - Staged roll out of classifying new and existing information
  - Managing and maintaining portable data classification
- Digital Rights Management
  - Understanding what digital rights are
  - Balancing digital rights with data classification
- Data Loss Prevention (DLP)
  - Identifying requirements and goals for preventing data loss
  - Peeling through the hype of DLP
  - Identifying practical DLP solutions that work
  - Managing, evaluating, implementing, and deploying DLP

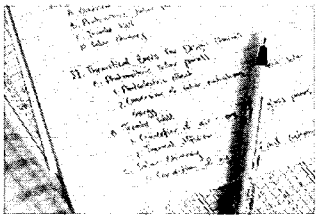

SEC 501 Advanced Security Essentials

This page intentionally left blank.

Introduction to Data Loss Prevention

- Protection of information:
  - Data in transit
  - Data at rest
  - Data classification
- Overall information protection and risk program

SEC 501 Advanced Security Essentials

Introduction to Data Loss Prevention

When dealing with information security in any form, there are generally two types of protection levels we must be concerned with:

- Data in transit (involves the endpoint security of devices used in communications)
  - Blackberry devices
  - VPNs
  - Etc.
- Data at rest (the security of media where stored)
  - Physical access/storage
  - Encryption of media

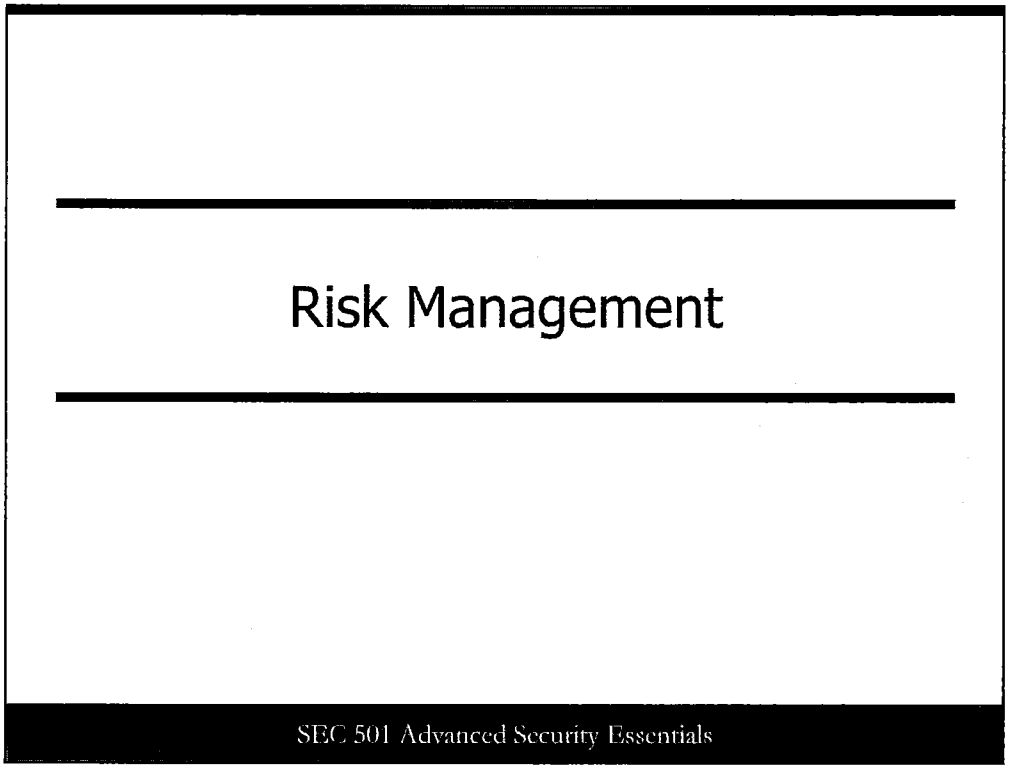

The following slides are designed to discuss and clarify the risk management process.

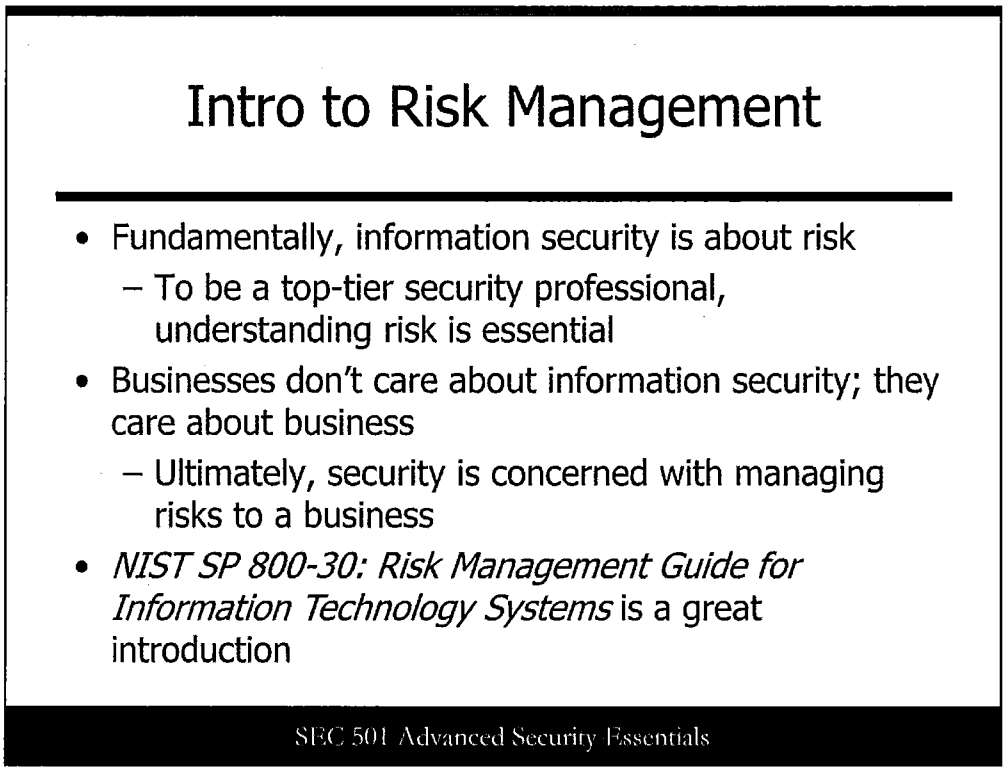

Intro to Risk Management

One of the most important concepts that an information security professional can understand is that of risk. Ultimately, the job of most security professionals boils down to risk management. Unfortunately, most information security professionals lack a keen understanding of risk management principles and simply make recommendations without truly appreciating the risk ramifications to their organization.

Security professionals must understand that the purpose of the organization is to fulfill its mission. The purpose of a security professional is to help the business make informed decisions about security issues that could potentially compromise the organization's mission.

A great primer to risk analysis and risk management is found in NIST's *Special Publication 800-30: Risk Management Guide for Information Technology Systems*.[1]

[1] http://csrc.nist.gov/publications/nistpubs/800-30-rev1/sp800_30_r1.pdf (http://cyber.gd/414_16)

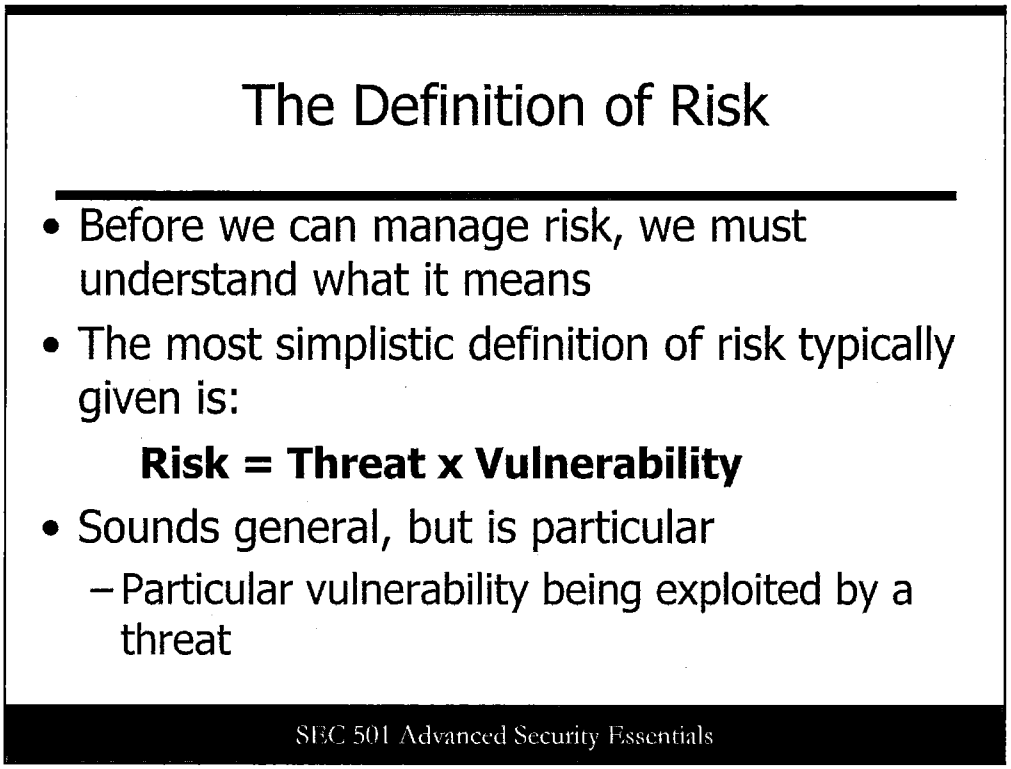

The Definition of Risk

Naturally, for the security professional to understand how to effectively manage risk, she must first understand risk. The simplest definition of risk associated with information security is:

Risk = Threat x Vulnerability

Appreciate that although the definition looks fairly general, it is actually particular. Risk is calculated for particular threat/vulnerability pairs.

On the surface the definition Risk = Threat x Vulnerability appears fairly simplistic and straightforward. However, actually calculating values can be quite challenging. To appreciate these challenges, we will parse the underlying concepts of threat and vulnerability and also fill in some gaps in this oversimplified definition.

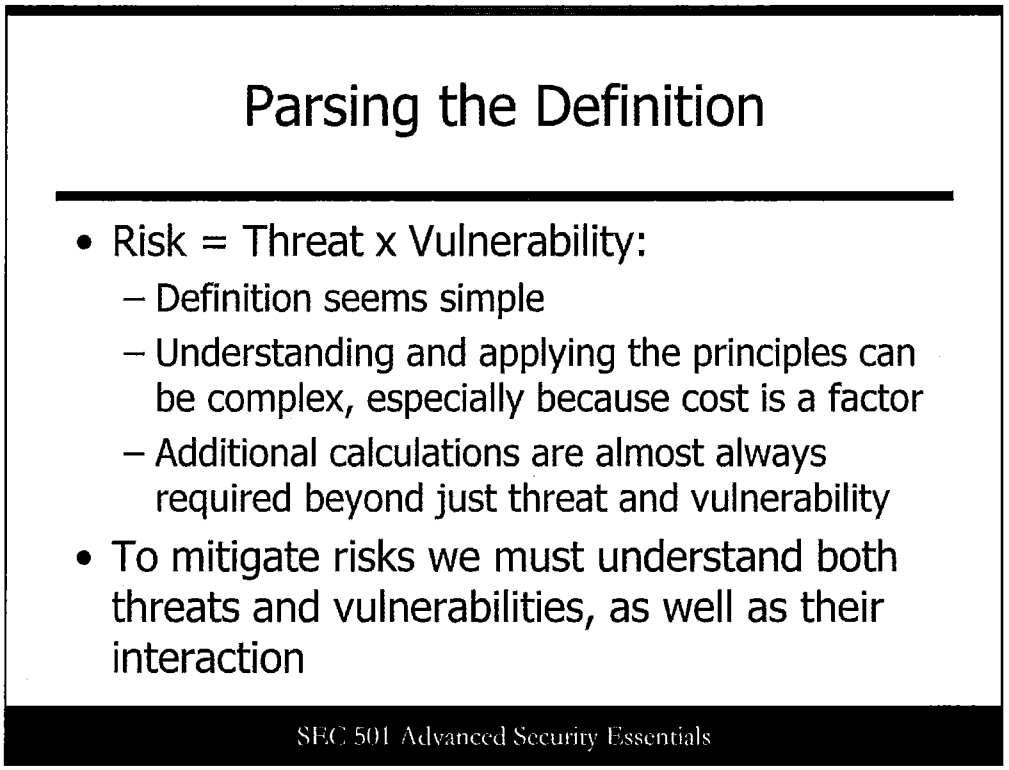

Parsing the Definition

Working with the simplified risk definition, Risk = Threat x Vulnerability, might give the impression that these calculations are easy to perform. However, this definition sometimes gives a false sense of simplicity. Yes, to decrease risk all the security professional has to do is to decrease the threats or vulnerabilities. Seems straightforward enough, and it would be if we had unlimited time and money. However, most organizations are limited on both of these resources, which forces decisions to be made about which threats or vulnerabilities to decrease, and how much they will be decreased with the various countermeasures.

In addition, there are important factors that inform the definition that are omitted in this simplistic definition, as we will see.

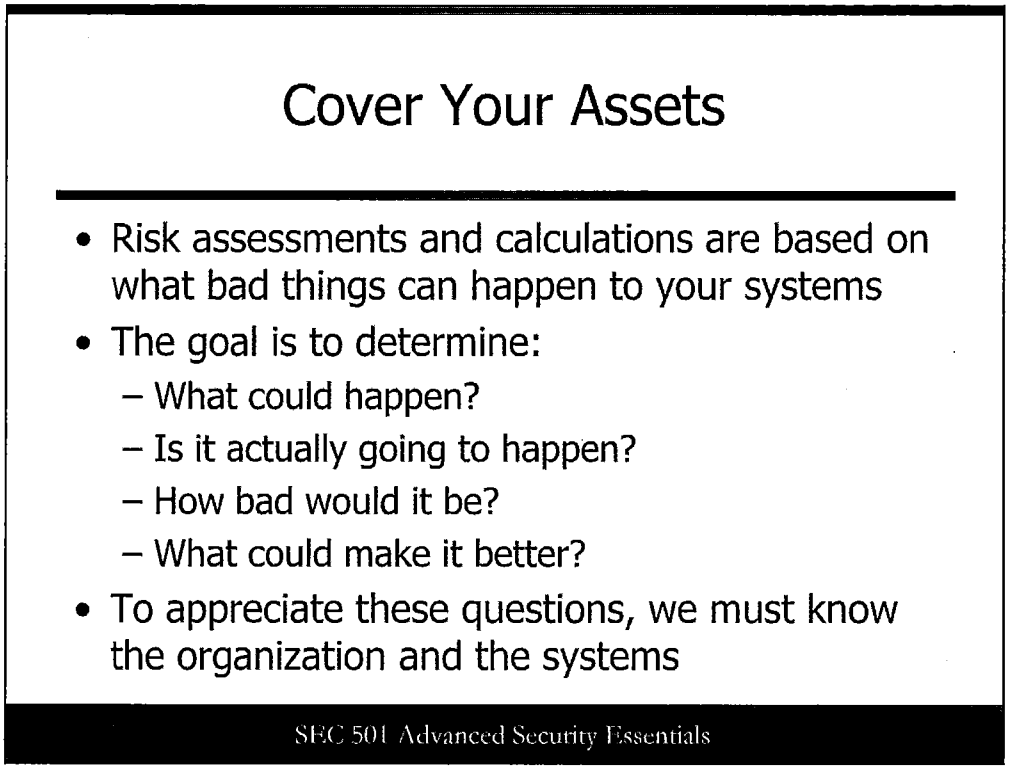

Cover Your Assets

Before we dive deep into the definition, let's take a step back and appreciate what we are hoping to understand and achieve through this process. Rather than precise words with specific meanings, let's use some simple questions to drive the process.

The goal is to determine answers to the following questions:

- What could happen?
- Is it actually going to happen?
- How bad would it be?
- What could make it better?

These straightforward questions illustrate most of what is done in risk analysis and risk management. Now we turn our attention to parse the more technical side of these questions.

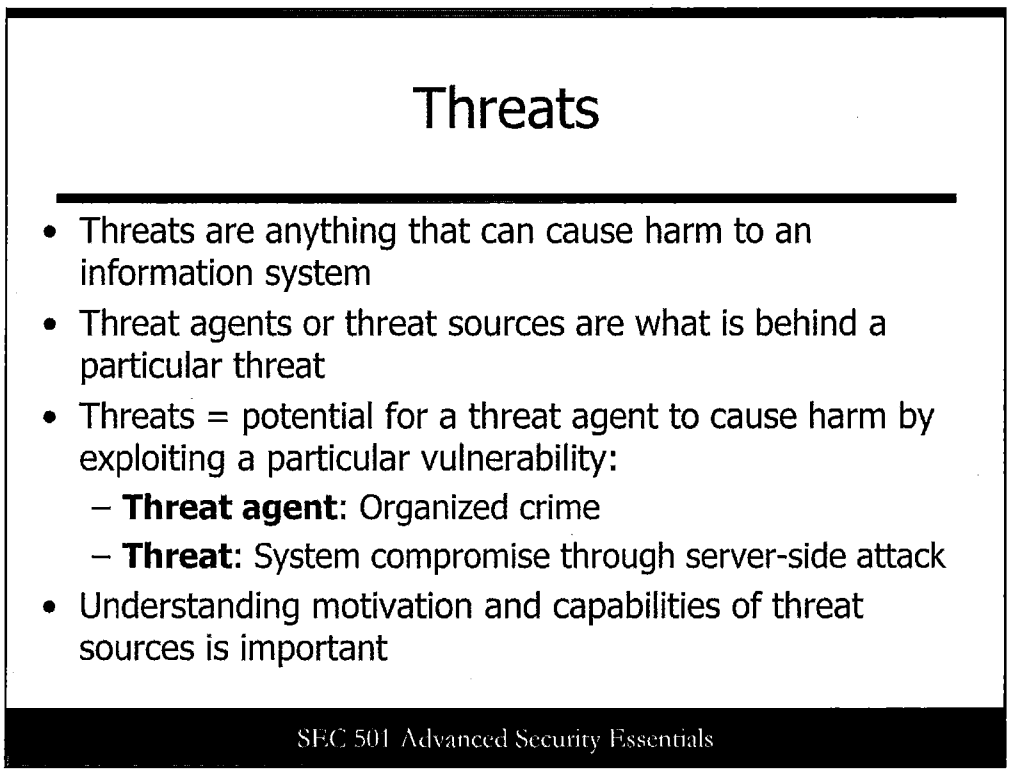

Threats

The first item in our definition to be parsed is the concept of a threat. A threat is simply something that can bring harm to an information system. Though our simplistic definition doesn't include this element, there is always a threat-source (aka threat-agent) that serves as the cause of the threat.

Let's use an example threat statement. Our web server will be DoSed via a server-side attack against the vulnerability associated with MS11-100.

MS11-100 is a patch for a vulnerability in ASP.NET that allowed a DoS to be introduced by targeting ASP.NET's hash table generation for POST variables. This attack would be most likely be carried out by sending an HTTP POST with an extremely large number of POST variables set to introduce a hash collision.

The threat is a denial of service condition on a web server. The threat remains regardless if this particular vulnerability exists on the web server. The risk for this particular threat vulnerability pair would be likely eliminated if the patch were successfully deployed.

We do not see anything about the threat source in the threat statement. It is actually fairly common for organizations to ignore the threat sources. However, the threat source becomes especially important when we try to determine another key concept, likelihood, which will be reviewed later. Key questions concerning the threat source are whether the source is motivated and whether the source is capable of introducing this threat.

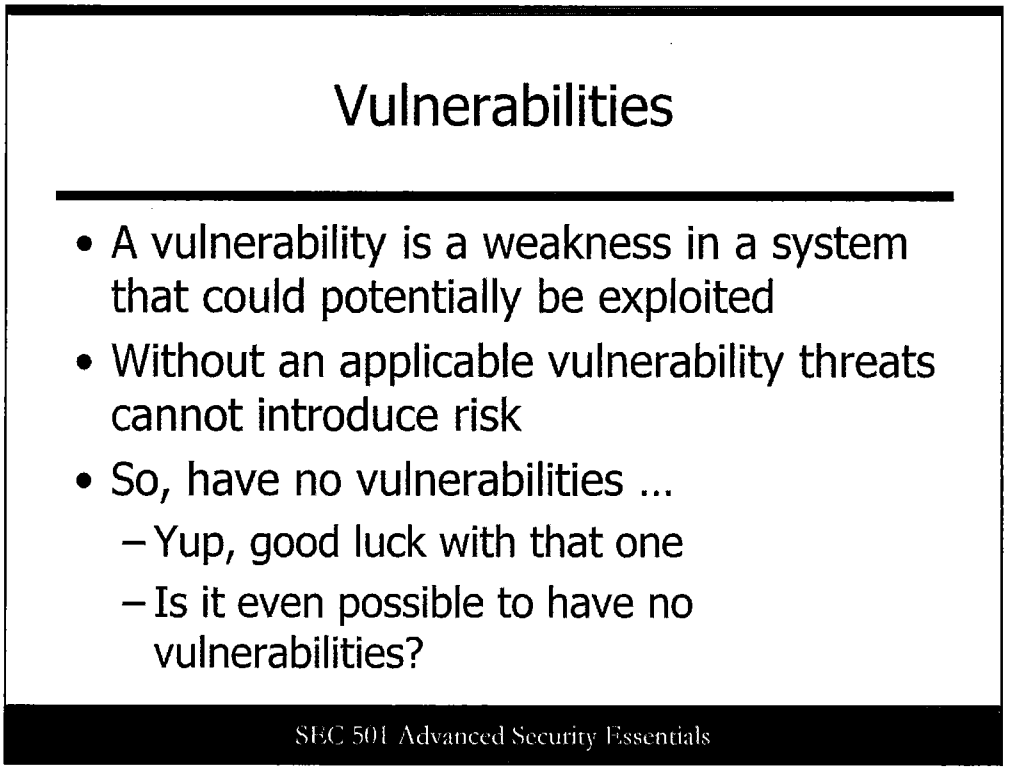

Vulnerabilities

Even if there are numerous motivated threat agents and there is no vulnerability, there is no risk. In the previous threat statement, "Our web server will be DoSed via a server-side attack against the vulnerability associated with MS11-100," a vulnerability is mentioned. If we are not vulnerable to the vulnerability associated with MS11-100, then we have no risk.

Potential ways in which this vulnerability would not exist include the server is patched; the server is not using IIS; the server is using IIS, but not ASP.NET; and the server is not Windows-based. Again, these are potential ways in which the vulnerability might not be present and is not necessarily true of MS11-100.

So, to not have any risk all, we have to do is have zero vulnerabilities. Sounds straightforward enough. Run a vulnerability scanner. Get everything patched. What is so hard about that?

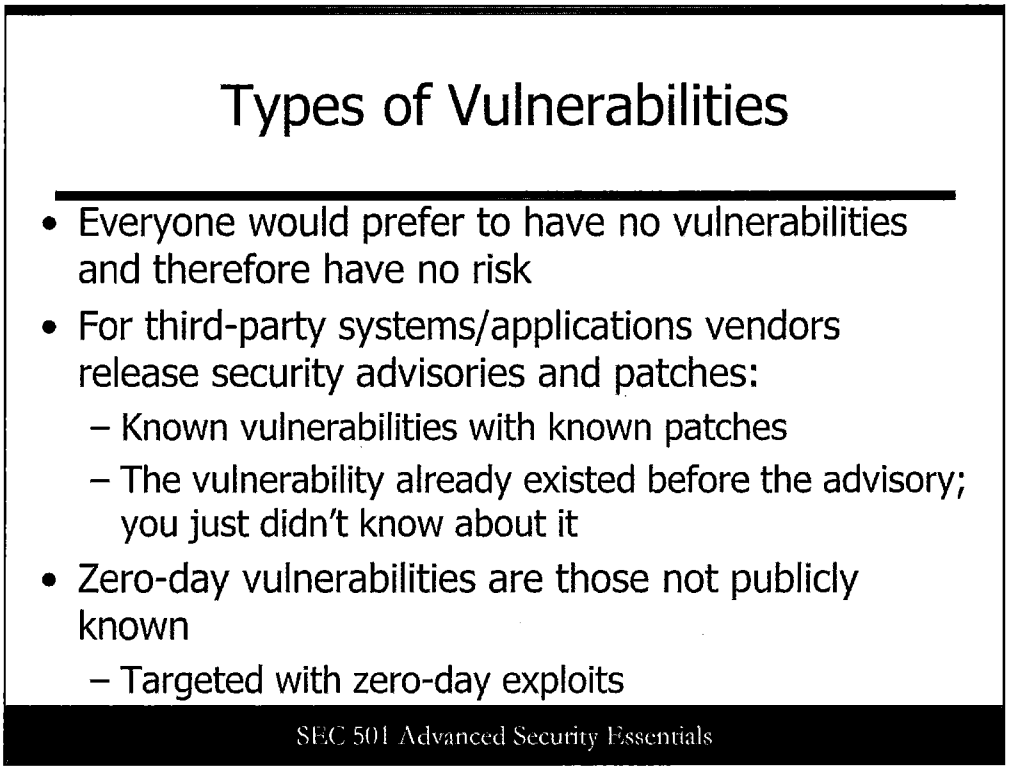

Types of Vulnerabilities

Although in theory patching every vulnerability might sound possible, reality is a different case. Patching Microsoft vulnerabilities is easier than almost all other vendors, and yet organizations are still often compromised by exploitation of these vulnerabilities. Then realize that the organization has to patch every known flaw on every system (including printers, access control systems, HVAC, and such). Sounds pretty tough, yet even if an organization were successful in patching everything, it would still have exploitable vulnerabilities.

Even if you patch everything you know to patch, you have still failed. Why? Those vulnerabilities that we patch (sometimes >10 years after the OS/application's release date) existed long before we ever had a patch. Someone, even if it were just the vendor, was aware of the issue in advance of the patch's release. Vulnerabilities for which there are no patches are known as 0day or zero-day vulnerabilities.

Although it is unlikely that your organization will be targeted with a zero-day vulnerability (outside of custom web applications), these vulnerabilities exist. What is the point? Why do we care? We need to appreciate that a modern information system's risk is never practically ever going to be zero.

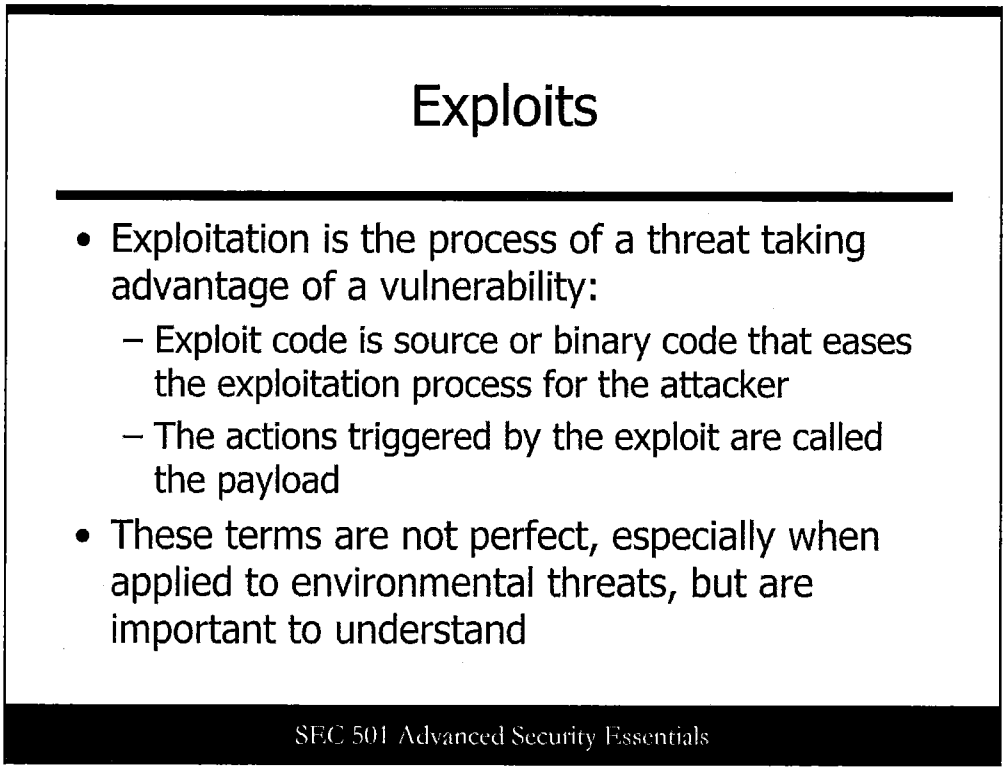

Exploits

Having already used this term, the meaning of an exploit is likely clear. An exploit is the means by which a threat exercises a vulnerability. An attacker (threat source) exploits a vulnerability. In addition to exploit used as a verb for understanding the risk equation, it is also necessary to understand the term, exploit code. Exploit code is source or binary code that eases the ability for an attacker to exploit a vulnerability.

When the concept of vulnerability scoring is introduced later, the existence of publicly available exploit code is one of the items that can increase a vulnerability's overall score.

Another concept related to exploits and exploitation is that of a payload. The payload, in exploitation terminology, is what action the attacker wants to carry out as a result of the exploitation. Getting a shell, adding a user, and exfiltrating files are some examples of payloads. Payloads are part of the post-exploitation portion of an attack.

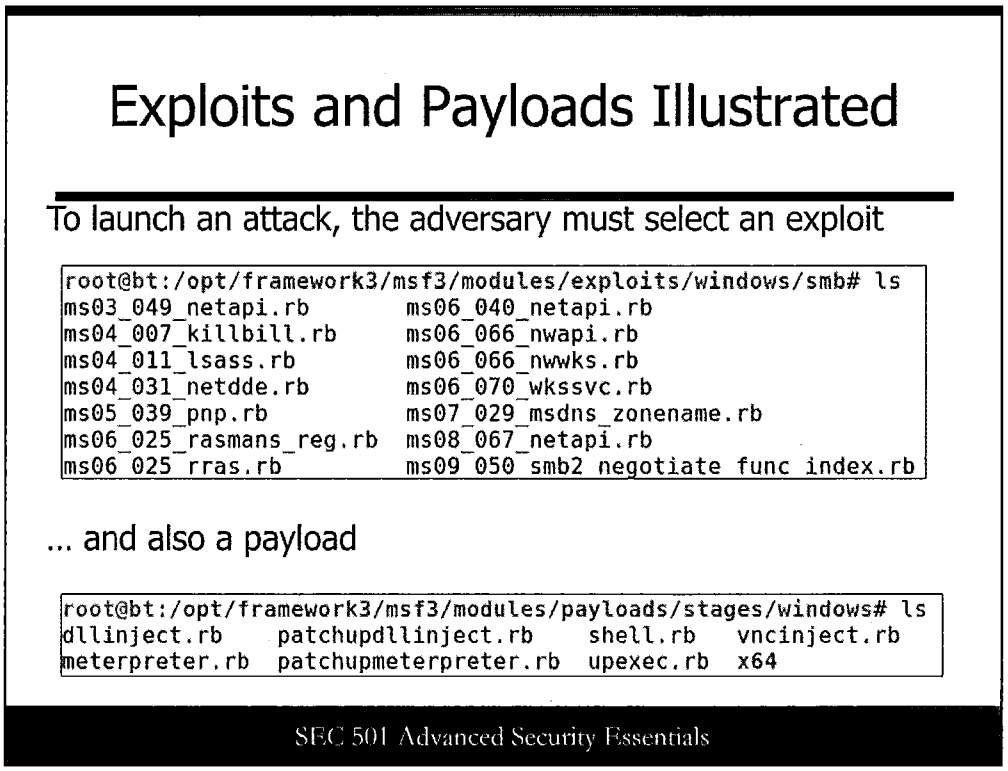

Exploits and Payloads Illustrated
To illustrate exploits and payloads a bit better, screen shots are provided. The directory contents shown correspond to Metasploit exploits and payloads.

As can be seen in the upper screen shot, these exploits are tied directly to particular Microsoft SMB vulnerabilities that have available patches.

In the lower screen shot, we see a few options of what actions the attacker might trigger: command shell access; VNC (remote GUI) access, uploading and executing a binary of the attacker's choosing; and the incredibly advanced Meterpreter payload.

For additional information on the outstanding open source Metasploit project, see http://www.metasploit.com.

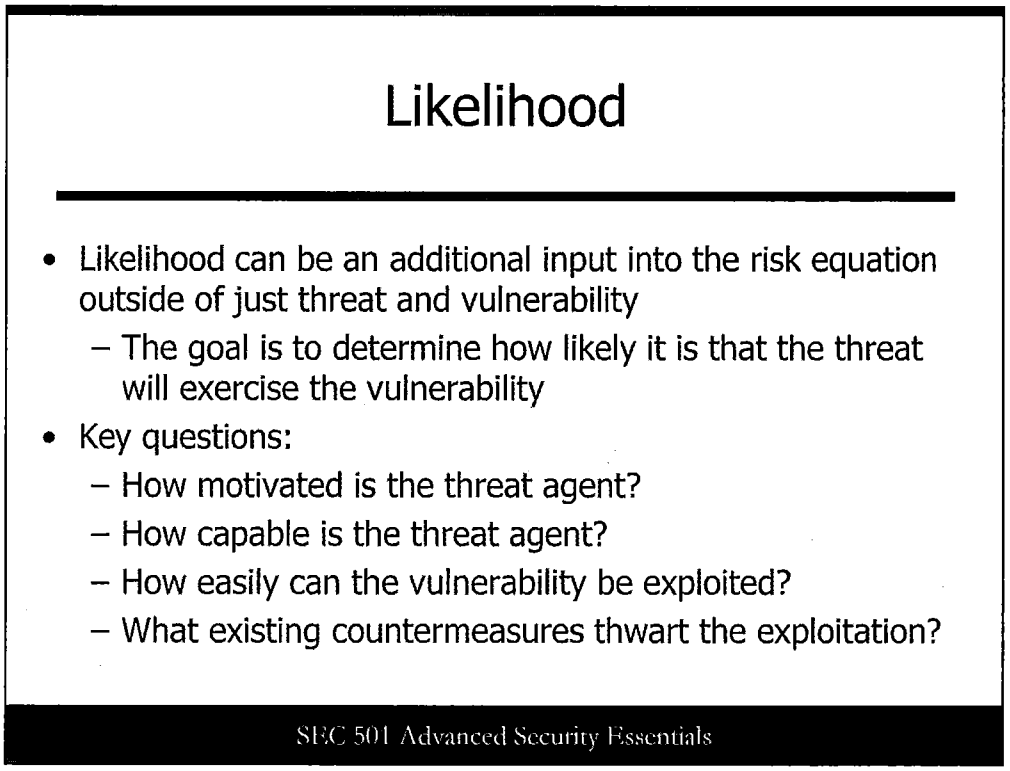

Likelihood

Merely understanding the concepts of threat and vulnerability is not sufficient for performing risk assessments. Likelihood is another key concept that helps inform our risk management.

Likelihood assessments attempt to determine how likely successful exploitation of the vulnerability will be. Several factors inform the likelihood of successful exploitation. These factors include threat motivation; threat capabilities; ease of exploitation; and existing controls and countermeasures.

Understanding how likely a scenario is can help to determine what an appropriate risk-based response will entail. The more likely a scenario, the greater the risk.

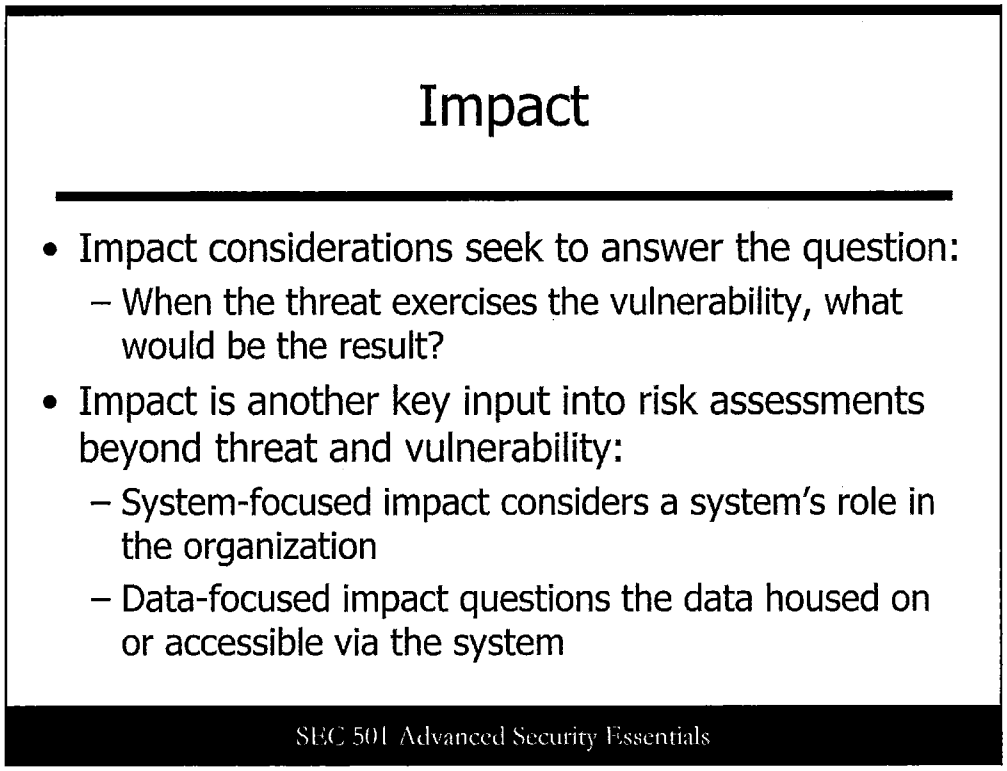

Impact

A final concept for understanding the risk equation is that of impact. In addition to the likelihood, impact is a critically important concept for determining risk that is not overtly stated in the simplistic Risk = Threat x Vulnerability equation.

Impact attempts to determine what the outcome of successful exploitation would be. Impact determination will necessarily take into consideration the information system in question as well as the data housed or processed by the information system.

The importance of impact is obvious. Two systems with the same vulnerability, accessibility, and subject to the same threat characteristics will not always warrant the same level of response from a security-perspective. The system's criticality to the organization will make a significant difference when determining what countermeasures are ultimately employed.

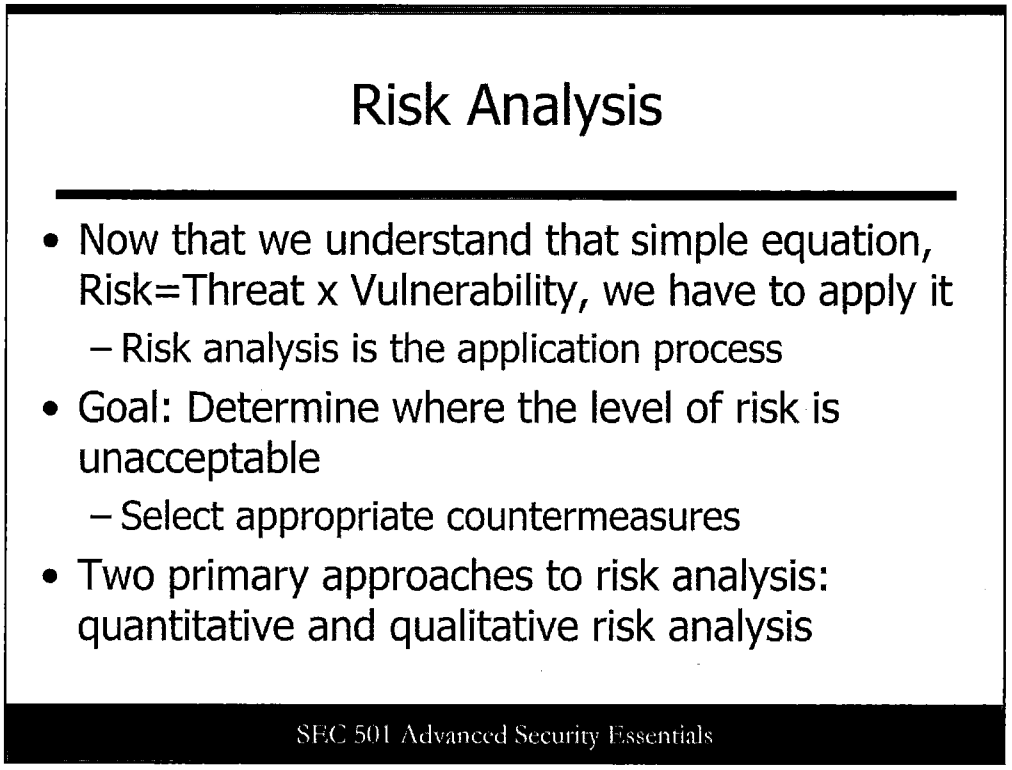

Risk Analysis

Now that the basic concepts that support the definition of risk are understood, we turn our attention to the process of risk analysis. We don't simply calculate risk to know our level of risk. We analyze risk so that we can understand it and make informed decisions about whether and which countermeasures need to be employed.

The two primary approaches to risk analysis are the quantitative approach and the qualitative approach. There is no right approach, as each has its own merits.

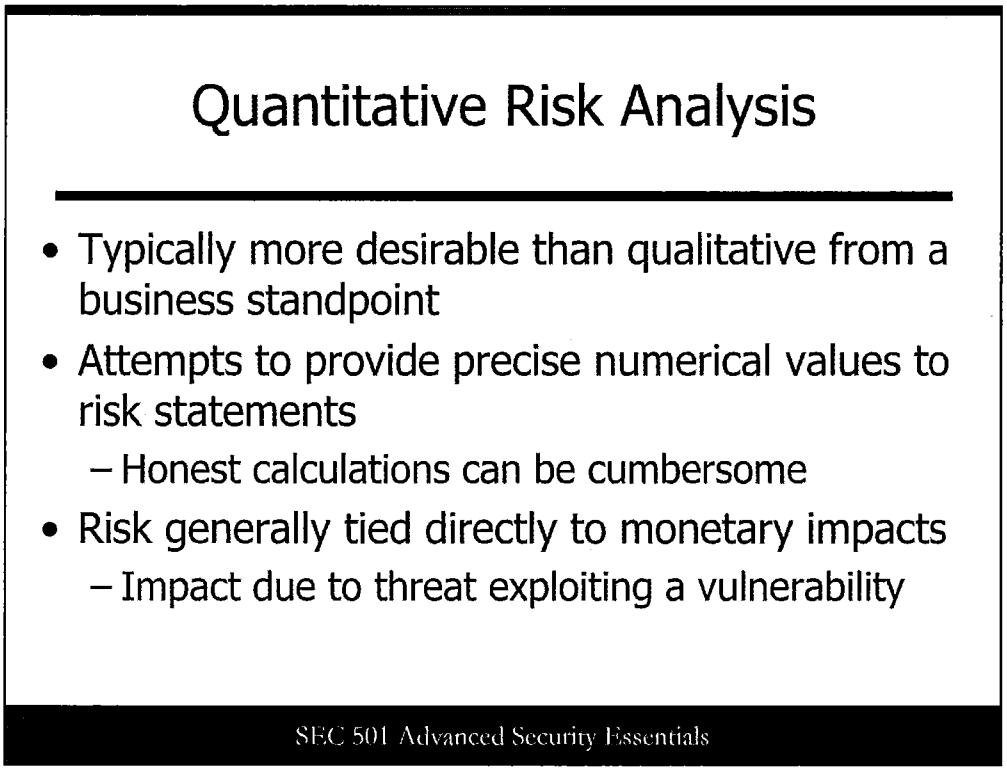

Quantitative Risk Analysis
Quantitative risk analysis is often thought to be preferable by those in business but is not always the best approach for an organization.

As expected, quantitative analysis is numerically based and is almost always tied directly back to money. For example, impact determination would be characterized by the cost to the business. Tying the results of risk analysis back to dollars and cents is quite appealing for most organizations.

However, performing a thorough analysis that yields honest calculations can be quite difficult for almost every organization. Determining with fidelity the value of the inputs into the risk equation is terribly problematic. And unlike many other industries' risk-based metrics, information security data is notoriously lacking and inconsistent.

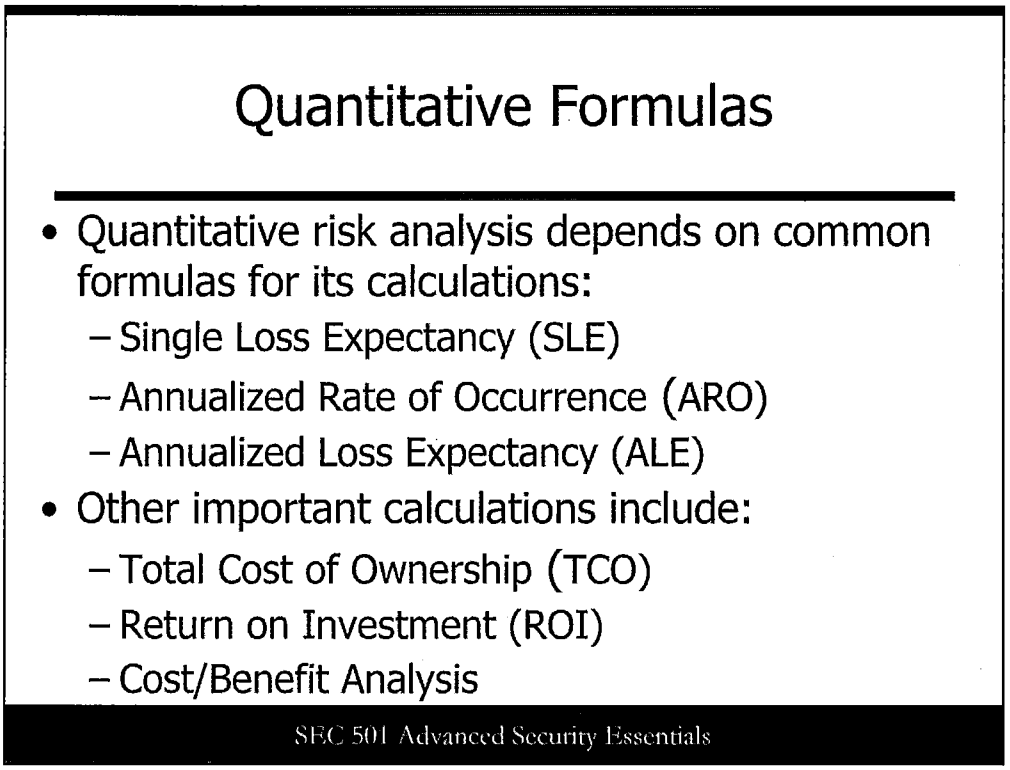

Quantitative Formulas

Quantitative risk analysis focuses on numbers, bringing with it a number of formulas and metrics that should be understood. The following are some of the key formulas used:

- Single Loss Expectancy (SLE) - SLE = EF x AV
- Annualized Rate of Occurrence (ARO)
- Annualized Loss Expectancy (ALE) – ALE = SLE x ARO

Additional calculations that are important to quantitative risk analysis as well as to other general security considerations are:

- Total Cost of Ownership (TCO)
- Return on Investment (ROI)
- Cost/Benefit Analysis

We dig deeper into these calculations shortly.

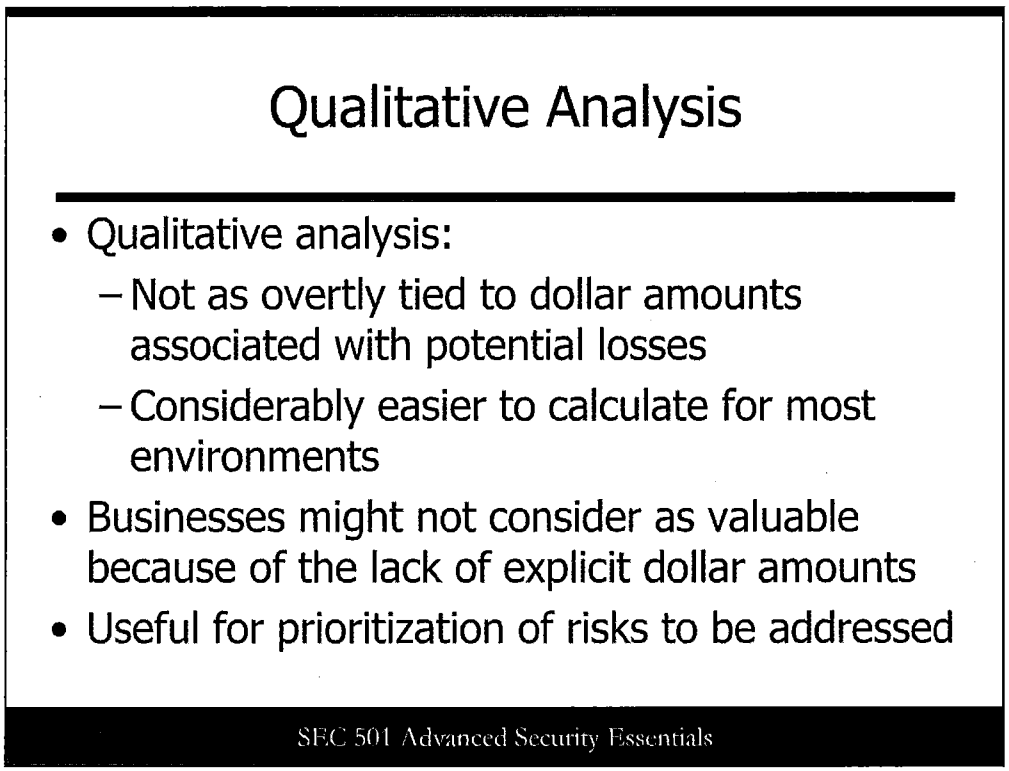

Qualitative Analysis
On the other end of the spectrum from quantitative risk analysis is qualitative risk analysis. The focus of qualitative risk analysis is not to produce detailed numbers directly related to actual monetary figures.

Qualitative analysis is not as focused on precise calculations of money, which can make it considerably easier to calculate. However, many businesses prefer the quantitative analysis's focus on money, as it is far easier to plug those numbers into budgets and projections.

Still, qualitative risk analysis should not be ignored simply because businesses would prefer to get dollar amounts. The truth is, a lot of the dollar amounts determined by quantitative analysis are often wild guesses. Given the relative ease with which qualitative analysis can be performed, it might actually be preferable.

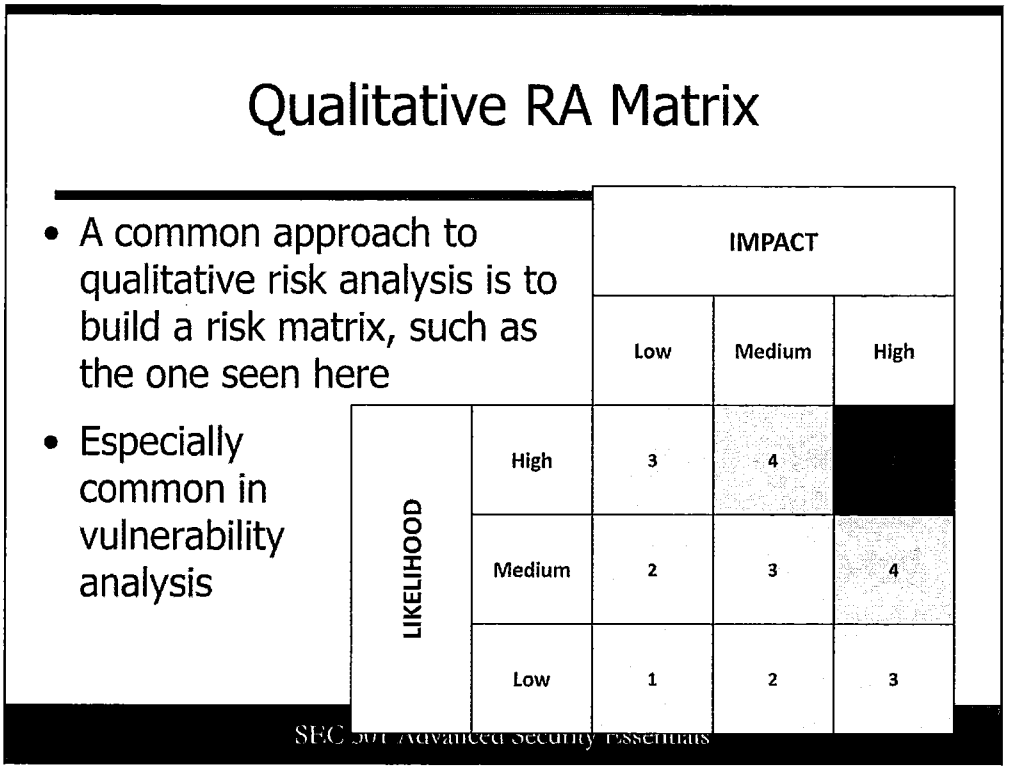

Qualitative RA Matrix

One of the key tools for performing qualitative risk analysis is the Risk Matrix. The Risk Matrix illustrates the continuum of risk (in this case from high to low) by plotting the Likelihood and Impact associated with a threat vulnerability pair.

Will populating the Risk Matrix yield dollar amounts associated with impacts that can be used directly in ROI calculations? No. But if your goal is to identify the most significant risks to an organization, this simple tool can prove extremely effective.

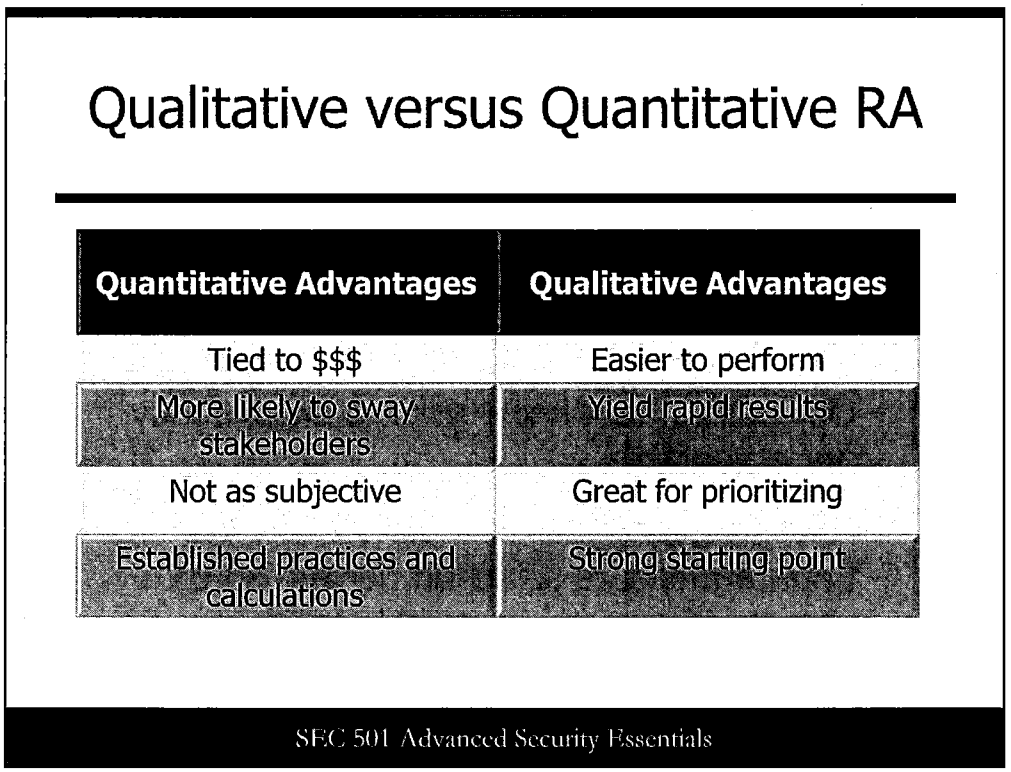

Qualitative Versus Quantitative RA
This chart highlights the relative advantages of quantitative and qualitative analyses.

Quantitative Advantages	Qualitative Advantages
Tied to $$$	Easier to perform
More likely to sway stakeholders	Yield rapid results
Not as subjective	Great for prioritizing
Established practices and calculations	Strong starting point

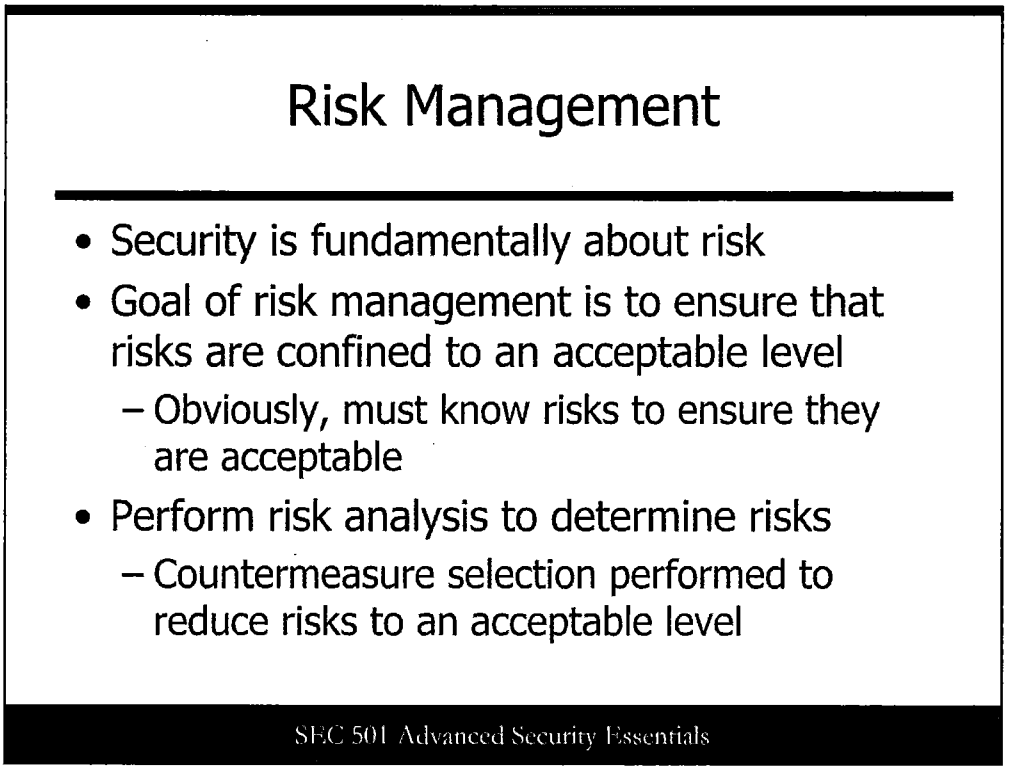

Risk Management

As previously discussed in the introduction to risk, much of security professionals' jobs are centered on dealing with issues of risk management. To that end, risk analysis is a key process that the security professional needs to know.

Though we have already discussed quantitative and qualitative risk analysis, we will continue reviewing risk analysis in more detail. Ultimately, the goal of analyzing risk is to understand the current state of risk and make informed decisions about where items need additional scrutiny.

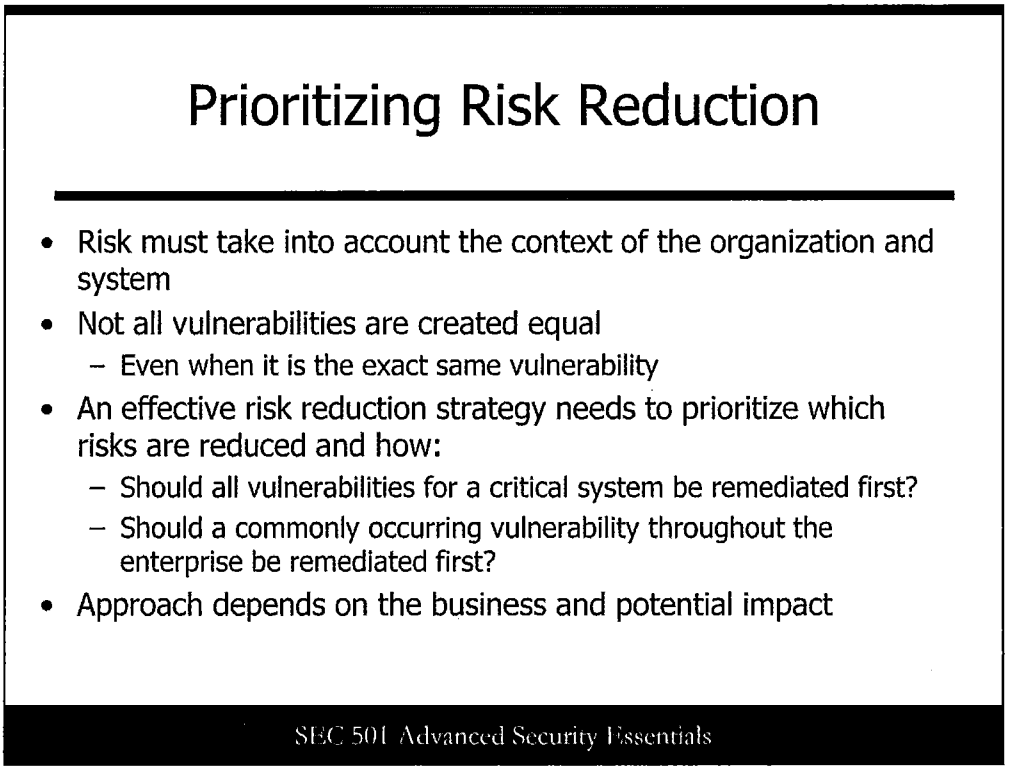

Prioritizing Risk Reduction

Naturally, it would be preferable if there were no risk at all. Unfortunately, organizations have neither unlimited time or budget to address even all known vulnerabilities. So, risk reduction must be prioritized. This is an additional output of our risk analysis: which are the most significant risks.

Effective risk management must prioritize a risk reduction strategy taking into account all the inputs into the risk analysis equation as well as the time and cost to implement countermeasures capable of eliminating or reducing risks to an acceptable level.

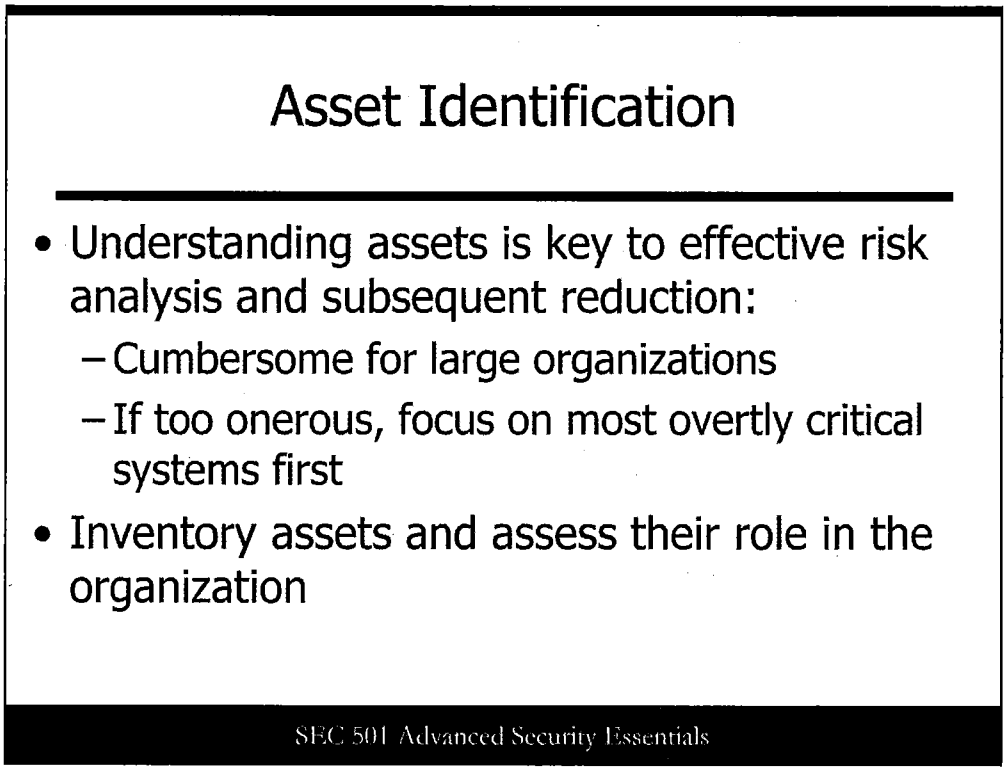

Asset Identification

To manage risk, the risks must be understood. For the risks to be understood, we have to appreciate the assets on which the vulnerabilities exist. Asset identification is a key phase of the risk analysis process.

Simply having an accurate inventory of information systems proves difficult for many organizations, let alone understanding the impact was the information system to be compromised. If too onerous, organizations would do well to first focus on asset identification for critical information systems.

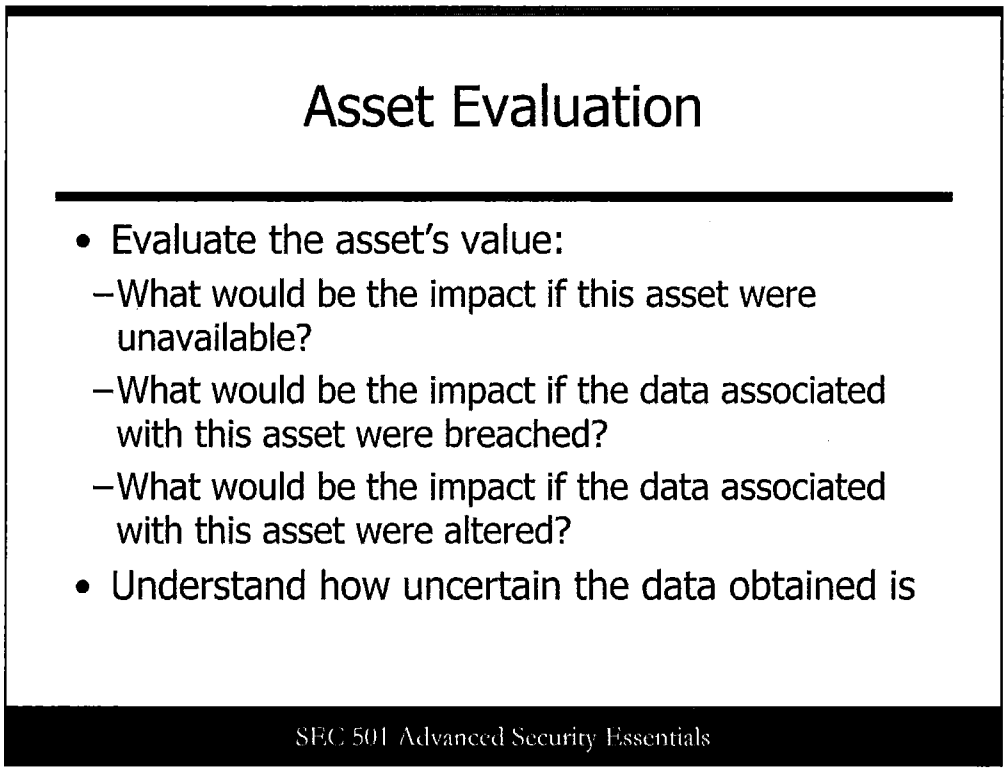

Asset Evaluation
Beyond merely identifying the information systems that exist in an organization, their role needs to be appreciated.

Key questions pertaining to the identified assets are:

- What would be the impact if this asset were unavailable?
- What would be the impact if the data associated with this asset were breached?
- What would be the impact if the data associated with this asset were altered?

An additional consideration is to appreciate the lack of certainty associated with the answers to these questions. This inherent uncertainty is one of the major challenges associated with quantitative analysis.

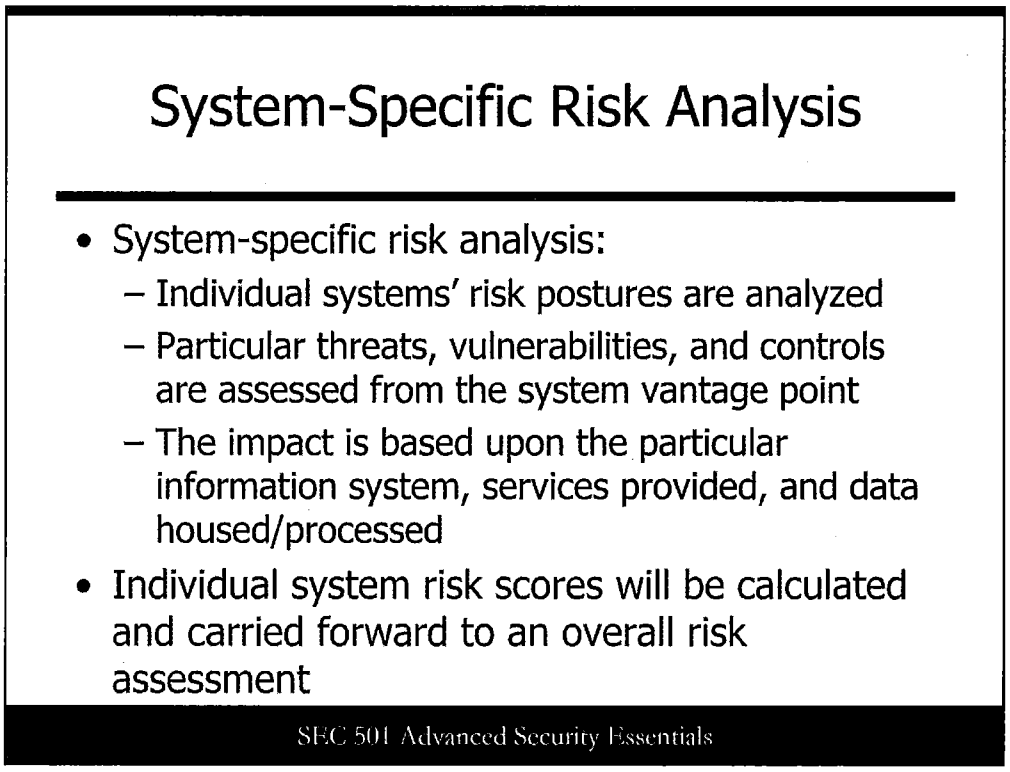

System-Specific Risk Analysis

Though discussed abstractly, particular threats and vulnerabilities are analyzed in light of specific information systems as opposed to in general. So the calculations that have been discussed from a risk analysis standpoint will have to be performed many times over.

Thankfully, some of the data can be reused after calculating once, but this is still an onerous process. Obviously, to understand a system's risk requires a detailed understanding of the asset's role and value to the organization first. This information will inform us about the potential impact associated with exploitation of vulnerabilities affecting this system.

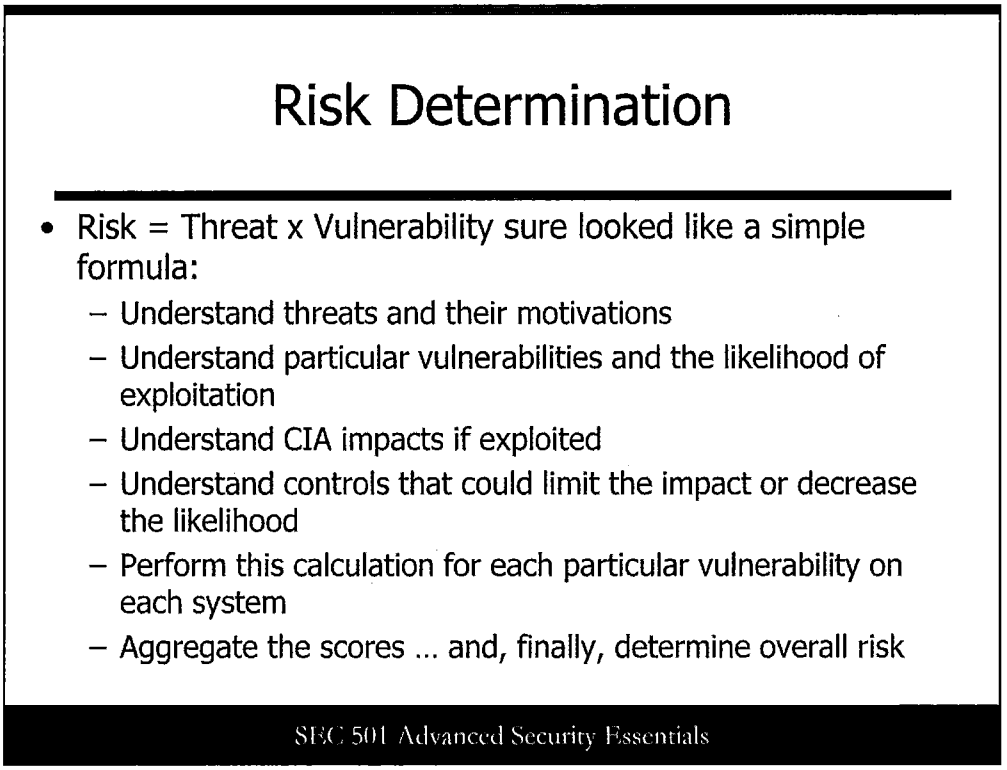

Risk Determination

Now that all the components have been identified and analyzed, actual risk determination can be performed.

Risk = Threat x Vulnerability sure looks like a simple formula, but now we can appreciate all that goes into this calculation.

Threat involves understanding threat sources, their motivations and capabilities. Also we have to understand particular vulnerabilities and the likelihood of successful exploitation. Presuming successful exploitation, we must understand CIA impacts. We must also assess the current controls that could limit the impact or decrease the likelihood. With all this, we can now perform the risk calculation for each particular vulnerability on each system. Then we just have to aggregate the scores … and, finally, determine overall risk. And then, we get to do something about it, or not.

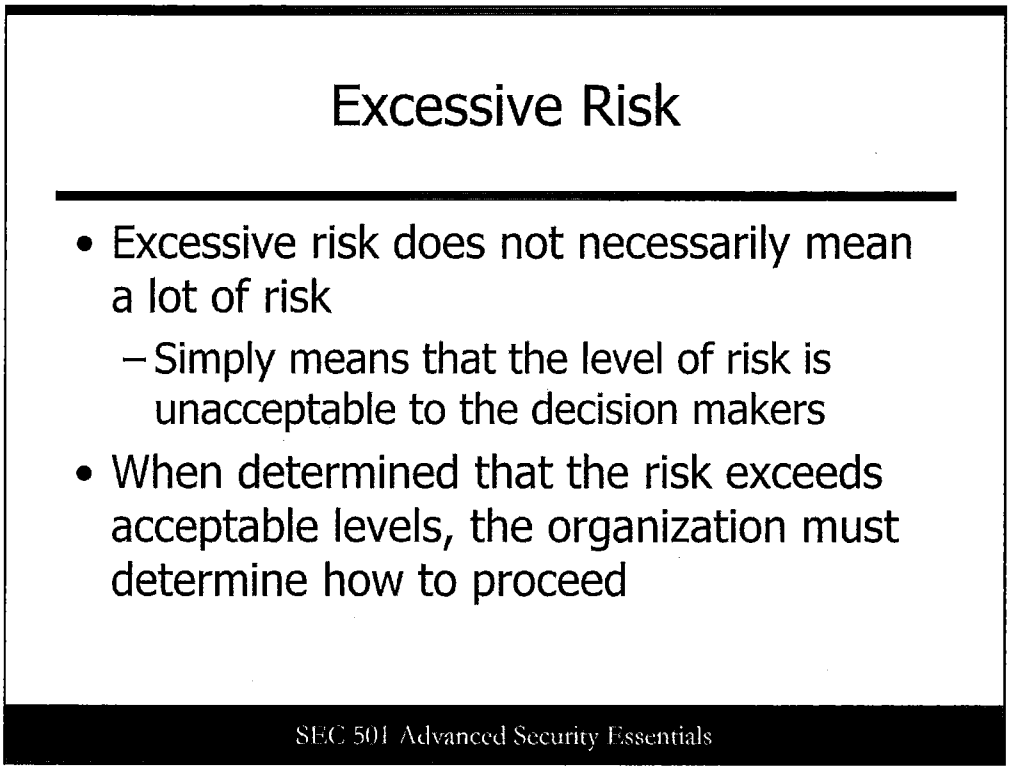

Excessive Risk
So, you have spent many sleepless nights and finally completed the individual and overall risk analysis. Management will review and determine whether the determined risk level is acceptable. If not, then the risk is excessive, which doesn't mean a lot of risk, but rather simply that the risk exceeds acceptable levels.

If risk is determined to be excessive, the organization must determine what the response will be. There are several different valid responses. Many expect that the default response to excessive risk would simply be to decrease the risk directly. This is one, but not the only, valid response to excess risk.

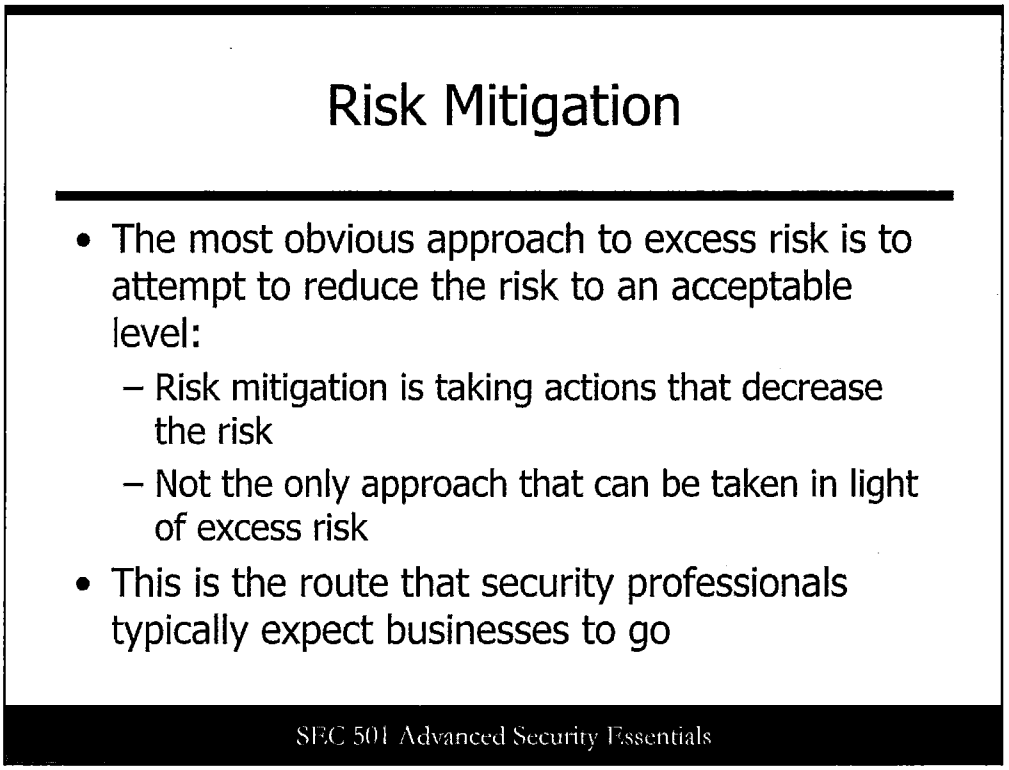

Risk Mitigation

The most obvious approach to excess risk is to attempt to reduce the risk to what is perceived to be an acceptable level. Although this tends to be the route security professionals advise and expect the business to pursue, it is not the only action that the organization can take.

Let's explore some of the other approaches as well.

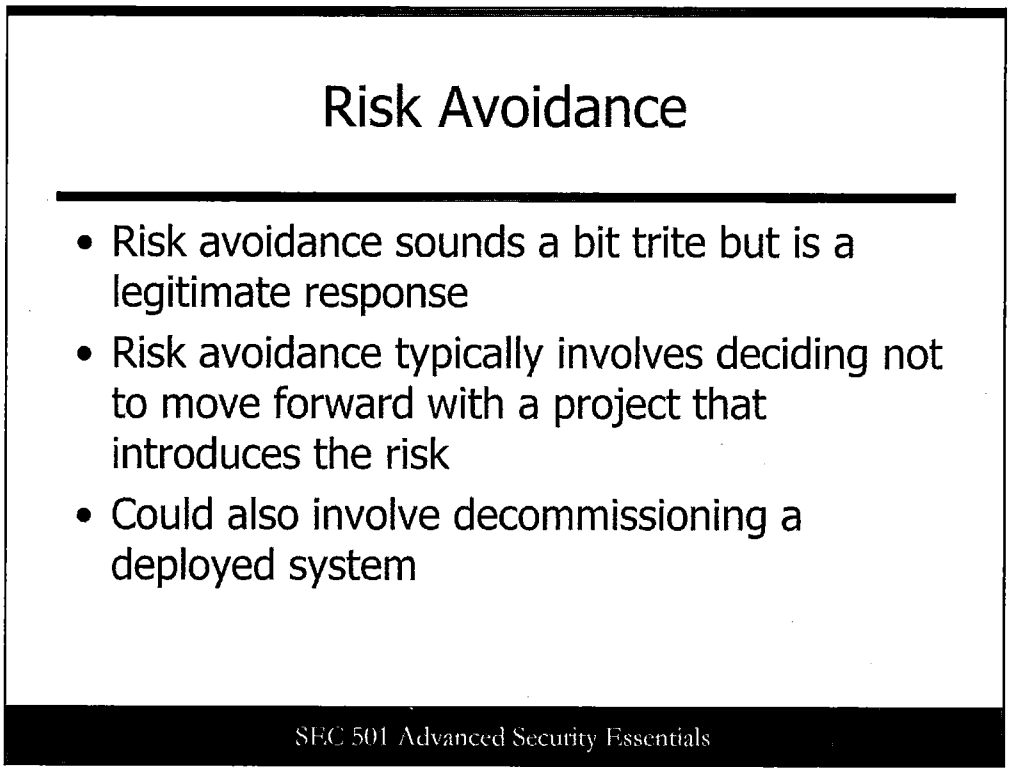

Risk Avoidance
Well, how about we just avoid that risky behavior. It sounds a bit childish to simply say, "Well let's just avoid that big scary risk," but it is actually a legitimate response.

Risk avoidance in an enterprise typically involves declining not to move forward with a project that introduces the unacceptable level of risk. This can involve choosing another option that does not include the same degree of risk, or simply doing nothing.

Risk avoidance with respect to systems can involve the decommissioning of a deployed information system.

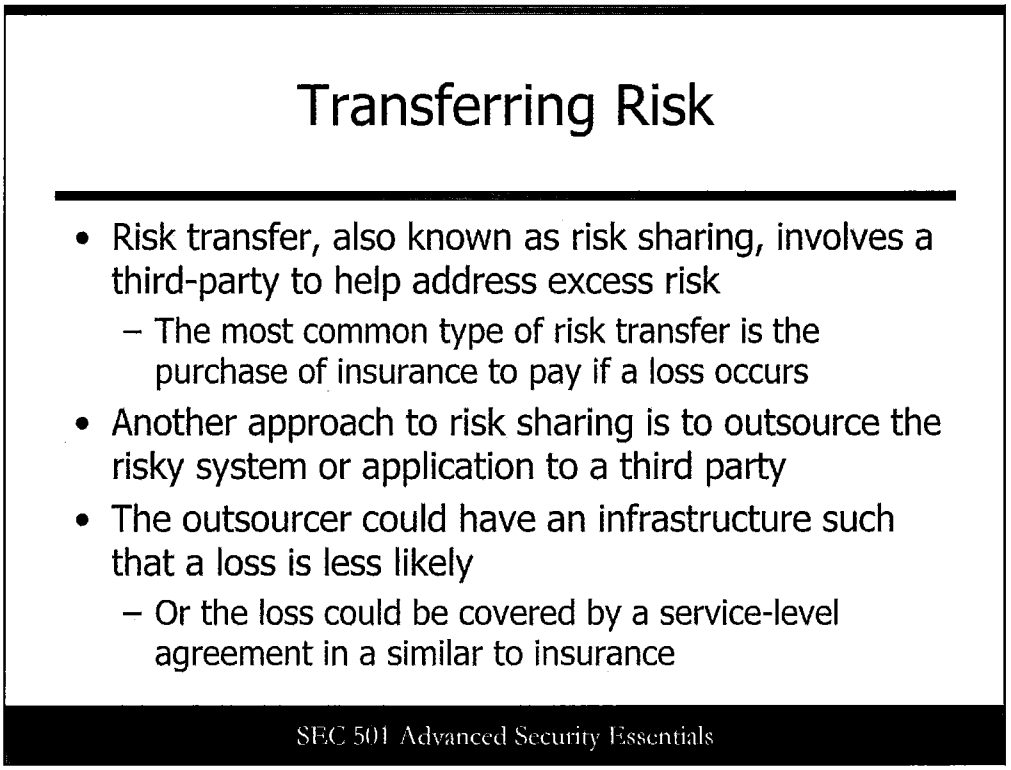

Transferring Risk

Another approach to dealing with excessive risk is called risk transfer, which is also referred to as risk sharing. The idea is to involve a third party to help address the risk. The most common type of risk transfer is through the purchase of insurance to pay if the loss occurs that is too likely for the organization to stomach.

Another approach to risk transfer is to outsource the risky system or application to a third party for development, management, hosting, or whatever the risky issue happens to be. The idea, in this case, is that the third party will be assuming the risk on behalf of the organization. It could be that the outsourcer, by nature, has more compensating controls that decrease the risk's likelihood of being realized. Or it could be that the loss is defined as part of a service-level Agreement and that the third party is more willing to accept the risk.

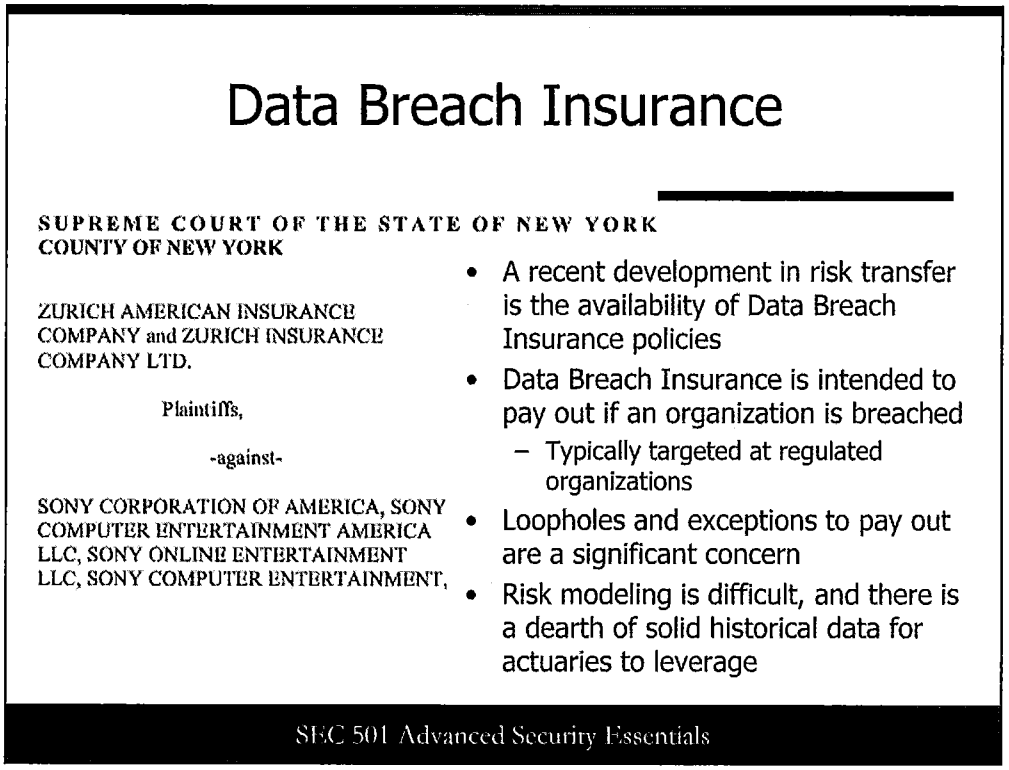

Data Breach Insurance

One type of risk transfer that is discussed with increasing regularity is that of data breach insurance. The idea is that organizations rather than deploying infrastructure and countermeasures sufficient to prevent a data breach (difficult to quantify), opt instead to insure against the potential loss.[1]

The business of insurance is to allow for one to pay a certain amount in regular premiums to offset the cost associated with a particular uncertain loss such as, in this case, a data breach. Data breach insurance is intended to pay should the organization be breached. Be aware that loopholes and exceptions pay outs are currently a significant concern. Also important, many companies mistakenly assume that their insurance covering loss of business records would cover them in the instance of a data breach, which is often not the case.

Risk modeling is difficult, and there is a dearth of solid historical data for actuaries to leverage.

The preceding image is taken from the lawsuit filed by Zurich against Sony after Sony's major breach in 2011.[2] Zurich was suing to not have to pay claims related to the breach as Sony did not specifically have a policy covering information security incidents.

[1] http://bits.blogs.nytimes.com/2011/12/23/insurance-against-cyber-attacks-expected-to-boom/ (http://cyber.gd/414_18)
[2] http://ace-insurance-litigation.com/document/zurich-american-insurance-company-et-al-v-sony-corporation-america-et-al-complaint (http://cyber.gd/414_19)

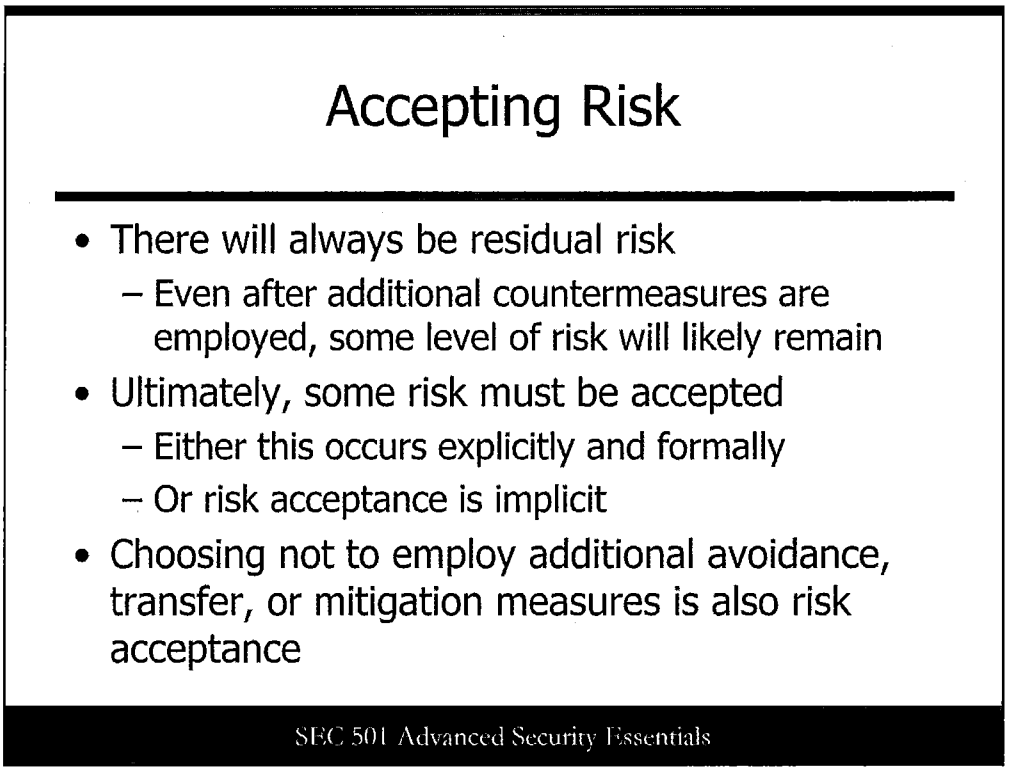

Accepting Risk
At some point, the organization will have to accept a certain level of risk. There will always be residual risk even if mitigating countermeasures are leveraged. Either this occurs explicitly and formally or risk is accepted implicitly, by choosing not to employ additional avoidance, transfer, or mitigation measures.

Accepting the risk does not mean that the organization simply did not perform the analysis and accepts whatever risk they might have; that is the ostrich approach to risk management.

Mitigating Risk

The most common outcome, and the one security professionals expect, is that risks that are deemed unacceptable will be mitigated. There are many and varied ways to mitigate risks.

Mitigation can be:

- **Threat-oriented**: Focused on reducing motivation of the threat agents
- **Vulnerability-oriented**: Reducing the vulnerabilities that the threat can exploit
- **Impact-oriented**: Reducing the overall impact that exploitation entails
- **Likelihood-oriented**: Reducing the likelihood that the threat can exploit the vulnerability

Some examples of how these might be accomplished: Reducing threat motivation could be accomplished through deterrents (such as increased rates of prosecution for crimes). Vulnerability-oriented mitigations are often the most common and can be achieved by patching or installing host or network-based countermeasures.

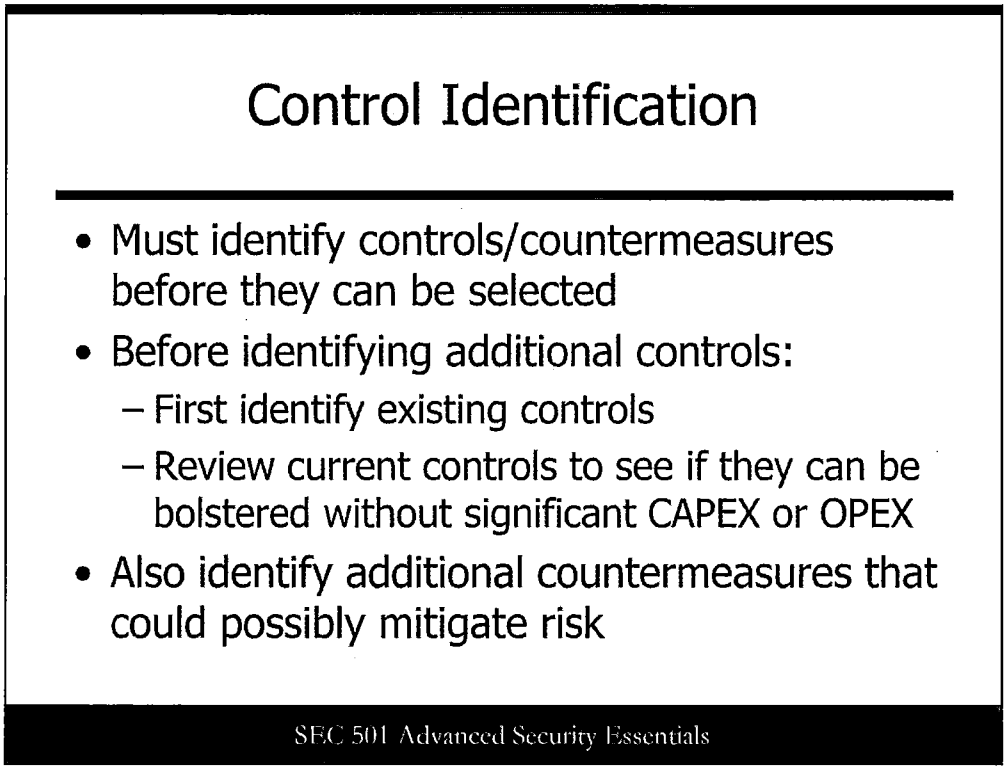

Control Identification

Before selecting potential controls and countermeasures for risk reduction, they must be identified. Naturally existing countermeasures must first be enumerated and analyzed. The analysis of these countermeasures should determine not only that they exist and are in working order, but also classify them by the type of control they represent: preventive, detective, deterrent, or directive.

Also, the existing controls should be reviewed with an eye to whether they can be bolstered without incurring significant capital or operational expense. Finally, additional countermeasures beyond the current should be identified for evaluation.

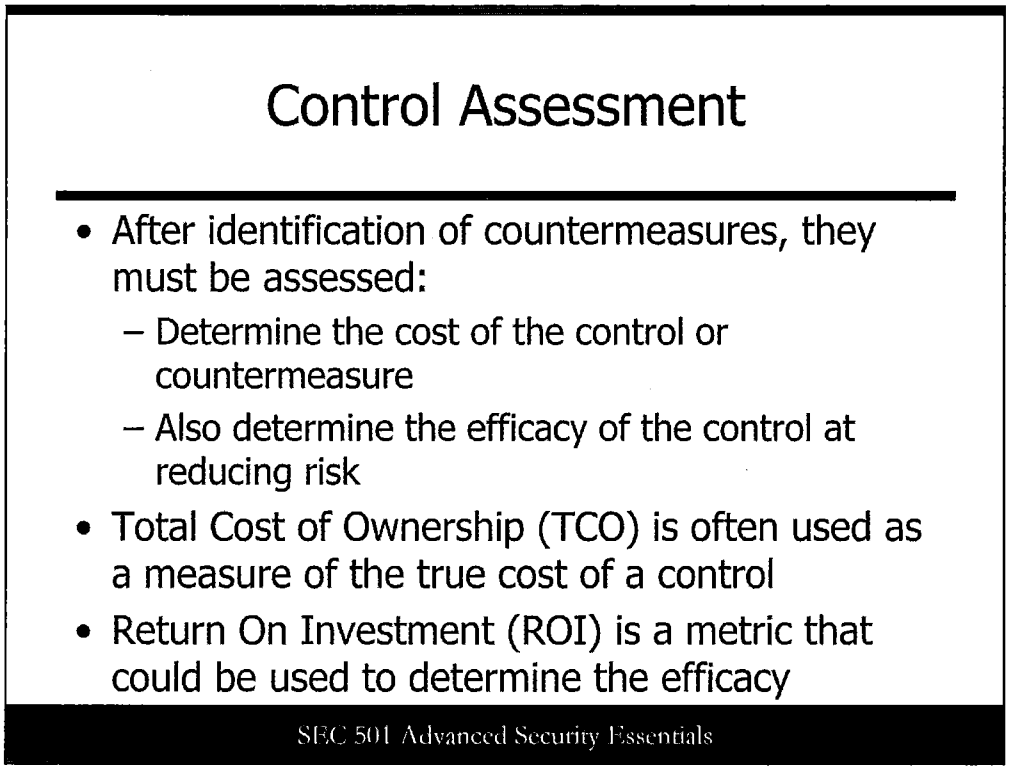

Control Assessment
After identification of additional controls, they must be assessed. The goal is to determine both the cost of the control or countermeasure as well as its efficacy. Effectively ,a cost benefit analysis is performed on the countermeasure to determine which countermeasure(s) to employ or whether the countermeasures should be adopted.

Two metrics that are often referenced for these types of assessments are Total Cost of Ownership (TCO) and Return On Investment (ROI). TCO attempts to capture the true cost of adopting something, beyond merely the capital expense. ROI attempts to determine how financially worthwhile something is based on how much money will be made based on the money spent. ROI is typically difficult for security countermeasures, as security will not make an organization money, but could only prevent future potential losses.

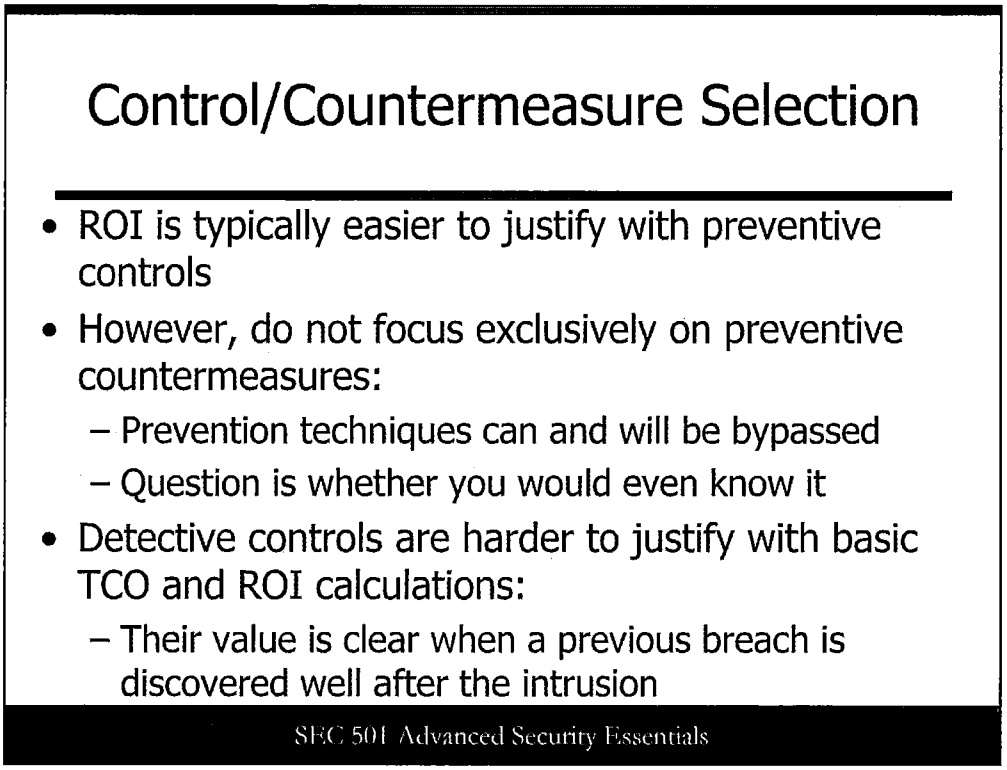

Control/Countermeasure Selection
If only ROI is employed to determine which control or countermeasure to employ, then typically preventive controls will almost exclusively be selected. Calculating a positive ROI for anything in security is difficult, as we have to justify based on reducing future potential losses. This is especially difficult for nonpreventive controls.

However, appreciate the all preventive controls can and will be bypassed. There is no security silver bullet that magically stops all attacks. Detective controls are vital. They become absolutely essential when performing incident or breach response.

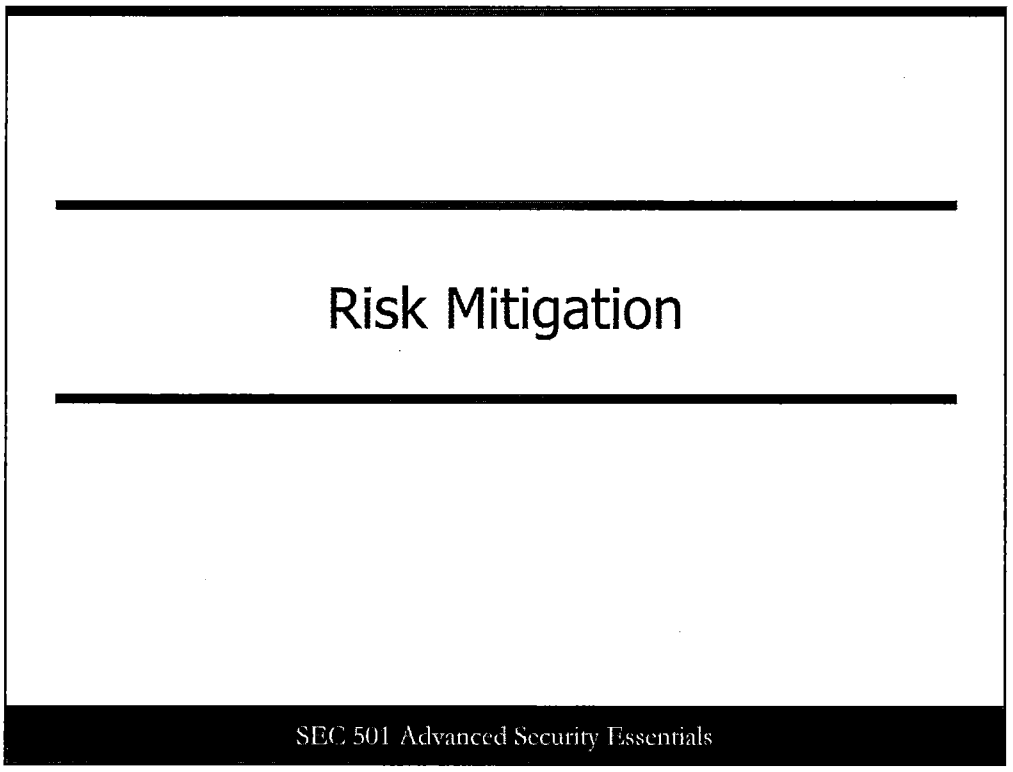

This page intentionally left blank.

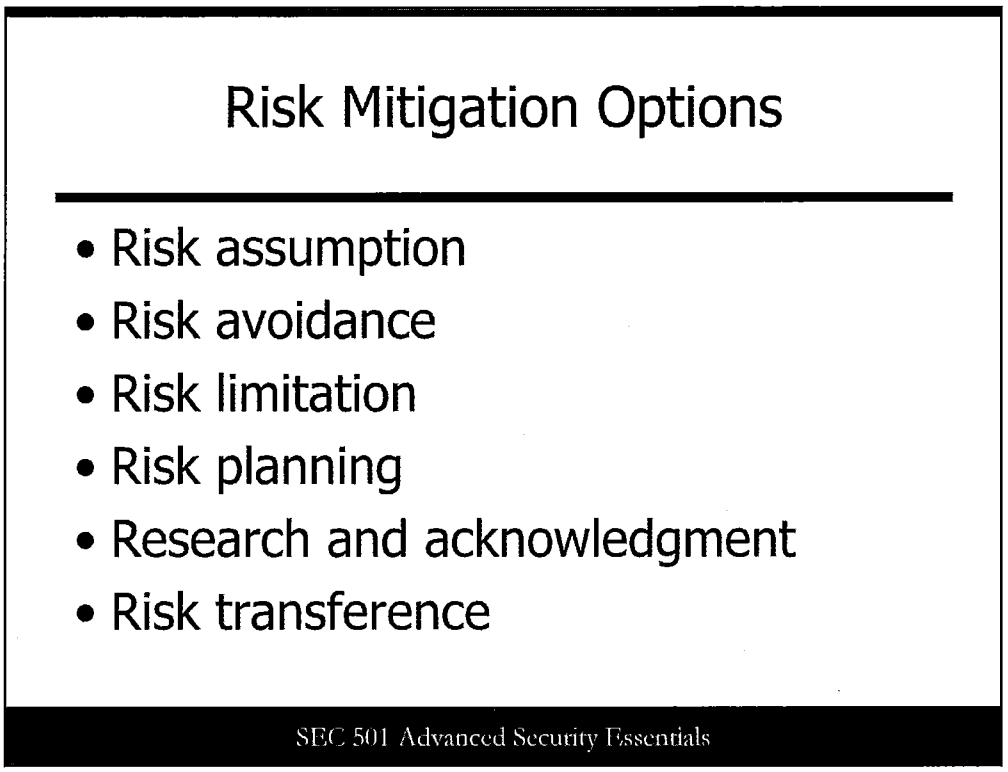

Senior management has the important task of identifying methods that will reduce any potential risk to the organization. This list represents action options they may use in this endeavor.

- **Assumption**: Accept risk as is, and implement available control measures in hopes of reducing or maintaining risk at an acceptable level.
- **Avoidance**: Remove or shut down operational entity with risk to avoid risk.
- **Limitation**: Implement controls that can reduce or minimize any adverse effects of threats toward vulnerabilities.
- **Planning**: This is a guidance step in that it involves prioritization, implementing, and maintenance of control measures.
- **Research and Acknowledgment**: Starts with acknowledging a vulnerability exists and researching for control measures to fix the problem.
- **Transference**: Offsetting loss through compensation of other operations. The risk is still there; other methods, such as insurance, is incorporated to cover the loss.

This is one way to manage risk through a step process to highlight milestones to reduce vulnerabilities.

- **Step one: Prioritize actions.** Itemization of vulnerabilities and security measures have been done, and now the analysis is focused on what to do first.
- **Step two: Evaluate recommendations.** Through paper or test implementation efforts, sample and evaluate recommendations to allow for selection of the best fix. The next step works hand in hand with this one: cost.
- **Step three: Cost-benefit analysis.** Selection of recommendations can now be done by weighing the risk and loss in comparison with the cost of the recommendation, which will lead to the most significant return or savings. It all comes back to money; you can't implement it if you can't afford it. If you can't afford it, and the risk remains high, new recommendations or methods of operation must be pursued.
- **Step four: Select control.** Purchase or develop the protective mechanism decided upon and prepare for the next step.
- **Step five: Assign responsibility.** This assignment concerns the management of the control and its implementation, not the risk of loss. Someone has to take control and ensure the control mechanisms are understood and implemented appropriately.
- **Step six: Develop a safeguard plan.** If the implementation of a control mechanism is conducted, and it unexpectedly fails, how will you recover? What is the fail-safe plan to return to normal operations while a fix is researched, redesigned, and implemented?
- **Step seven: Implementation.** Put the plan into operation. Monitor it, adjust it, document it, and continue to manage it.

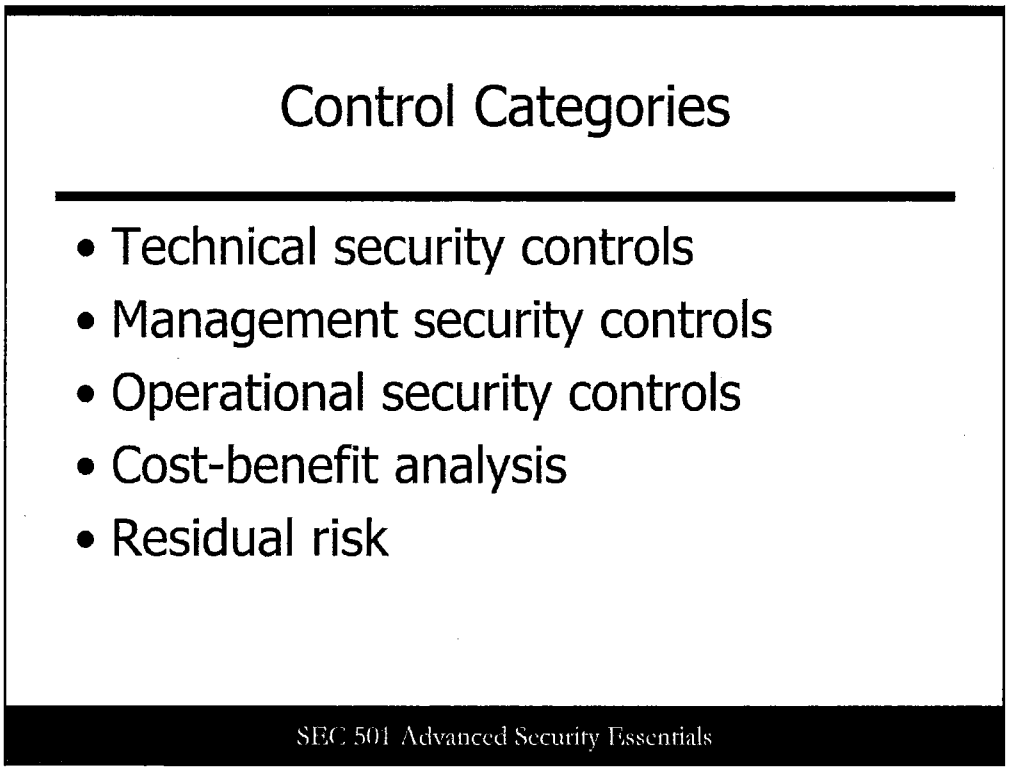

Categories of control are means of labeling and documenting your efforts. They allow for the shifting of funds based on their purpose within a security posture.

- **Technical:** Simply stated, are those of hardware or software security-related functions.
- **Management:** Policies, guidelines, checklists, and reporting. This also includes training issues.
- **Operational:** Auditing, test and evaluation, management issues such as updates and patches, and monitoring of the controls.
- **Cost-benefit analysis:** Implementation and selection of controls throughout the life cycle, which may warrant changes in controls due to cost efficiency and reduced risk.
- **Residual risk:** After all the approved and implemented control measures are in place, this is the risk that still remains.

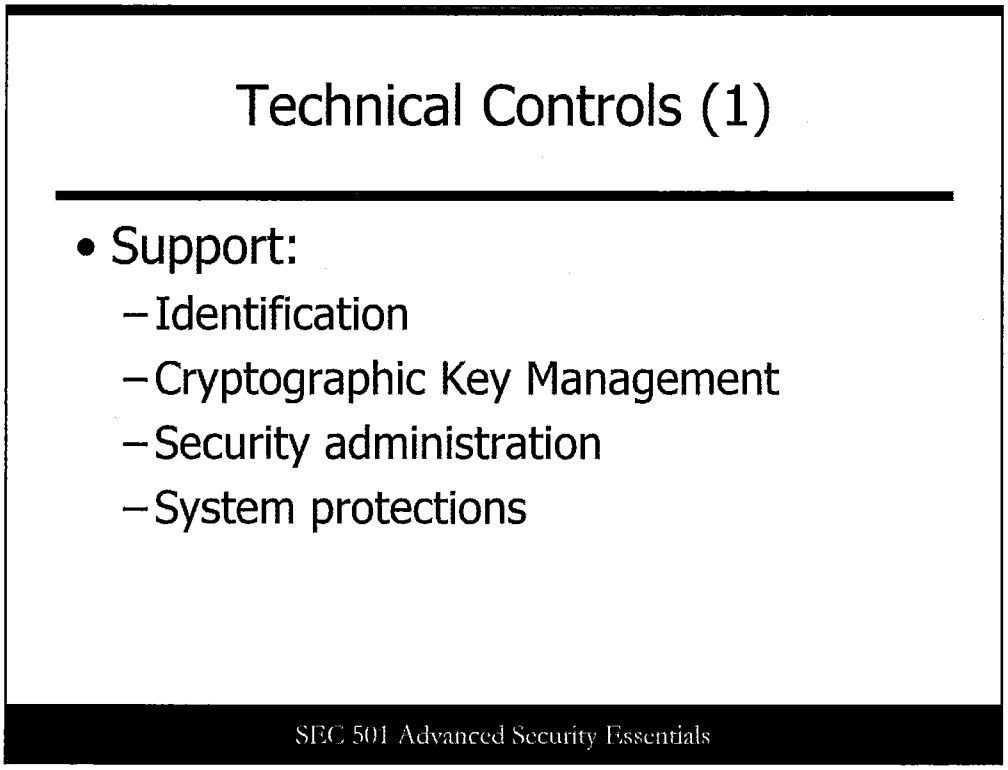

Technical Controls (1)

Technical controls center around support functions, preventive controls, and detect and recovery. Let's break these down further.

Support controls are normally already installed within operating systems and applications. They should have been included in the design process and allow users and managers alike to incorporate various security functions. Of these controls, the most obvious is the one of identification. This control measure separates the employees, who should have access, to the outsider, who should be denied access. It also includes key management, tokens, and most security applications and is essential for other security controls Media Access Control (MAC) and Discretionary Access Control (DAC) to be implemented.

System protection warrants additional explanation. Implementation of data classification steps to manage least privilege, process separation, security layering (in respect to applications), trust issues, and object reutilization (object reuse) are all system protection issues.

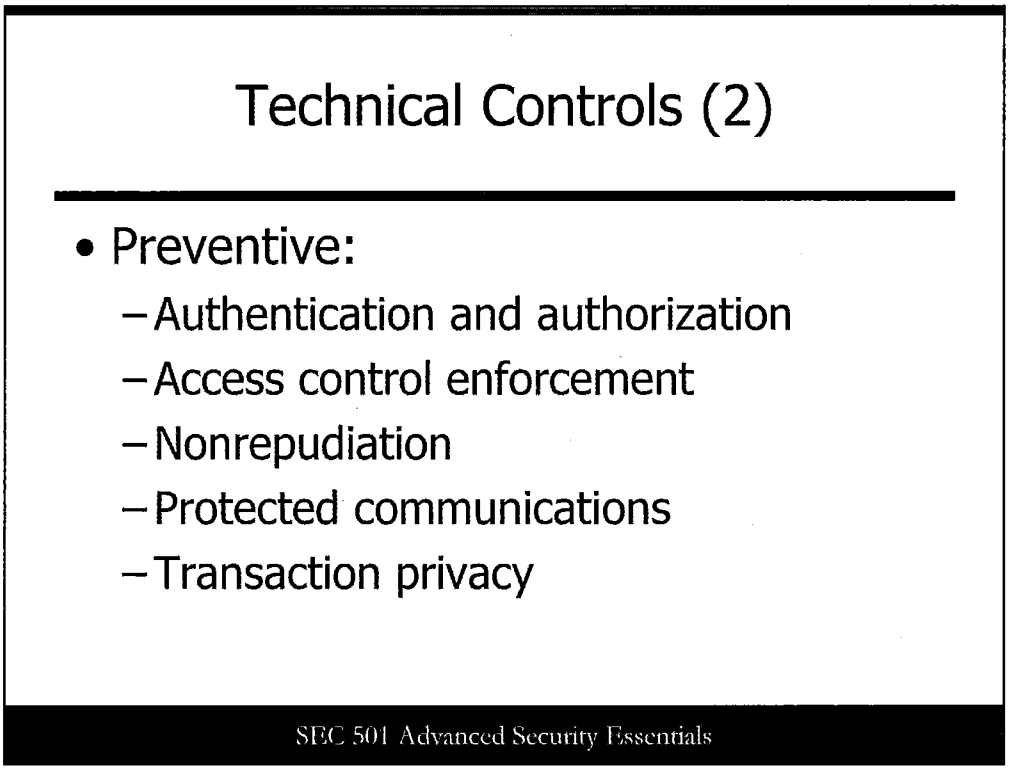

Technical Controls (2)

Preventive controls, as their name states, are control mechanisms and processes, used to prevent a loss before it occurs. These controls work with Support Controls in that the identification that has already occurred is now validated through the use of authentication and authorization. Authentication handles the security level validation, whereas authorization assigns or applies appropriate access allowances to users or groups.

Nonrepudiation is a security step to remove all plausible deniability from the senders and receivers. This control falls under preventive controls due to its use in the of assurance of identity of the sender or receiver.

Protected communications include all data encryption methods (for example, hashing and all the algorithms associated with hash, Internet Protocol Security (IPsec), Data Encryption Standard (DES), Triple DES).

Transaction privacy are controls that protect against the loss of private information in respect to transactions, not data-at-rest. Examples of control mechanisms of this type are secure shell, Secure Socket Layer (SSL), and Secure Hypertext Transfer Protocol (SHTTP).

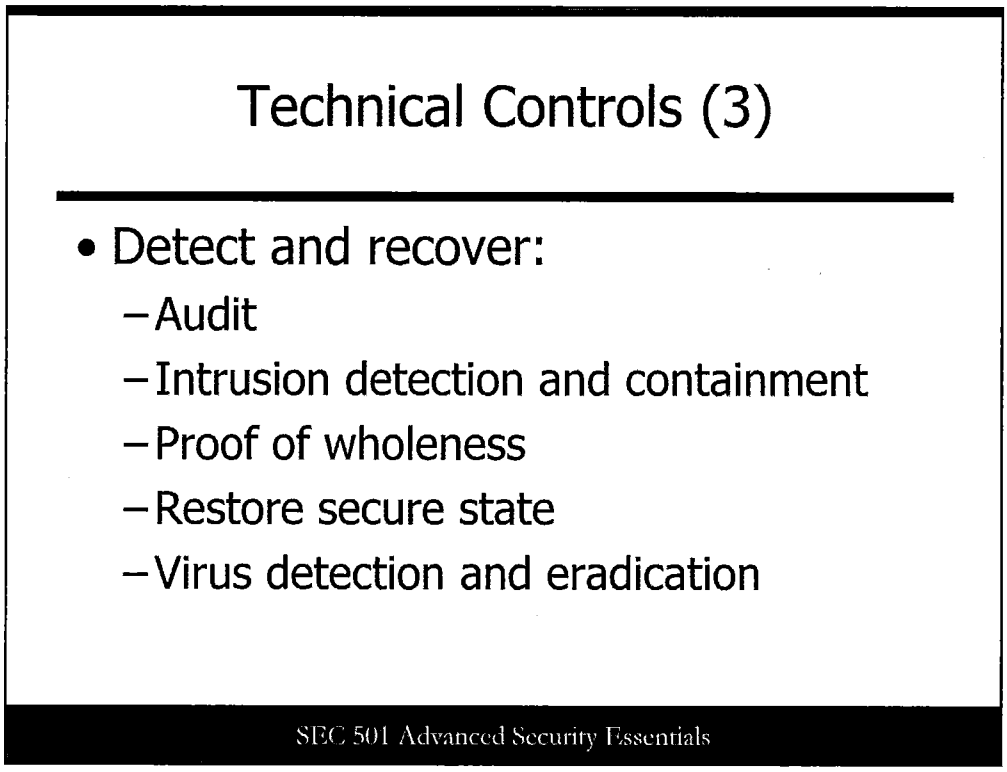

Technical Controls (3)
Detect and recover controls are reactive measures put in place to alert or recover from nefarious or irregular activity or loss. These controls are the assurance controls over the smooth and secure operations we want to maintain.

Audits are used to identify irregular activity in logged (after-the-fact) activity to assist in the identification and level of activity. They help paint the picture of the steps that occurred during an event. Due to their evidentiary nature, audit logs and their mechanisms need to be protected as well.

Intrusion detection and containment: Intrusion Detection Systems (IDS) are reactionary in nature and are usually rule-based or anomaly-based. Upon identifying abnormal operations, these devices/applications begin logging and alert security and/or management of the activity.

Proof of wholeness refers to the integrity or irregularity of traffic or data and sends alerts to individuals and/or logs. This information is later used to determine corrective actions and identify loss. These controls can be extremely important when referring to the integrity of firewalls or other protective devices because weaknesses at those devices affect the security of the whole agency.

Restore secure state: This is like a snapshot in time. If a security breach or event occurs, these control mechanisms enact measures to place the compromised systems back in a secure state. It may be reverting to a previous, uncompromised state or one that restores a clean image and enacts updates/patches and places back online. Some comprehensive restore states require continuous monitoring because it may have been an update that placed the systems or data at risk.

Virus detection and eradication is just that. Active applications or appliances designed to detect and remedy irregular activity that pose a risk of infecting or reducing the integrity or security of a system or data.

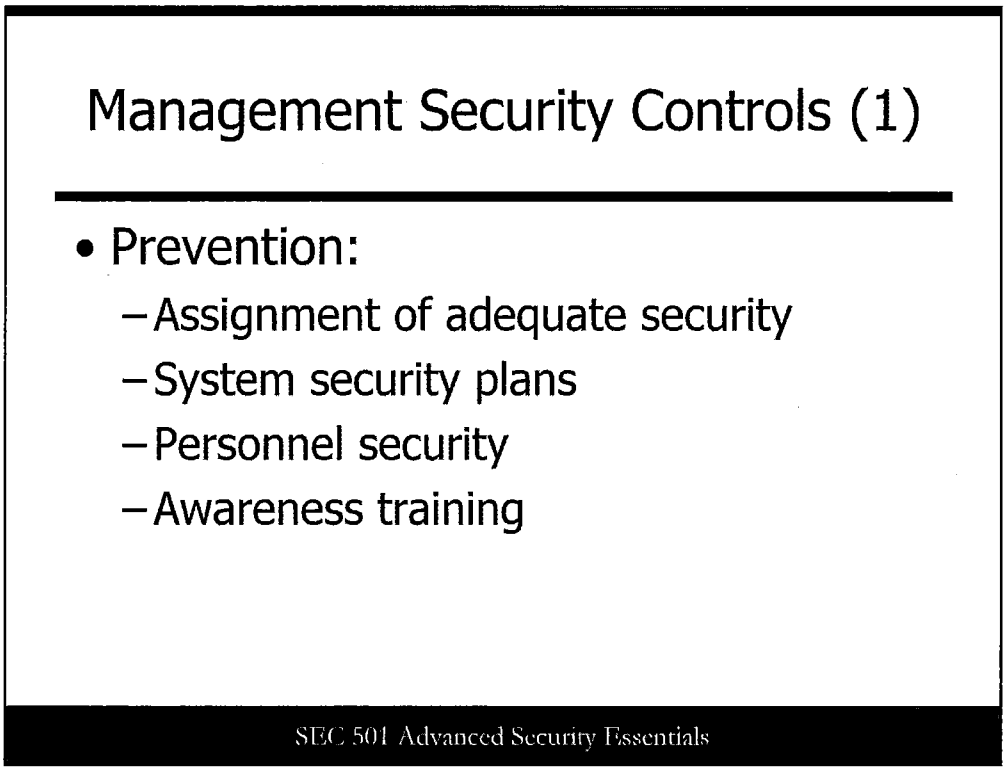

Management Security Controls (1)

Roles and responsibilities are an important security control. Ensuring appropriate management of security controls is an important function, and periodic review is essential to ensure existing procedures and protection are still sufficient to reduce any potential risk.

System security plans are one of the best documentation efforts any organization can develop. They offer identification of assets, especially critical ones, evaluate and assign appropriate security measures to protect such assets, and can be used often in the effectiveness of control measures, replacement requirements, and capability inputs. In addition, they are an invaluable management tool. NIST SP 800-18 "Guide for Developing Security Plans for Information Technology Systems" provides a good guideline.

Personnel security as a management control involves separation of duties, least privilege, and assignment and termination of user accounts.

The cornerstone of security is awareness training. All the protective measures in the world can be used and incorporated, but if the personnel are not made aware of their function, their use, how to identify security issues, how to implement security, and how to report incidents, then all the protective measures in the world cannot prevent loss.

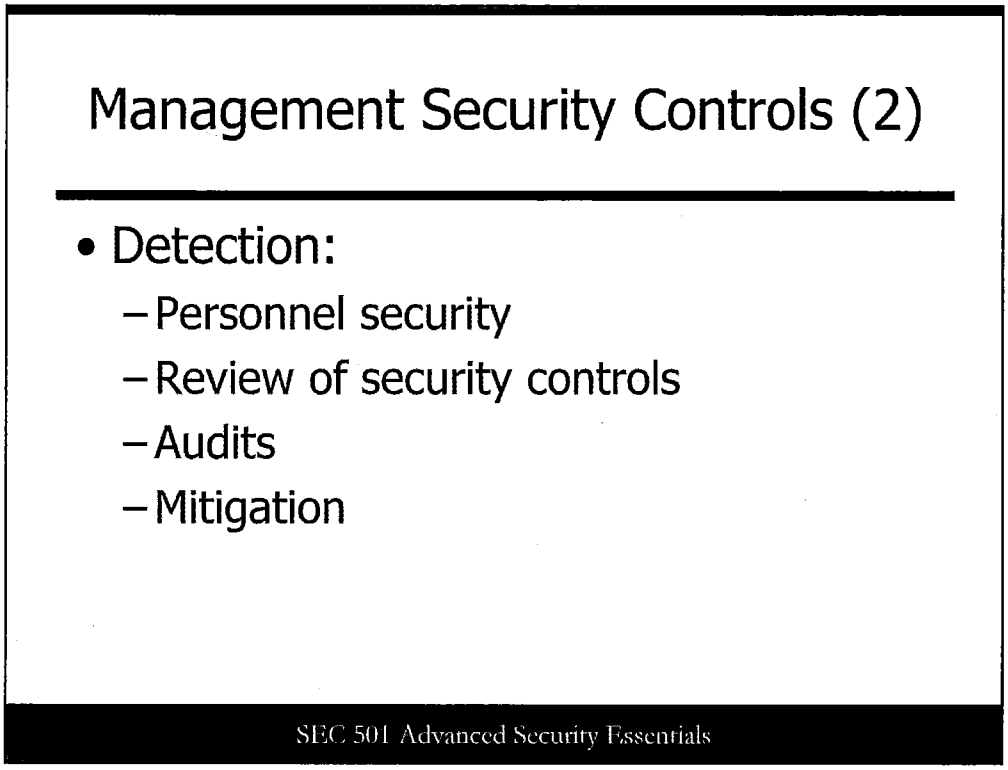

Management Security Controls (2)

Detection controls within management responsibility.

Again we see personnel security. This time, as a detective control, so we are referring to in-depth background investigations, clearances, and rotation of duties.

Review of security controls speaks for itself; security is a living process, and without periodic reviews to ensure detection and prevention, they cannot remain effective.

Audits provide the checks and balances of our deterrent operations. Audits provide the analysis tools in detecting past activity, verification of user access, and assistance in the identification of activity that has occurred.

Mitigation is a managerial control, in that it refers to the identification and action of the management toward the handling of risk. If a threat poses minimal risk, and management determines through analysis or review of existing controls that sufficient or more controls are needed, the appropriate action occurs.

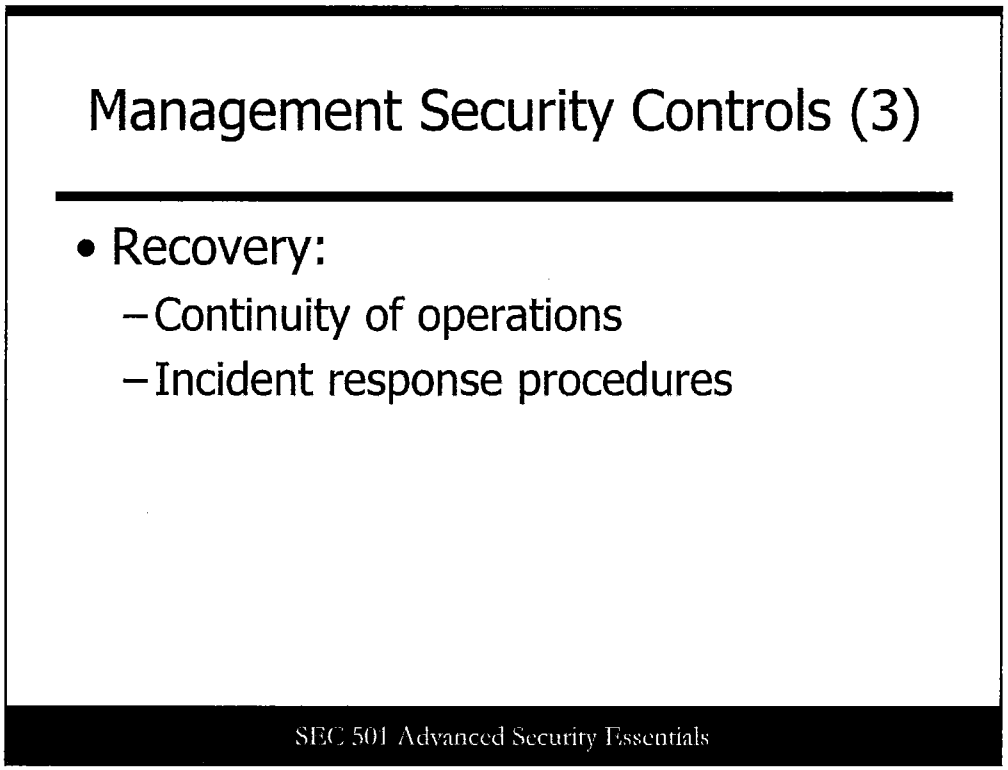

Management Security Controls (3)

Recovery: After an attack, or loss of data, or significant event, operations must continue. How are we to do this? With effective recovery controls, we can ensure critical assets are returned to operation. An effective recovery plan includes the following:

Continuity of operations: Sometimes referred to as a COOP, this is a plan on how to handle a disaster or loss of resources after the fact and is a guide to getting operations back in a readiness state.

Incident response procedures: What steps are taken when an incident occurs, whether it be a virus, a network attack, or a data compromise occurring? NIST SP 800-94, "Guide to Intrusion Detection and Prevention Systems (IDPS)," affords guidance on the inclusion of IDS in your security. Incident response procedures should include steps to identify, report, and respond to attacks against one's devices, network, or data.

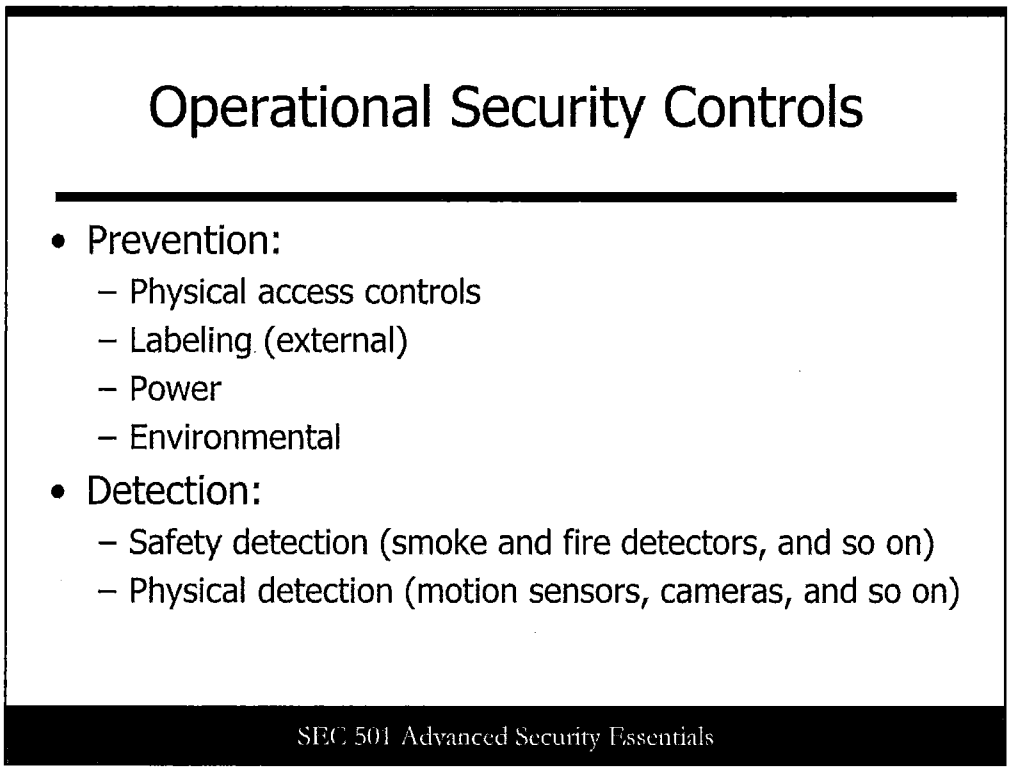

Operational controls are ones we deal with every day and provide the security of operations. These controls cover everything from virus detection and uninterruptable power supplies (UPS) to environmental controls such as HVAC systems.

For prevention controls, we include:

- Security gates, guards, and dogs
- Heating, ventilation, and air conditioning (HVAC)
- Fire suppressant
- Labeling of assets (classification and responsible agents)
- Off-site storage (recovery)
- Safes and locks

For detection controls:

- Smoke and fire sensors
- Motion detectors
- Heat/thermal detectors
- Vibration alarms
- Voltage regulators

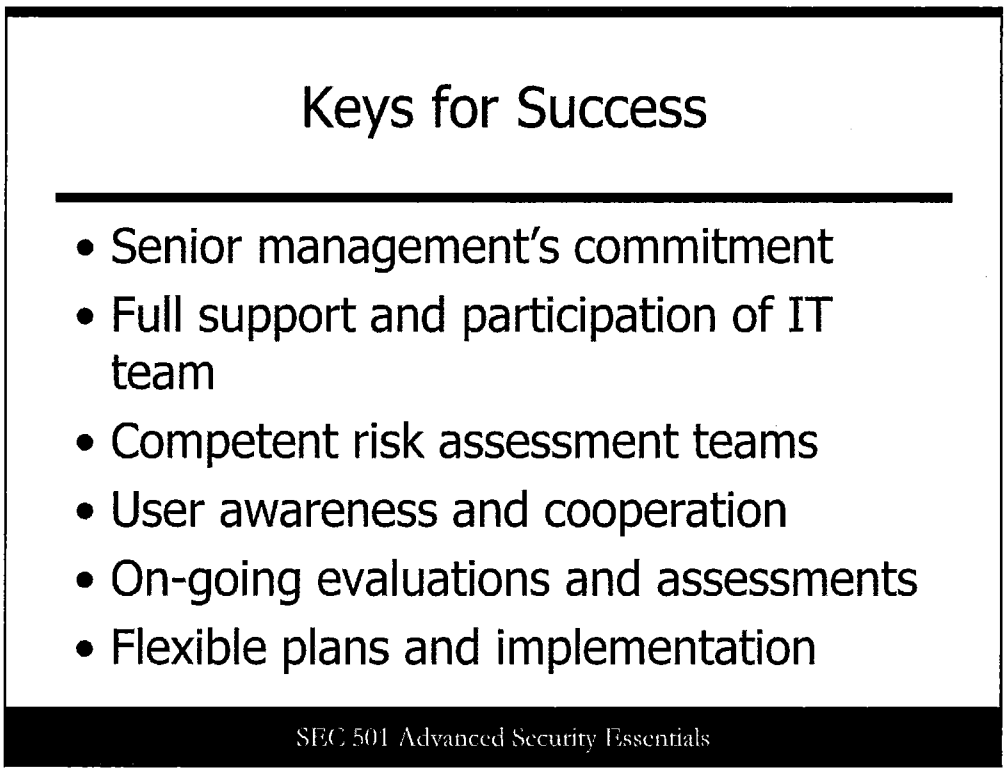

Identifying risks is important, but it is critical that risk has buy-in from the executive team. Risk management cannot just be an area that IT Security focuses in on; there needs to be a corporate-wide program with buy-in from all business units. Coordinating across the entire organization will lead to a successful security program.

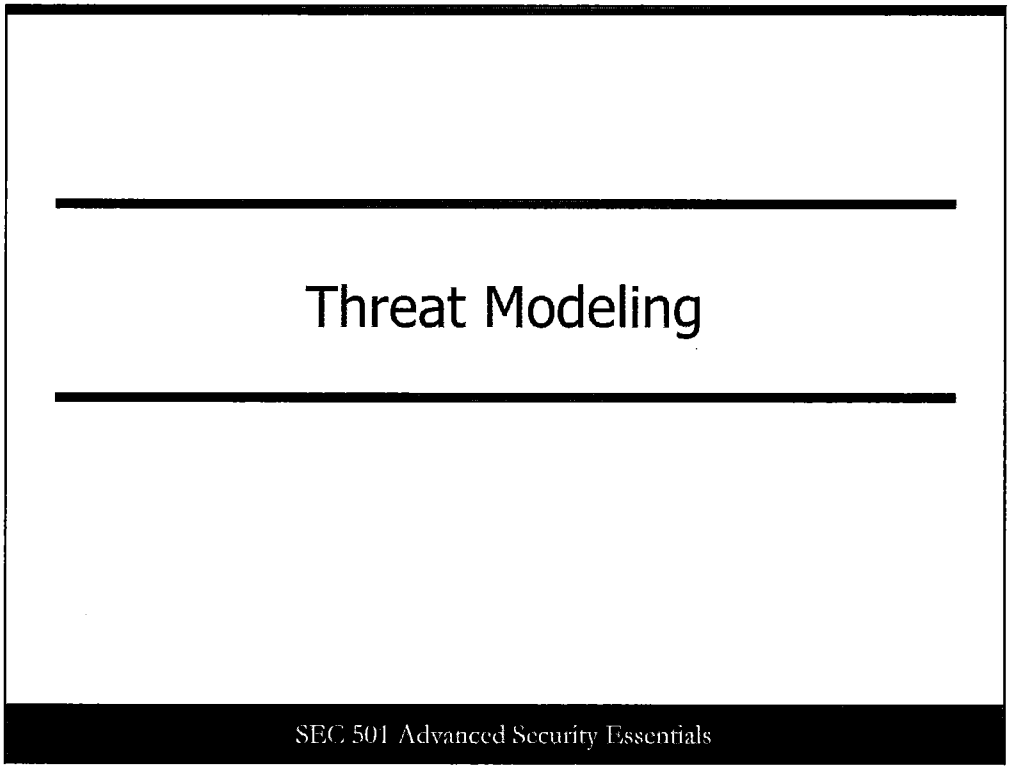

This page intentionally left blank.

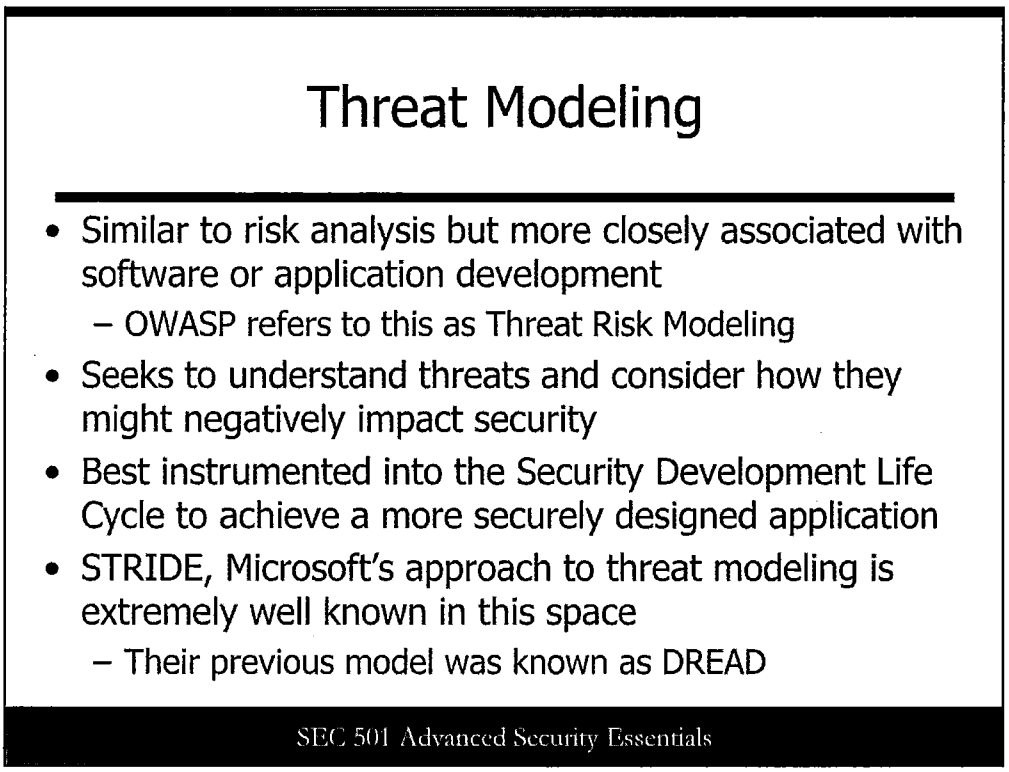

Threat Modeling

The concept of threat modeling is quite similar to that of risk analysis. However, threat modeling is much more closely associated with software or application development. OWASP actually even refers to threat modeling by the name Threat Risk Modeling.[1]

The organization most well-known for incorporating threat modeling is Microsoft. Its current approach to threat modeling is known as STRIDE.[2] Their previous model was called DREAD.

[1] https://www.owasp.org/index.php/Threat_Risk_Modeling (http://cyber.gd/414_115)
[2] https://msdn.microsoft.com/en-us/library/ee823878%28v=cs.20%29.aspx (http://cyber.gd/414_116)

Threat Identification
The goal of threat identification is to appreciate the applicable threat sources, understand their motivation, and also determine their capabilities.

The lack of reliable data that details the threat landscape makes this a rather challenging portion of the risk analysis process. Most organizations outside of the government and intelligence sector do a fairly poor job (or don't attempt at all) in the threat identification portion of risk analysis.

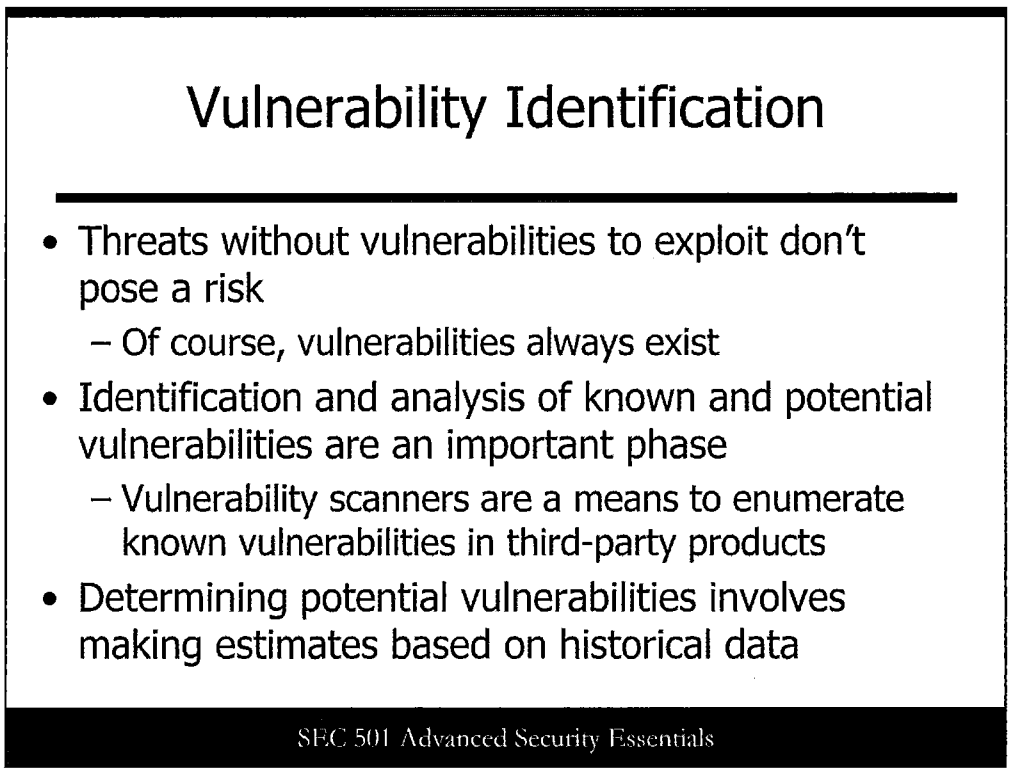

Vulnerability Identification

Even if it is accepted that there are motivated and capable threat sources that pose an active threat, without associated vulnerabilities those threats can exploit, this is no risk. Unfortunately, vulnerabilities always exist.

The identification and analysis of vulnerabilities in common third-party products are fairly straightforward and largely commoditized. Numerous vulnerability scanners can all do a passable job at vulnerability identification in enterprises.

The more significant challenges in vulnerability identification come when dealing with custom-developed software, and web and mobile applications. Identifying these vulnerabilities is significantly more difficult and more likely to result in both false positives and false negatives than typical network vulnerability scanners. Also rather cumbersome, to say the least, is assessing potential unknown vulnerabilities.

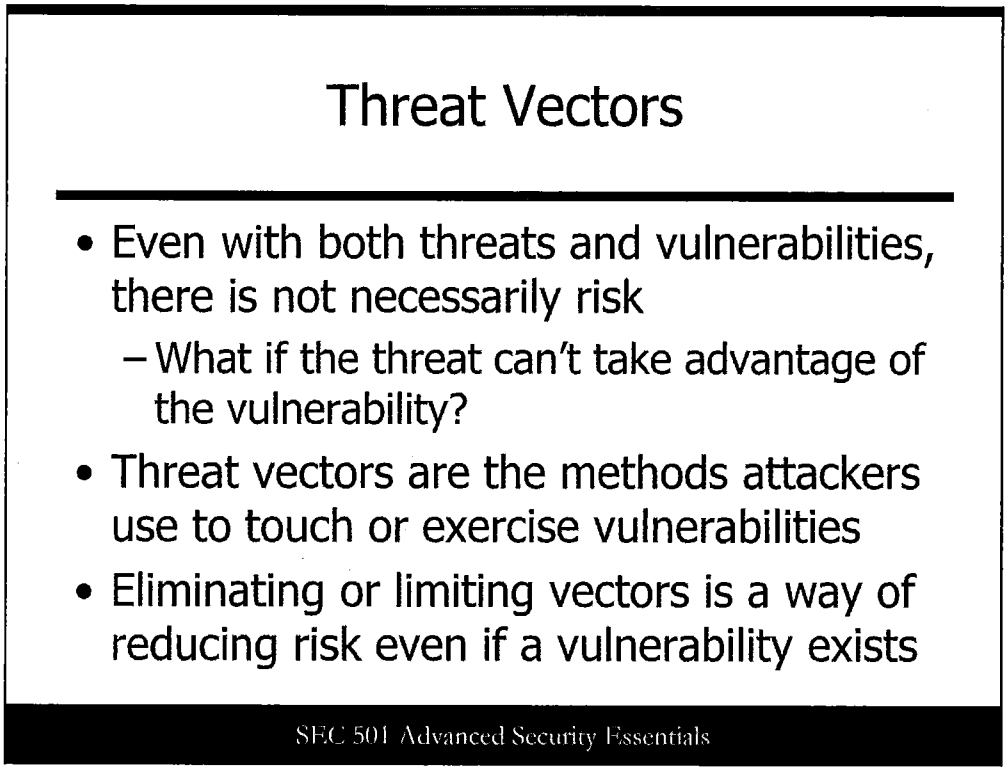

Threat Vectors

Another consideration that changes how we appreciate the system risk is the concept of a vector, or threat vector. Even with capable motivated threat sources that target a vulnerability that exists, the risk might be negligible or nonexistent for that particular threat statement. Wait; we can have both a capable threat and an exploitable vulnerability, but no risk?

Yes, because the mere presence of a threat and vulnerability does not mean that there is a way that the threat can exploit, or take advantage of, the vulnerability. There must be a means for the threat to exercise the vulnerability for there to be risk. This is the concept of a vector, or threat vector.

Imagine an internal Windows NT system not even at the latest Service Pack. (Although this might sound implausible, NT boxes still have not all been decommissioned.) It would be hard not to exploit this system; you look at it sternly and it is likely to blue screen. Certainly there are attackers that could exploit this system and would be motivated to do so, but unless they can actually get their exploits to the system, there is no real risk.

The elimination or limitation of vectors is an additional means of reducing risk, even without changing the threats or vulnerabilities themselves.

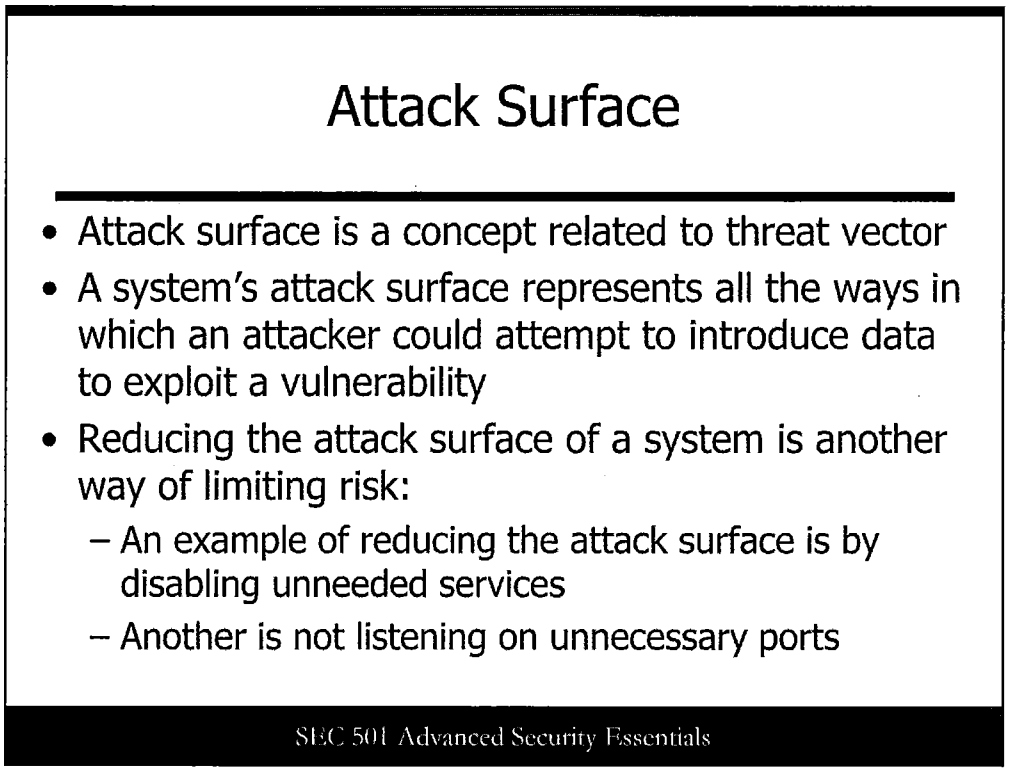

Attack Surface

A risk concept closely related to that of threat vector is the concept of the attack surface. A system's attack surface refers to all the various ways in which an attacker could attempt to introduce data with the goal of exploiting a vulnerability.

Reduction of a system's attack surface is an additional way of reducing risk. Importantly, reducing the attack surface is one of the means of reducing risk associated with unknown vulnerabilities.

For example, imagine that a Windows workstation is running the SSDP Service (because that is the default and few people know what SSDP actually does). Assuming this service is unneeded, then disabling this service would reduce the attack surface. The attacker cannot target the system via this service. Also, even if there is a vulnerability in SSDP, then this system would not be vulnerable, even if unpatched, because of the attack surface reduction.

Security Configuration Management or hardening are means to reduce the attack surface of systems by ensuring that only the necessary features are enabled on systems.

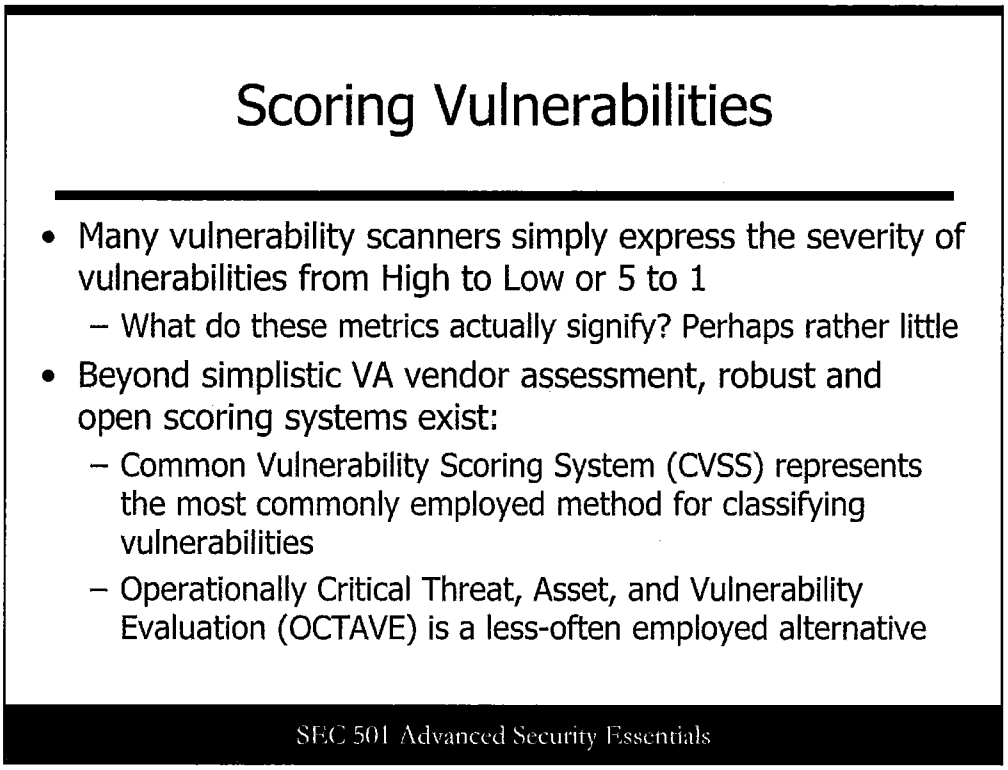

Scoring Vulnerabilities
All vulnerabilities are not created equal. Organizations require a consistent means of evaluating the thousands of vulnerabilities announced by vendors. Vulnerability scanners will do their best to help guide the prioritization, but often leave something to be desired.

Beyond the scanning vendor's scoring, robust open vulnerability scoring systems exist. Although there are others, such as Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE), version 2 of the Common Vulnerability Scoring System (CVSSv2) is by far the most commonly referenced.

Many vendors actually provide their assessment of the CVSS score for their announced vulnerabilities.

Additional information on OCTAVE is available at http://www.cert.org/octave/ (http://cyber.gd/414_117).

Vulnerability Reporting

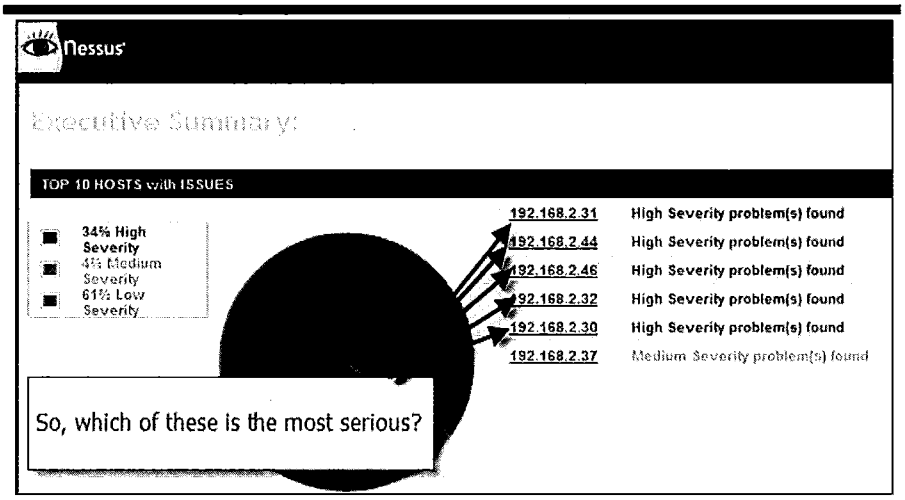

This page intentionally left blank.

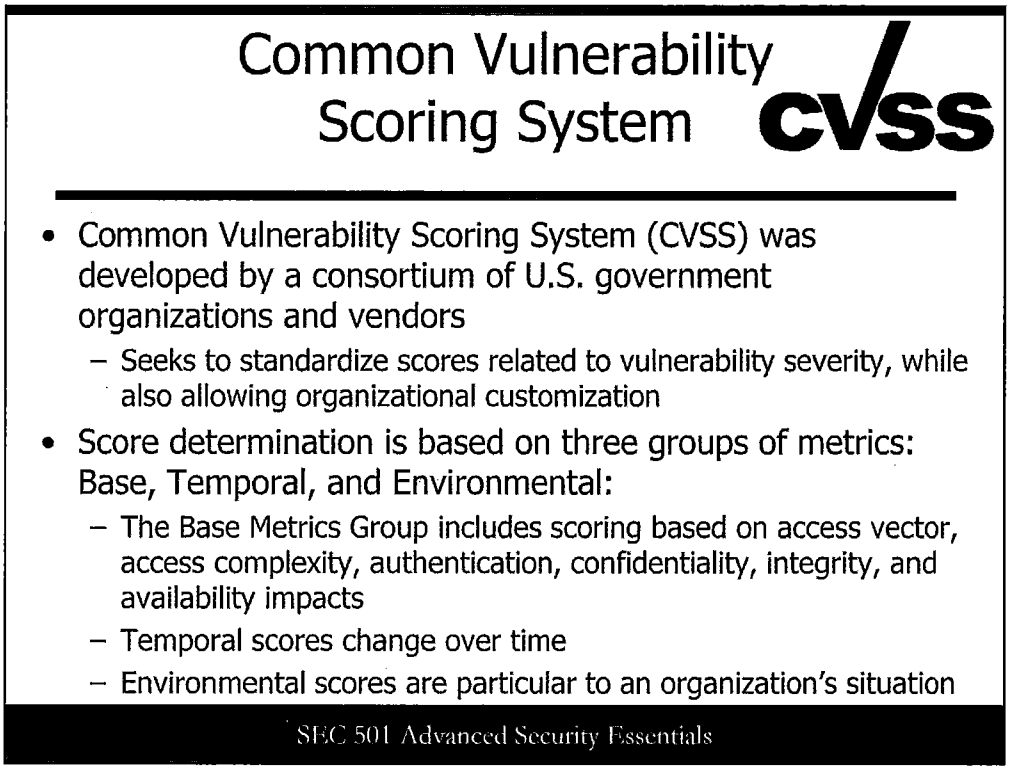

Common Vulnerability Scoring System
Numerous organizations have been involved in both the initial development of CVSS as well as versions. The U.S. government was involved in some capacity as well as major software, hardware, security, and vulnerability scanning vendors.

The goal of CVSS is to provide an open standard method for comparing the relative severity of vulnerabilities. CVSS accomplishes this through the use of three groups of metrics: Base, Temporal, and Environmental.

The Base Metrics are the standard scores and are required for a CVSS score, whereas the Temporal and Environmental scores allow for additional precision. The Base Metrics Group includes scoring based on Access Vector; Access Complexity; Authentication; Confidentiality, Integrity, and Availability Impacts. Temporal Scores are those that will change over time, for example, the availability of both patches and exploit code. The Environmental scores are those particular to an organization's situation.

Additional information on CVSS is available at http://www.first.org/cvss (http://cyber.gd/414_118).

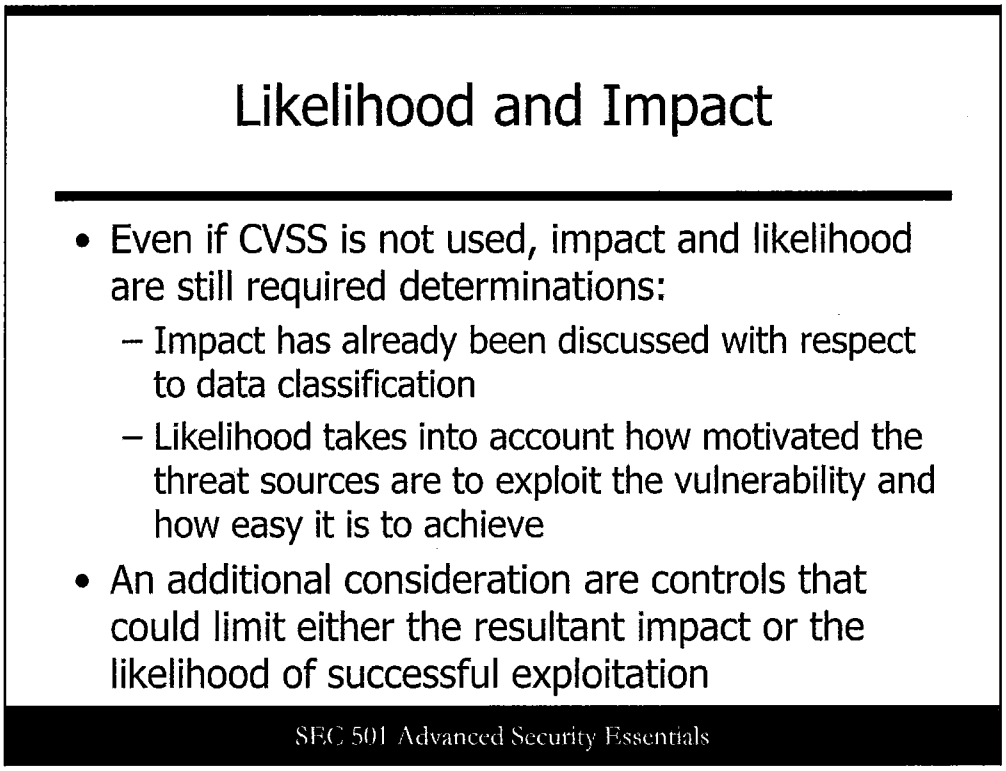

Likelihood and Impact

Although CVSS is a robust approach that allows for organizational customization, regardless of whether a formal scoring system like CVSS is used, impact and likelihood are key metrics to determining risk. Impact has already been discussed with respect to potential loss of CIA.

Likelihood is an attempt to determine whether successful exploitation is likely. The motivation and capability of threat sources will affect the likelihood. The availability of both patch and exploit code, which indirectly impacts the threat's capabilities, will also affect the likelihood. Other items like the vector needed, level of access required, and whether interaction is necessary for exploitation would also affect the likelihood metric.

Finally, security controls or countermeasures currently deployed can have a nullifying effect that limits impact and/or likelihood associated with exploitation.

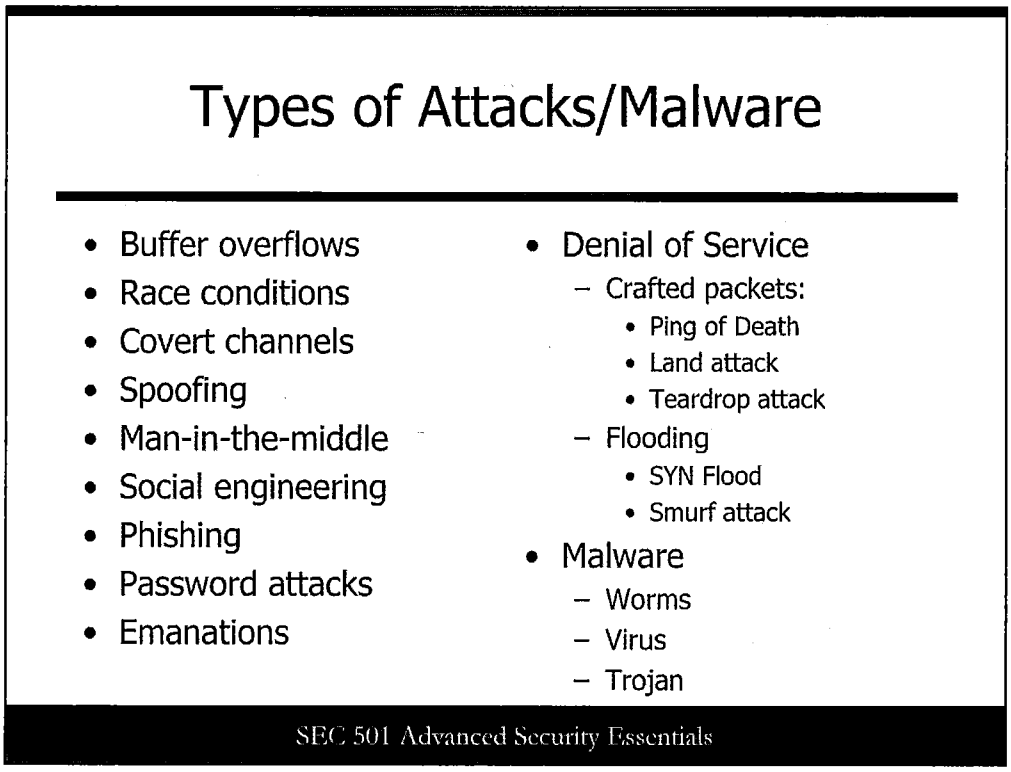

Types of Attacks/Malware

- Buffer overflows
- Race conditions
- Covert channels
- Spoofing
- Man-in-the-middle
- Social engineering
- Phishing
- Password attacks
- Emanations
- Denial of service
  - Crafted packets:
    - Ping of Death
    - Land attack
    - Teardrop attack
  - Flooding
    - SYN Flood
    - Smurf attack
- Malware
  - Worms
  - Virus
  - Trojan

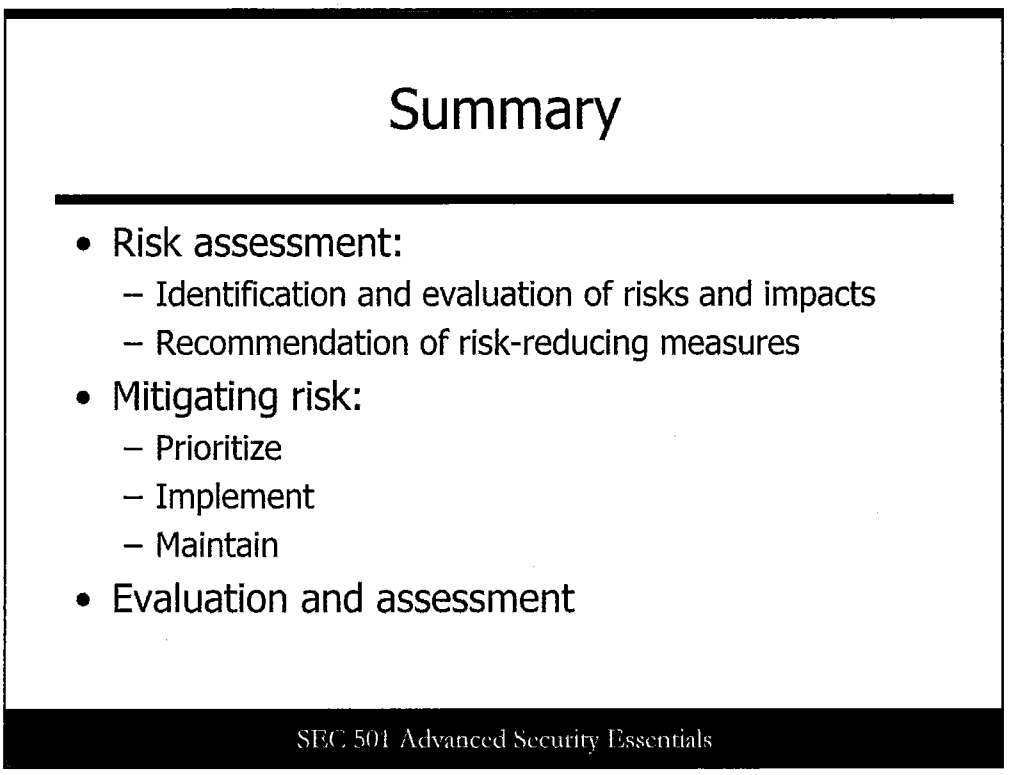

Risk management encompasses three processes: risk assessment, risk mitigation, and evaluation and assessment.

The risk assessment process includes: identification and evaluation of risks and risk impacts and recommendation of risk-reducing measures.

Risk mitigation refers to prioritizing, implementing, and maintaining the appropriate risk-reducing measures recommended from the risk assessment process.

Evaluation and assessment involves putting a risk management program to the test. Review it, test it, practice it, evaluate it, and publish it.

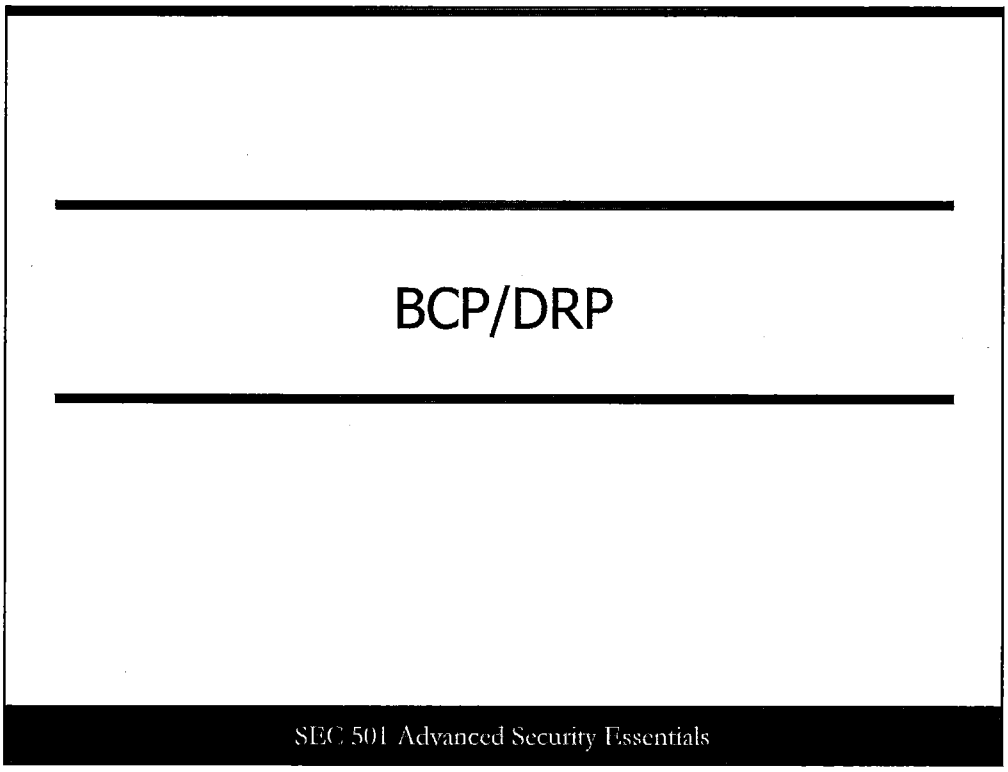

This page intentionally left blank.

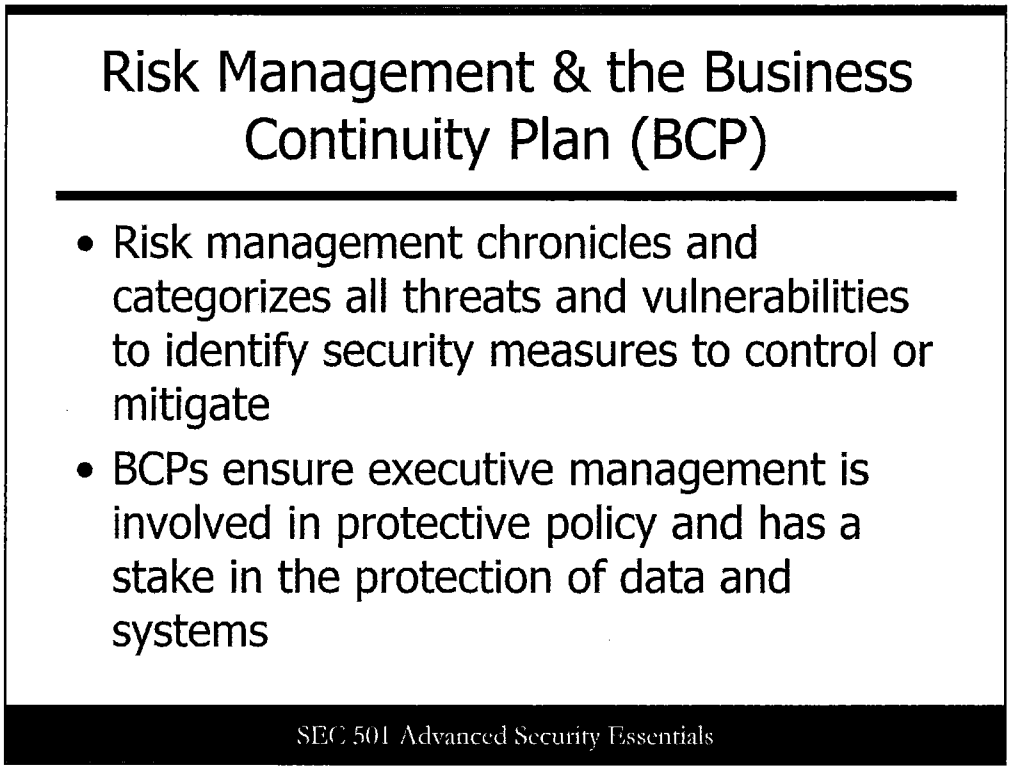

The inclusion of a thorough risk management process into a Business Continuity Plan (BCP) and also into a Disaster Recovery Plan (DRP) is a good practice. The risk management documents and supports most actions contained with a BCP and DRP.

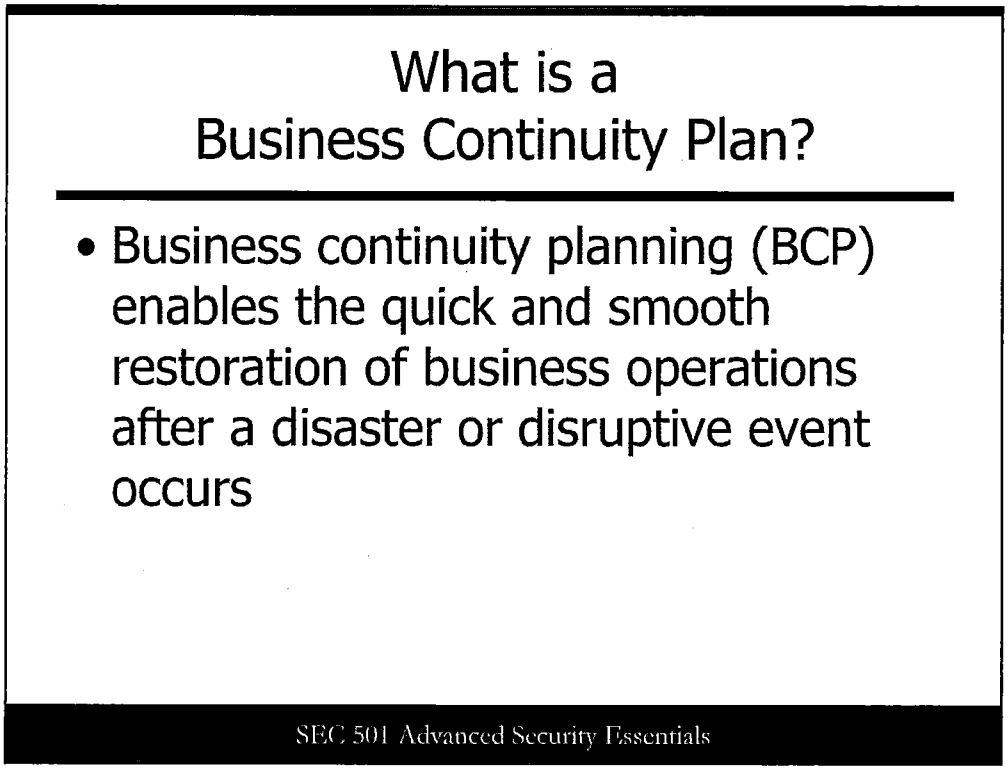

What is a Business Continuity Plan?

Contingency Planning Within Your Policy
A critical aspect of security policy for your organization is planning for contingencies.

Overview of Contingency Planning
First, we define what a Business Continuity Plan (BCP) and Disaster Recover Plan (DRP) are and explain why an organization needs them. Subsequently, we dive into the process for developing them.

What Is a Business Continuity Plan (BCP)?
According to the National Computer Security Center, a Business Continuity Plan (BCP) is, "A plan for emergency response, backup operations, and post-disaster recovery maintained by an activity as a part of its security program that will ensure the availability of critical resources and facilitate the continuity of operations in an emergency situation."

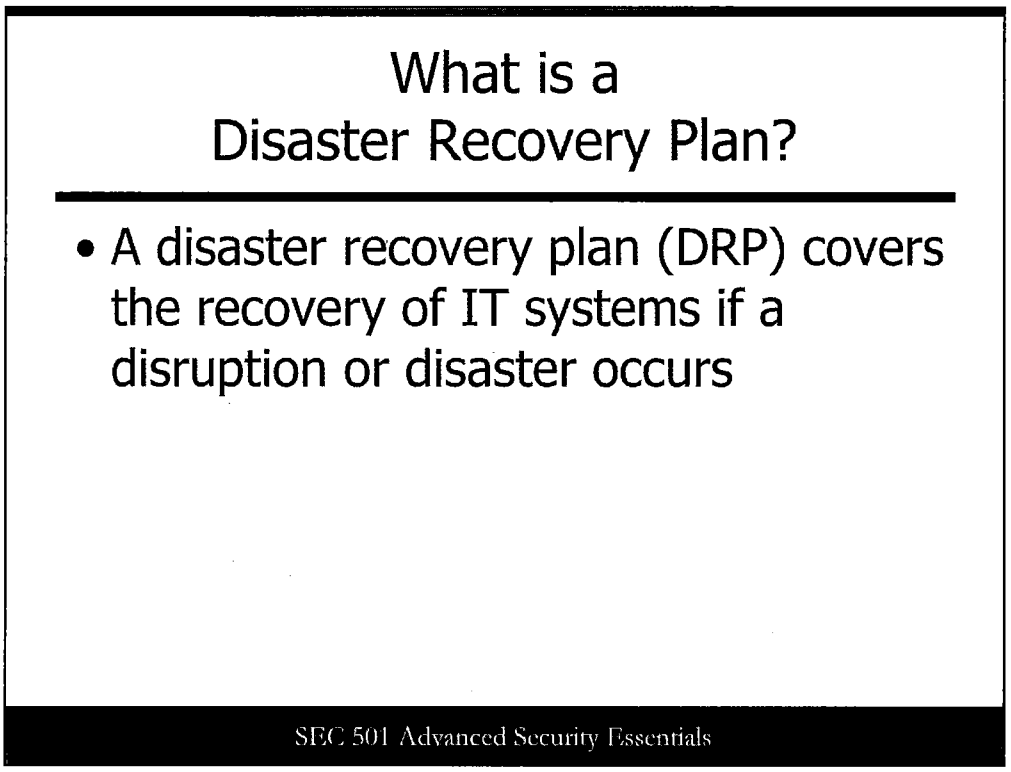

What is a Disaster Recovery Plan?
A disaster recovery plan (DRP) covers the recovery of IT systems if a disruption or disaster occurs. It provides the capability to process essential organizational applications, even if they are not operating at 100% efficiency, and the ability to return to normal operations within a reasonable amount of time.

Although the terms BCP and DRP are sometimes used interchangeably, business continuity planning and disaster recovery planning are two distinct plans that tackle different areas of the recovery process. Business continuity planning deals with the restoration of the business processes or the continued operation of business processes. Organizational processes can continue without computers. For example, checks can be written by hand. With a continuity plan, the company can reduce the impact a disaster has on the normal business operation. The disaster recovery plan covers the restoration of critical information systems that support the business processes.

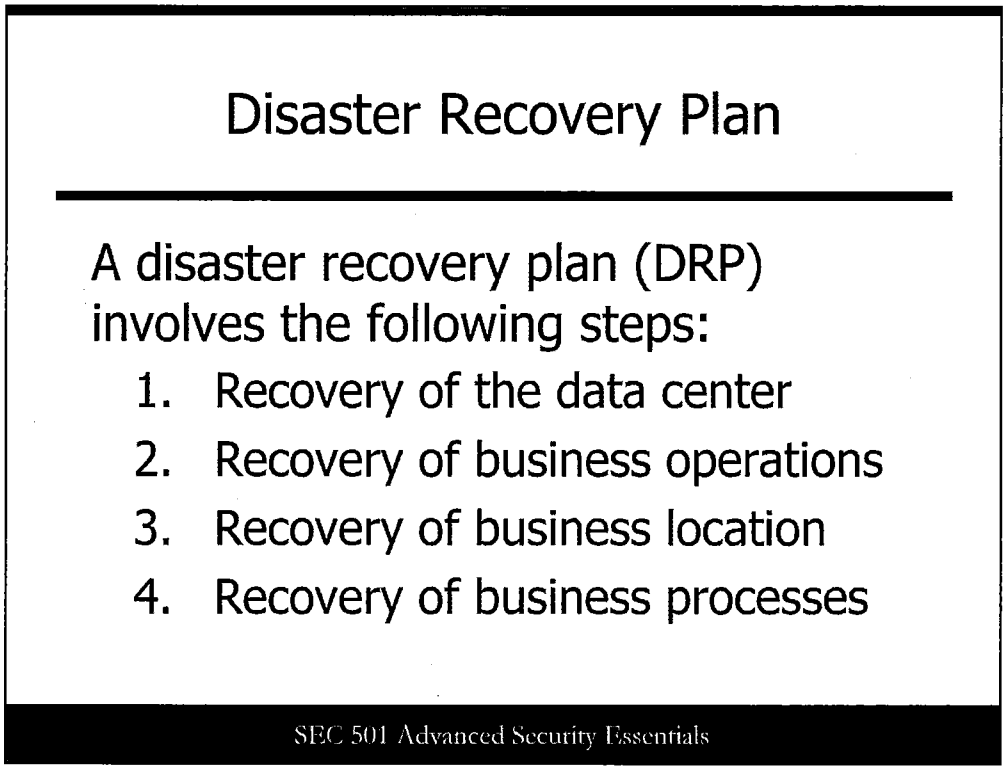

Disaster Recovery Plan
Disaster recovery planning involves the following steps:

1. The recovery of the data center: Because the DRP relates to the restoration of the information systems, it should address bringing the data center, one of the critical areas, back online.
2. The recovery of business operations: This is sometimes referred to as user contingency planning. If a critical computer system is down, this part of the DRP deals with the alternative methods of continuing with the business operations. For instance, if your main payroll system were inoperable, a contingency plan could be to issue the payroll checks manually.
3. The recovery of the business location: Part of the business resumption plan, this section deals with the steps required to recover the actual physical business location. Often a disaster is partial, and recovery of the premises might consist first of patching together what is left, followed by backfilling what has been lost.
4. The recovery of business processes: Also part of the business resumption plan, this section handles the recovery of all the various business processes so the company can resume normal business operations. This is the paramount step. The entire purpose of the plan is not about computers, networks, and data, but about the timely continuity and restoration of business processes.

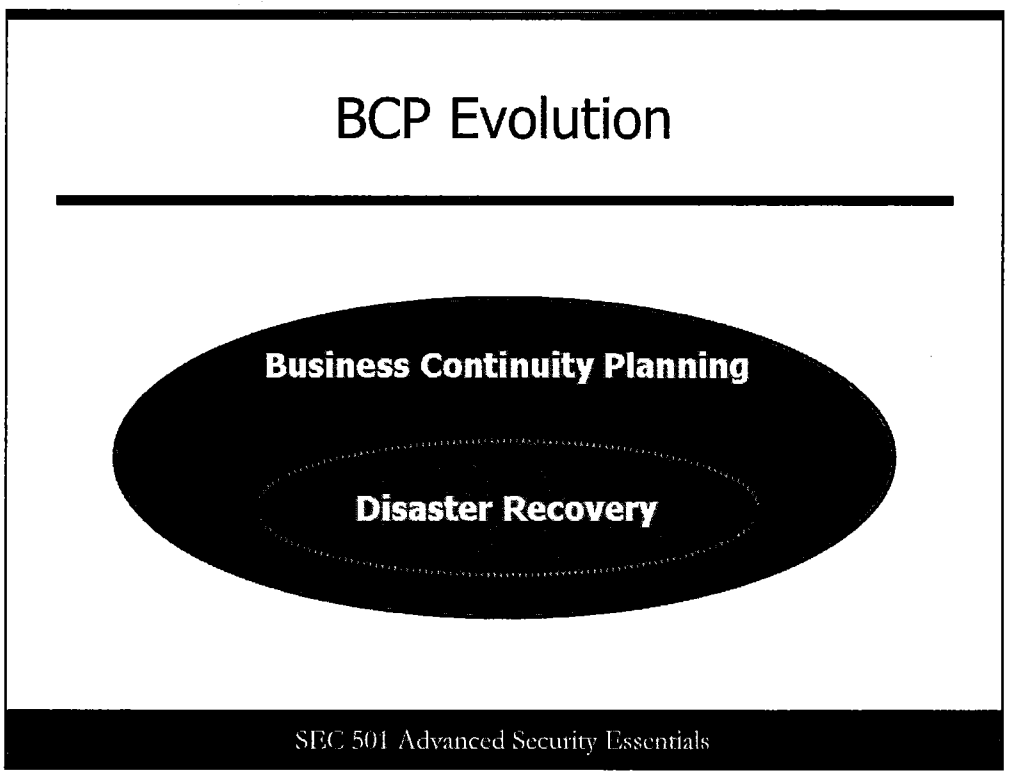

BCP Evolution
In many instances, the terms disaster recovery planning and business continuity planning are used synonymously or, at least, are not clearly differentiated, thereby causing confusion among individuals introduced to this field for the first time.

Disaster recovery planning is an integral part of business continuity planning, but it does not encompass the entire discipline. The complexity of business operations forced disaster recovery planning to evolve into business continuity planning as planners recognized that more was required to ensure business survival than could be encompassed by disaster recovery planning alone.

An easy distinction between disaster recovery planning and business continuity planning is:

- Disaster recovery is short-term focused.
- Business continuity is long-term focused.

This diagram illustrates the temporal relationship between disaster recovery planning and business continuity planning.

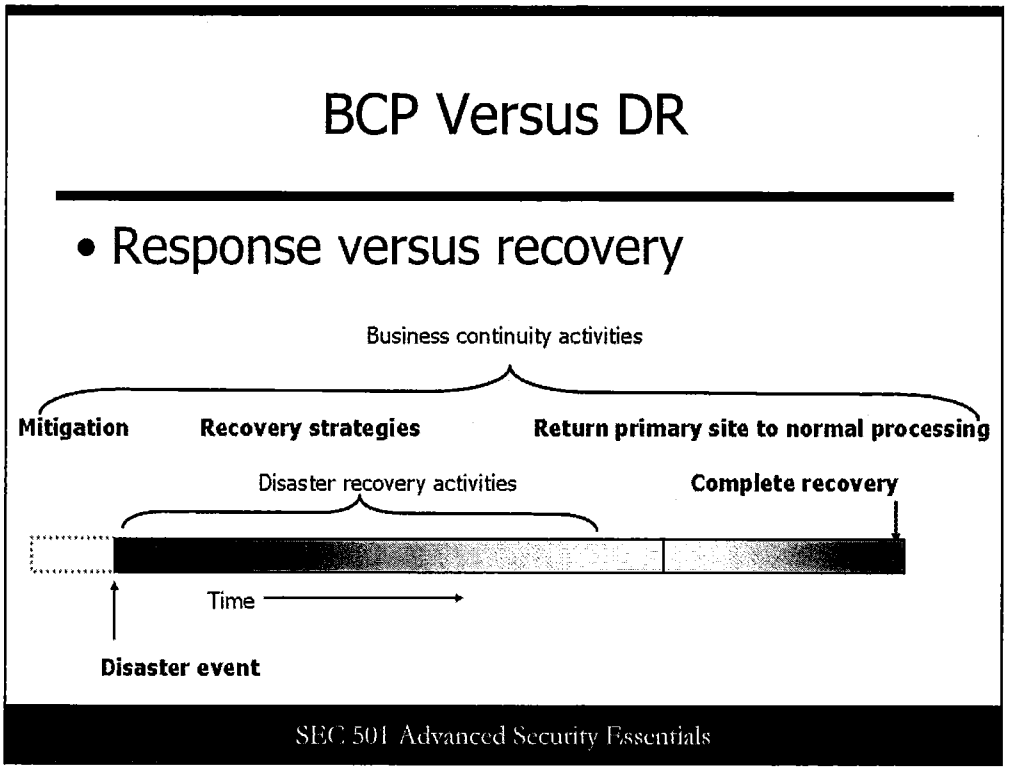

BCP Versus DR

Disaster recovery provides a response to disruption, whereas business continuity planning implements the recovery. The preceding figure shows that the disaster recovery activities have a short time span, but business continuity activities are much more pervasive and long-lasting.

The goal of BCP/DRP is to make the response time to a disruption as short as possible and the time required for completely recovery.

During disaster recovery activities, that is, when a disaster strikes an organization, almost all normal business activities are heavily modified, reduced, or completely suspended. Only critical business processes resume and usually at an alternative site.

As repairs are completed, normal business activities resume as the business continuity plan dictates. Recovery is complete after all normal business processes return to "business as usual."

Business continuity activities form an umbrella over a crisis situation, while disaster recovery activities are a *subset* of business continuity activities.

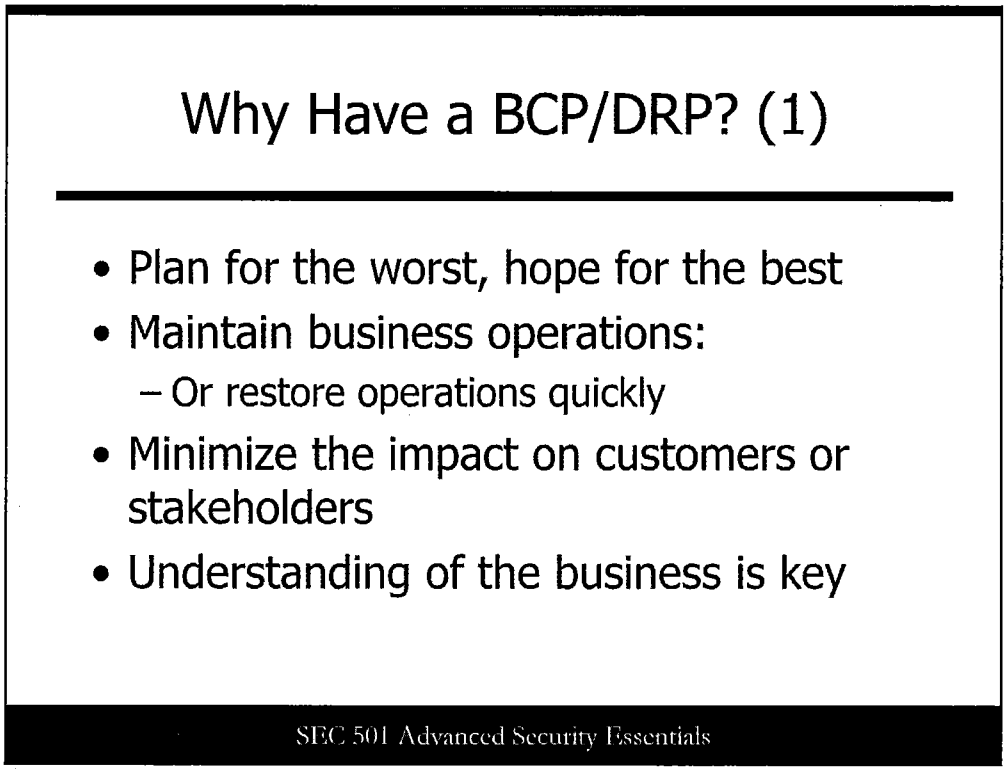

Why Have a BCP/DRP? (1)
Continuity planning might be likened to insurance: It is an expense consciously made to significantly reduce the impact of something bad occurring. You hope that nothing bad occurs. Even if it does not, the insurance premium and the expense of continuity planning were not wasted because they purchased certain assurances as a key component of the organization's risk management. As the slide says, "Plan for the worst; hope for the best."

Going through the process of identifying your customers (both internal and external) helps to define which business processes are the most important to your organization's survival. You must know who your customers are to properly prioritize business process restoration. For example, if a hospital's surgical center had a backup generator for the floor but the rest of the hospital was without power, many patients in intensive care could die as a result. If air conditioning were unavailable, more people could die from the heat even if they were not originally at risk (as recently happened in Europe). Power outages would prevent cafeteria workers from fixing patient meals, and staff would not even be able to use most of the vending machines. In addition, if power to the ambulance dispatch center were also impacted, requests for assistance could be affected; external customers might also suffer. This is an unlikely scenario, but it gets the point across: Look at all the stakeholders, not just the most visible.

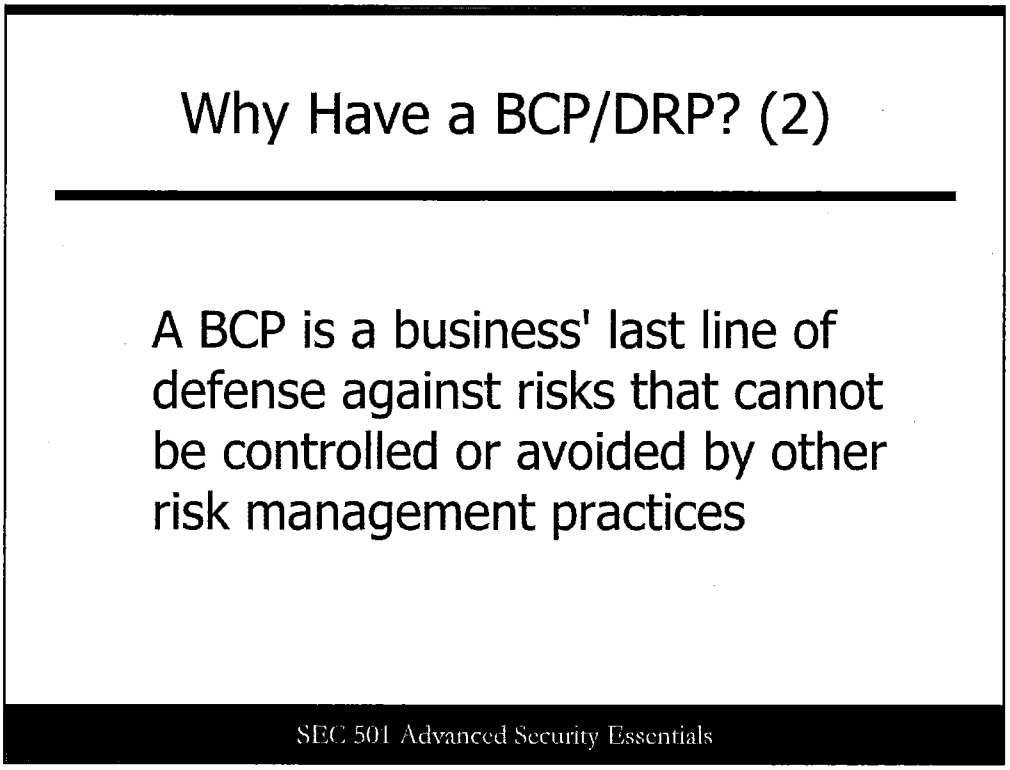

Why Have a BCP/DRP? (2)
A BCP is a business' last line of defense against risks that cannot be controlled or avoided by other risk management practices. Business continuity and disaster recovery planning are not the places for the attitude "ignorance is bliss."

Organizations are often so reliant on key resources, such as technology, that they cannot operate without them. The most key resources of all are human; that is, people who operate the organization's processes, including its recovery processes. The most important aspect of the BCP is to protect the lives of your employees. Yes, the name of the document is called a Business Continuity Plan. In essence, however, unless you operate a company without people, employees are your most valuable asset. Consider the situation of the bond trading company Cantor Fitzgerald. During the September 11, 2001 terrorist attacks on the World Trade Center in New York, Cantor Fitzgerald lost more than 700 of its 1,000 employees. Even if it had been able to get their computer systems up and running, it would not have had the personnel to continue business.

Business Continuity Plan Phases

Project Initiation Phase: Preplanning phase, estimate scope, gain management support, form implementation teams, obtain resource requirements.

Current State Assessments: Determine current state of operation, may include pen testing, Business impact assessments, and benchmarking.

Design and Development Phase: Bases of baseline develop recommendations and action plans regarding next steps of recovery, BCP and DRP crisis management plan. Begin recover resource acquisition process.

Implementation Phase: Work with process owners to develop a clear short-term and long-term plan. Develop testing and training strategies, and the enterprise crisis management plan.

Management Phase: Day to day management of the continuity plan, oversight, and roles and responsibilities.

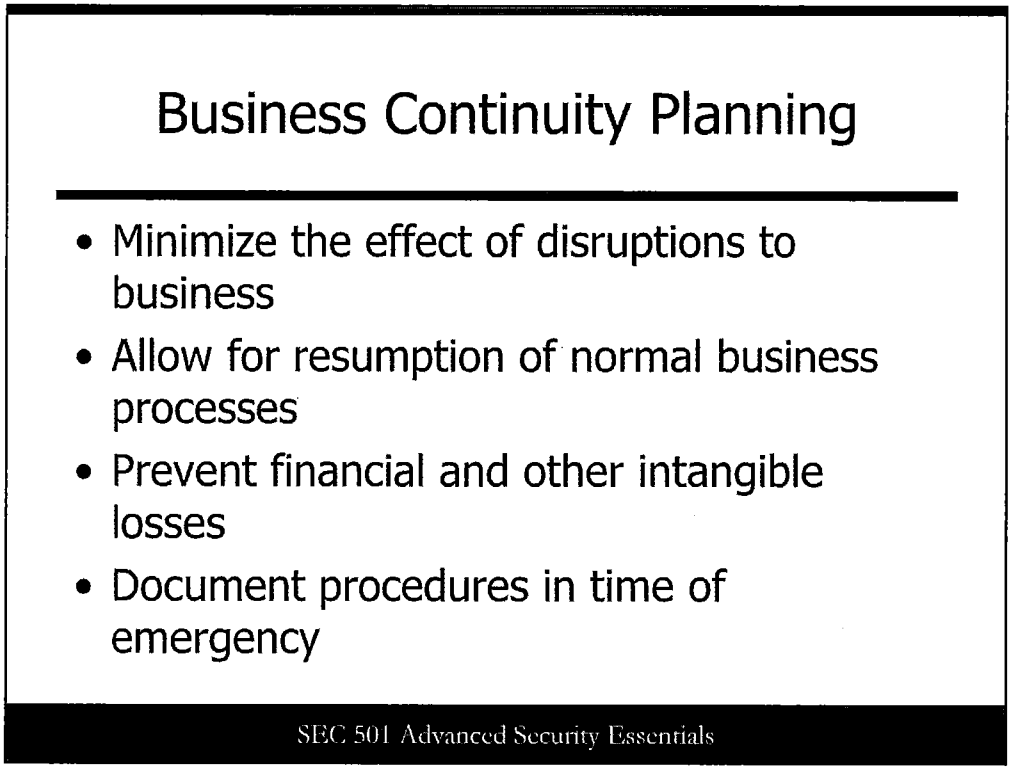

Business Continuity Planning

Business continuity planning encompasses the processes and actions necessary to ensure that the most essential business can continue to operate if unforeseen events occur.

First, to minimize the effects of disruptions, we must first identify the following:

- Events/circumstances/actions that can cause a disruption (that is, threats)
- Cost-effective means of preventing disruptions and their impact (mitigation strategies)
- Ways of transferring or reducing risk (such as buying insurance)
- Recovery strategies that minimize downtime and loss of productivity

Why is continuity planning necessary? To prevent loss to the business of profits, assets, market share, perception of clients, and so on. Regarding loss, financial losses are often the predominate concern in the commercial sector. However, the loss of life, property, or national security might be of greatest concern in some fields, such as health care, and in many government agencies. Financial loss is usually the easiest to justify to management. Less tangible losses such as loss of good will are more difficult to quantify, but are just as important.

After plans have been developed to ensure continuity of operations, these plans must be documented and distributed to the proper personnel. This avoids confusion, ensures that all participants know what roles they play and what responsibilities they have, and provides an efficient way to quickly return to "business as usual."

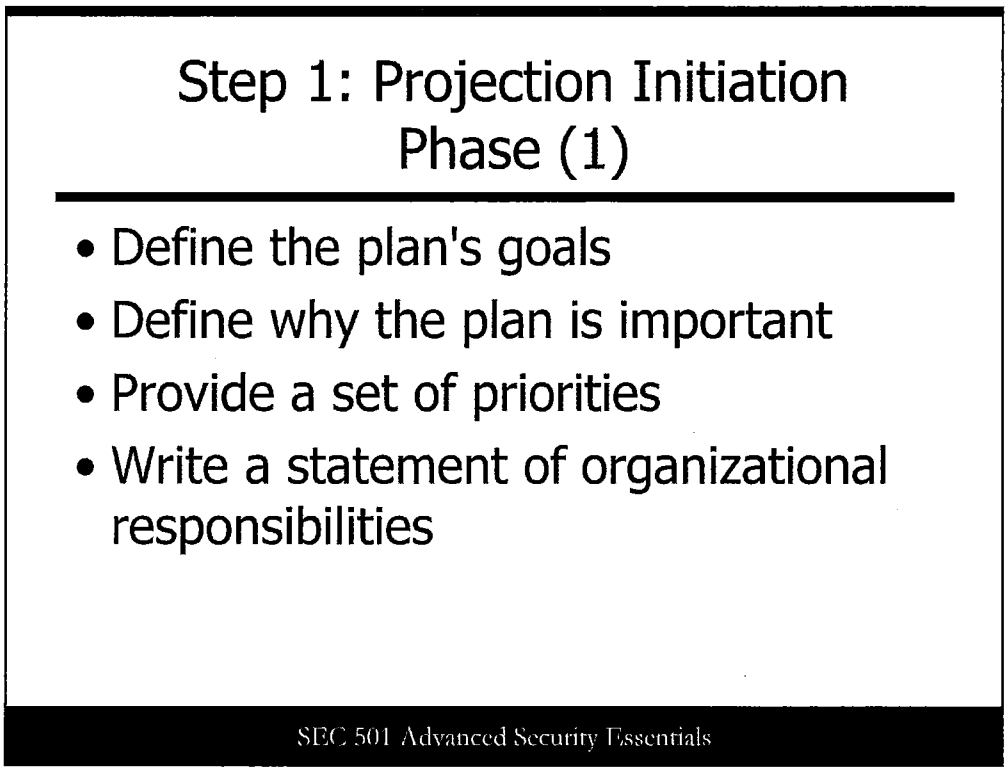

Basic Elements of Continuity Planning (1)
Project Initiation
The key components of a Business Continuity Plan follows:

- **Assess**: Identify and triage all threats (BIA).
- **Evaluate**: Assess the likelihood and impact of each threat.
- **Prepare**: Plan for contingent operations.
- **Mitigate**: Identify actions that might eliminate risks in advance.
- **Respond**: Take actions that are necessary to minimize the impact of risks that materialize.
- **Recover**: Return to normal as soon as possible.

The Continuity Plan should define the goals of the plan and why the plan is important. Writing a plan without understanding the goals is almost as bad as not having a plan at all. The plan should provide a clearly defined set of priorities and a statement of urgency based upon the Business Impact Analysis (BIA). While developing the plan, you should ask the following questions:

- Why is a business continuity plan important to my organization? That is, what would be the impact if we had a disaster and no plan for how to handle it?
- What are the most important business functions (and/or IT systems) that must be recovered quickly before you lose the ability to continue business operations?
- Is the plan's goal simply to further prevent damage or to actually stop and fix the problem?

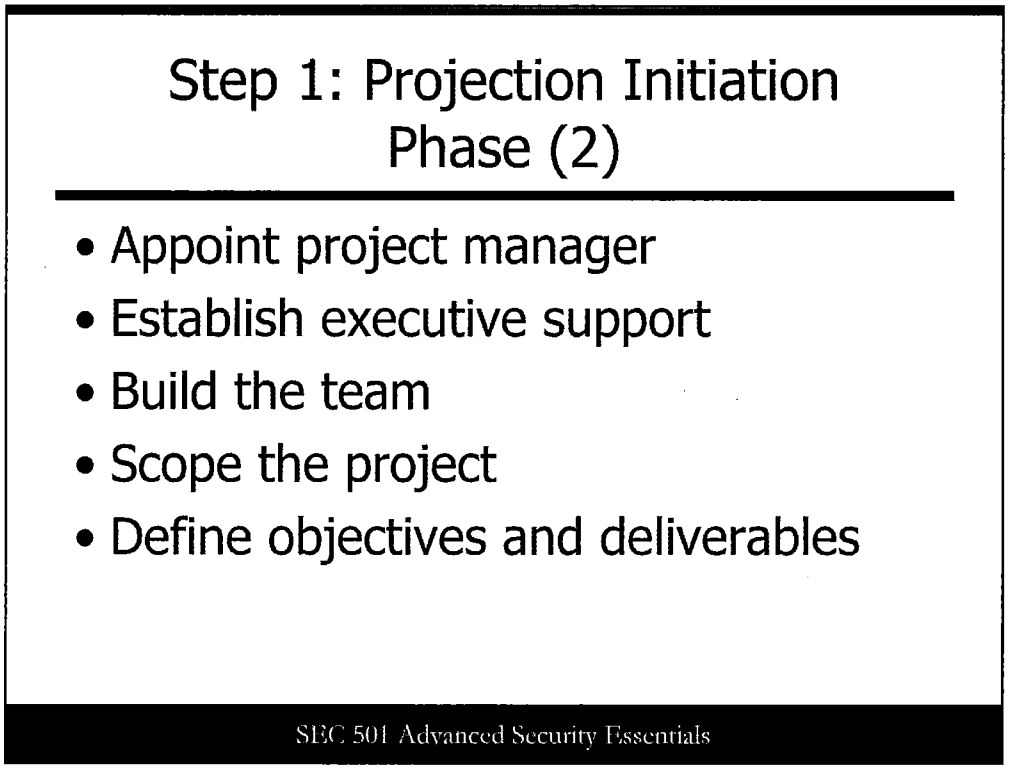

Project Initiation Phase (2)
Project initiation involves the following tasks:

- Appointing a project manager
- Establishing executive support
- Building a team
- Scoping the project (prioritizing)
- Defining the objectives and deliverables

Appoint Project Manager
As in any business project, you must appoint a project manager as the lead for the business-continuity plan. This individual acts as the enterprise-wide point of contact for continuity planning issues and interfaces with executive management. As the security officer, you are the ideal candidate to take on this role, but do not be surprised if the company makes this a new-hire, full-time position that is separate from the security officer. Whoever takes on the role must be enthusiastic and motivated about building the continuity plan. It is no small task and has many hurdles to overcome.

The project manager sets up the initial meeting of continuity team members and provides an overview of the business continuity process and all that it entails. In addition, the project manager does the following:

- Constructs a project budget that delineates costs to the company for creating the plan, including manpower costs, planning tools, reference materials, outside consultants, training, and equipment.
- Provides monthly update reports detailing activities, any difficulties that were encountered, subsequent resolution, and progress against the project timeline.

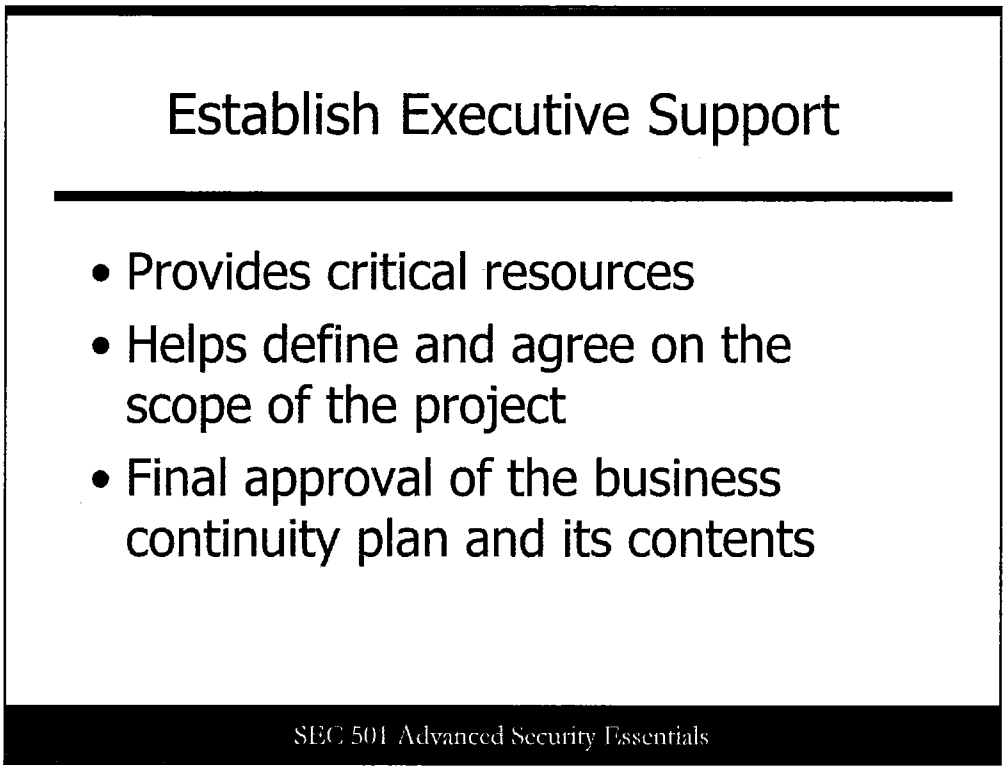

Establish Executive Support

At this point, establishing executive support should be a no-brainer because it ensures money, resources, and people. Most important, however, is that a business continuity plan is not yours; it is the CEO's. Ultimately, the business continuity plan is a statement of commitment from the CEO to the employees, customers, and colleagues about his willingness to provide reliable services and products in a timely and assured manner. It states, "I, as the CEO, will do my absolute best to continue to serve you."

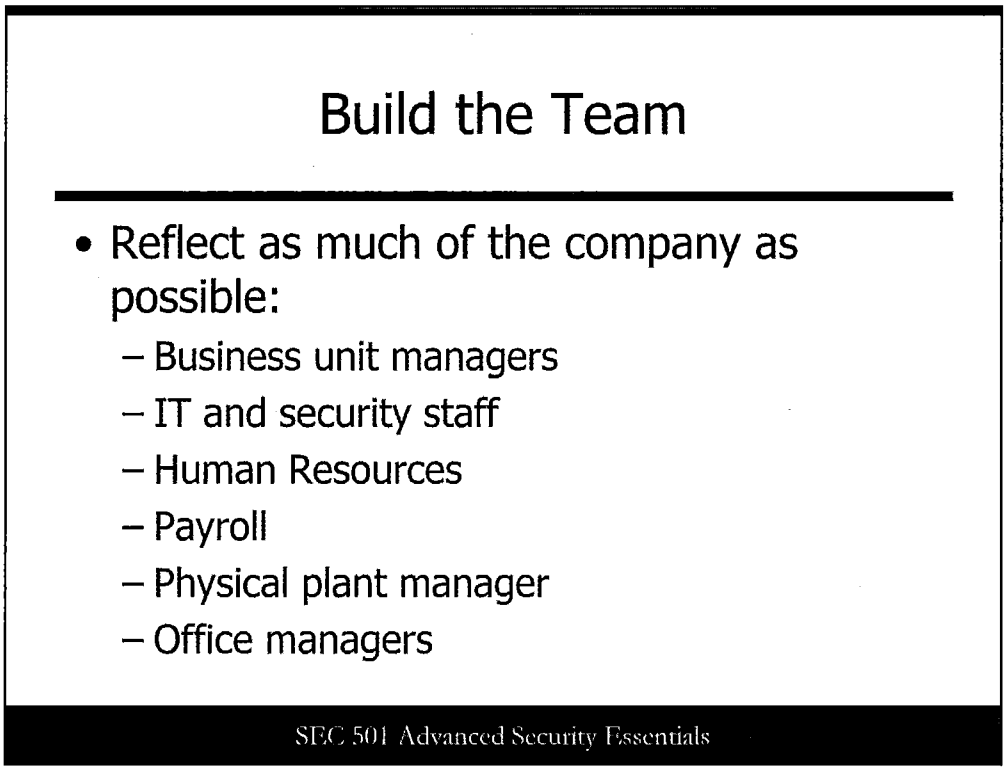

Build the Team
It is important to include as much of the company's hierarchy as realistically possible. There is no telling where a disruption might occur, and it is important to rely on the expertise of each member to add value to the business continuity plan. In reality, each employee has a vested interest in the company's capability to handle disruption because paychecks depend on it. Possible members include business unit managers, IT and security staff, Human Resources, payroll, the physical plan manager, and officer managers.

The temptation is to include too many people on the continuity team or fail to employ careful team selection and training. To do so might impede the process in particular and the value of the continuity plan as a whole. Include team members based on their positions in the company, the importance of those positions, and the business functions they represent, in addition to their ability to work in a collaborative and team-oriented manner. Project managers should be discerning but not repressive in the selection process. The business continuity plan's success depends on the success and ability of the team who creates, plans, and executes the plan.

Your team will eventually have the following categories:

- **Executive team:** Business unit managers, senior managers, and executive managers in the business unit who are responsible for recovering critical functions.
- **Management teams:** People in the command center who are responsible for managing, controlling, and guiding the recovery efforts.
- **Response teams:** Responsible for executing the recovery procedures and processes. Typically, you assign one team per critical business function.

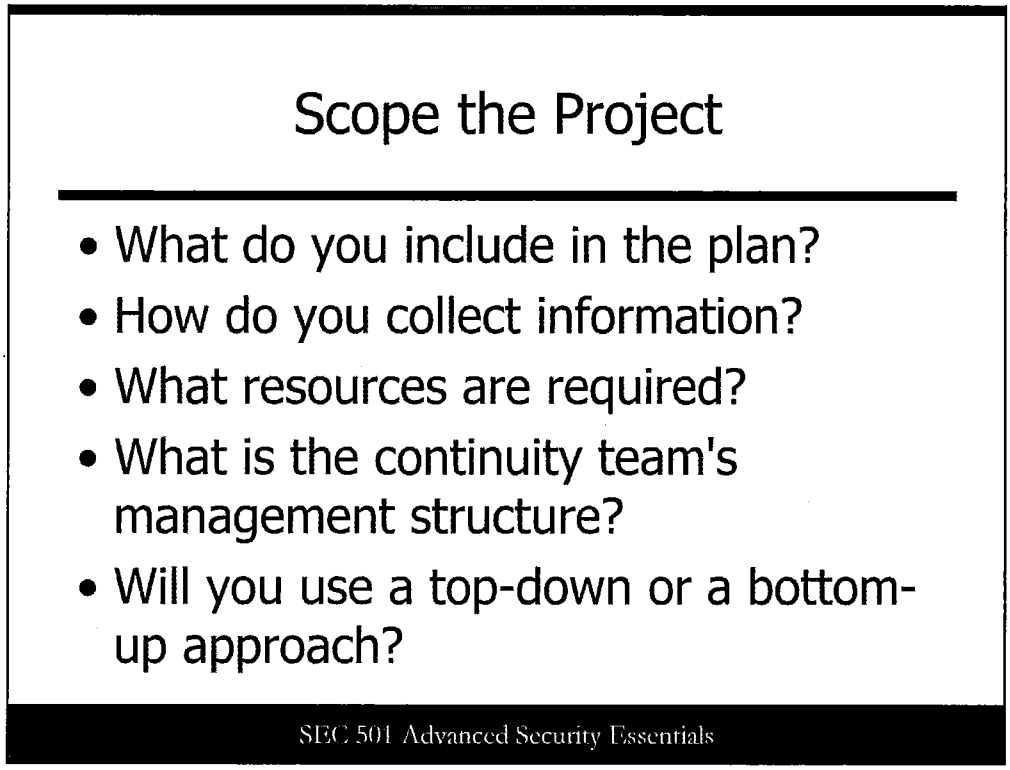

Scope the Project
A little bit of forethought goes a long way. As the size of the company increases, so does the necessity to detail the required work and the resources necessary to build an effective business continuity plan. "Detail" is the name of the game. You are creating a map of the company; therefore, the more detail, the better. You should closely examine the company's operations and describe them in as much detail as possible. Who gets paid when? Who are the vendors? Who are the customers? What is the company's workflow? On whom do you depend for raw material, paper, water, and electricity? How many offices do you have, and where are they located? How many employees do you have? What is your turnover? Who is your insurance provider, and what are you insured against?

In these cases, you are not looking for vulnerabilities or weakness; you are simply trying to build a map of the company so you know your assets. If you do not know what you have, you might not know what needs attention. Of course, detailing the continuity team's needs is just as important. How often does it meet? How does it handle all this information? How does it categorize this information? How much work does it have? What resources does it need to make this possible? This step gives the list-happy people in the group an opportunity to contribute to the effort. A final important point is to document the continuity team's management structure. Will it be a flat organization with individuals acting on behalf of their business units, or will you adhere to a strict hierarchy with information flowing up and down the chain in a predefined manner?

A top-down approach is a great way to create a company's first business-continuity plan. Beginning with senior management, the team interviews business unit leaders, middle management, IT management, and end users. In this manner, the continuity team can see the company from many different perspectives and build cross-company relationships that are necessary during the risk-analysis portion of the process. Again, you are not yet looking for vulnerabilities or weaknesses; you are simply

getting a feel for the environment. The bottom-up approach is better for the continuity plan's maintenance rather than first-time development. After you document all aspects of the organization, the bottom-up approach works better as a sanity check and a measure of the effectiveness of the current continuity plan.

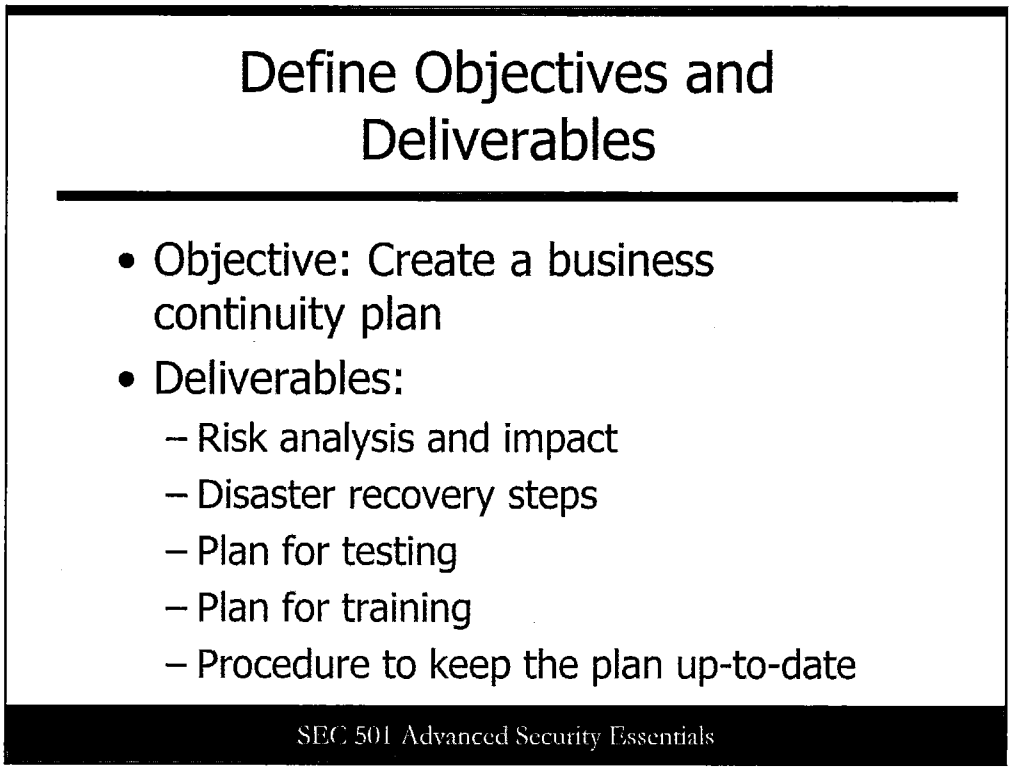

Define Objectives and Deliverables
The objective of the business continuity team is fairly self-explanatory: to produce a continuity plan that enables the corporation to recover as quickly as possible from unforeseen disruptions that interrupt the normal flow of information and or business.

The deliverables can be just as simplistic as the objective statement: a spiral-bound business continuity plan for all involved parties. However, such a statement is a bit too broad. The business continuity plan is actually a number of documents rolled into one, and you should list each as a deliverable. There is no telling how complex the plan will be or how many situations the continuity team might face, so you should account for each section of the plan. In addition, depending on the sensitivity of the information, you might exclude the risk analysis from the document for purposes of confidentiality.

This sort of documentation is usually required for federal agencies, but it is also beneficial to commercial agencies. Having clearly defined processes, procedures, roles, and responsibilities eliminates confusion, reduces misinterpretation, helps facilitate the training of new personnel, and enables management to understand the environment in which the business operates. The documentation can also serve as justification for additional funding.

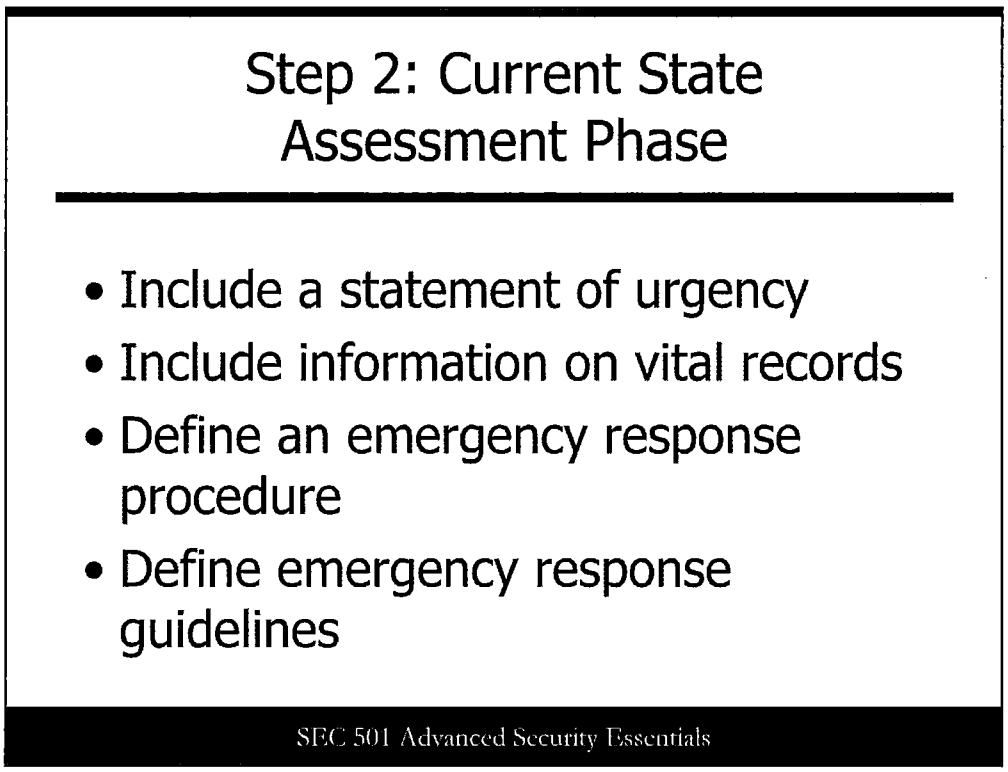

Basic Elements of Contingency Planning
Current State Assessment
The plan should clearly define organizational responsibilities, emergency response procedures, and emergency response guidelines. It is critical that roles and responsibilities be clearly and unambiguously defined. Time is too precious during disaster recovery for turf wars, inadequate empowerment, or lack of clarity regarding who is in charge. Security controls might need to be altered during recovery, and the empowerment of certain personnel might have to be increased. Definitions for when and under what circumstances such changes from the normal begin and end must be clear. This should all be documented clearly to avoid any misunderstandings, security violations, or delays in recovery.

Information about vital records must be included in the plan. This is often the bulk of the documentation and might include lists of people and how to contact them, inventories of equipment, software and data , including offsite backups and how to obtain them; vendors and how to contact them; information about emergency services; network diagrams; media contacts; and so on. This is where the planners' imagination becomes evident. We become used to the information within our systems and take access to it for granted. Think about this question: "When I cannot use my system(s), what information that I normally rely on is no longer available?" Some needs might be different during recovery than normal. For example, electronic communications needs can change as a result of users or systems being relocated. But these needs can also be inventoried. Site-knowledgeable personnel is the most important recovery component.

This is especially so because disasters are usually only partial losses, and initial recovery consists of patching together what is left, including business processes. This takes knowledgeable people.

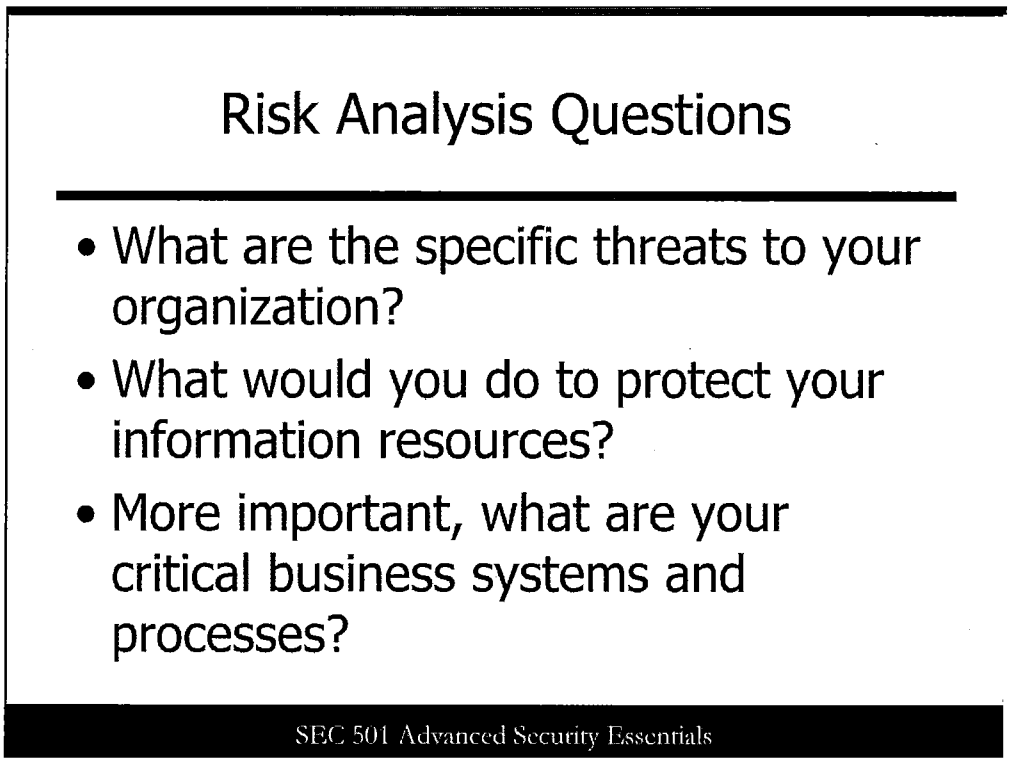

Risk Analysis Questions
Here we focus on risk analysis as a component of contingency planning. For that purpose, risk analysis consists of the following steps:

- Identifying your critical business systems and processes
- Identifying the specific threats to your organization, especially to those critical systems and processes
- Evaluating the vulnerability of an asset and the probability of an attack or disruption to occur
- Determining what you would do to protect your information resources
- Weighing the loss of assets versus the cost of implementing mitigating controls

Risk can be expressed as follows:
$Risk_{(due\ to\ a\ threat)}$ = Threat x $Vulnerability_{(to\ that\ threat)}$

Or more completely, as expressed in ISO 17779 and many risk management methodologies :
$Risk_{(due\ to\ a\ threat)}$ = Threat x $Vulnerability_{(to\ that\ threat)}$ x Impact

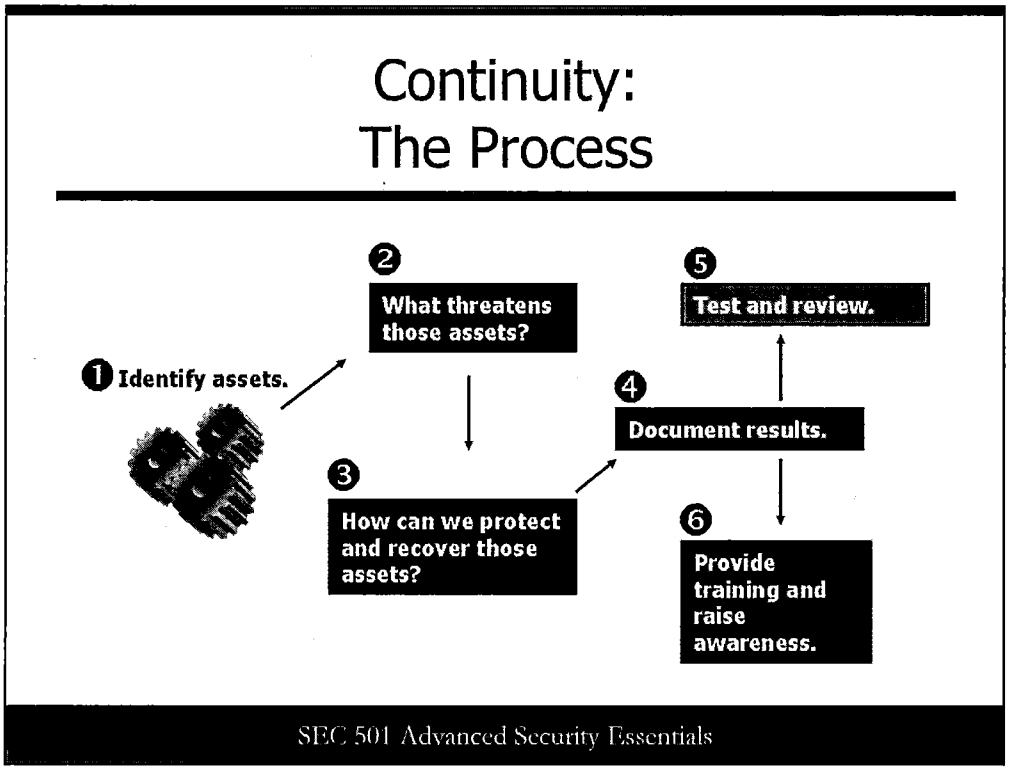

Continuity: The Process

The nice thing about pictures is that they make everything look so simple. Although the preceding figure makes the process of continuity planning look straightforward, the details are numerous. However, keeping the big picture in mind can help you maintain perspective, regardless of how complicated the actual construction of the business continuity plan becomes. Note that this process has many variations, but the preceding figure represents the most general ideas.

The following list outlines the business continuity plan steps you must take for each step:

1. Identify assets.
2. What threatens those assets?
3. How can we protect and recover those assets?
4. Document the results.
5. Test and review.
6. Provide training and raise awareness.

The good news is that in the process of building your security program, you have probably already accomplished steps 1, 2, and 3. What you probably have not covered is how to recover if protection fails. Step 4 involves creating the business continuity plan itself: the documentation that you will print and distribute to all appropriate personnel.

Step 5 involves testing and reviewing the continuity plan. There are many different ways to ensure that the continuity plan actually works, ranging from inexpensive and nonintrusive to cost-prohibitive and expensive. It is important to create a continuity plan that is testable, but some industries might not test their plan because of operational constraints. Surprisingly, this is not uncommon with casinos, trading

floors, and other high-paced entities. However, most corporations can test at least some part of the plan, if not the entire plan. Testing makes sure that the assumptions and ideas with the continuity plan are sound and sufficient. Reviewing the continuity plan is another critical aspect of continuity planning. If the plan does not properly reflect the current state of the corporation, the continuity plan might be, for all intents and purposes, useless.

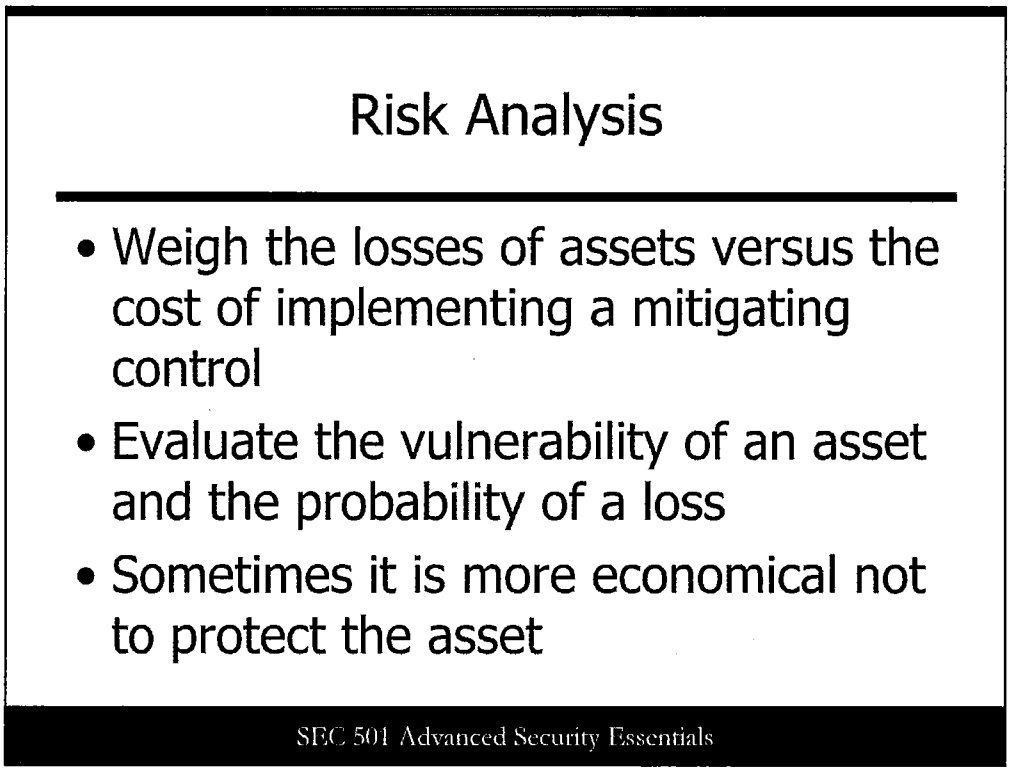

Risk Analysis
After you understand the risk, you can do one of four things:

- **Risk avoidance**: When you decide not to become involved in the risk situation.
- **Risk acceptance (also termed risk assumption or risk retention):** When you acknowledge and accept that the risk is something that could happen. You intentionally or unintentionally retain or assume the responsibility for loss or the financial burden of loss within the organization.
- **Risk transfer**: When you shift the responsibility or burden to someone else. An example would be getting insurance to cover the damage.
- **Risk reduction**: When you apply the appropriate controls to mitigate the effects of the disaster, thereby reducing the risk.

Next, we look at the impact component of business risk as we perform a business impact analysis, or BIA.

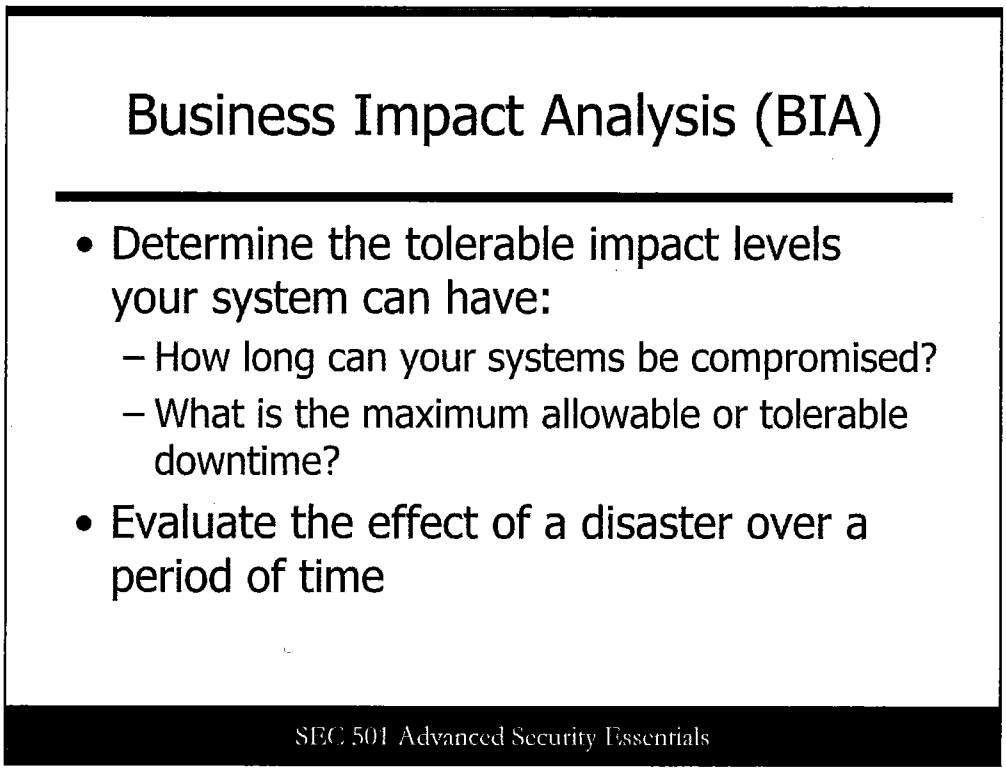

Business Impact Analysis (BIA)
After risk analysis in the BCP-DRP planning process life cycle comes business impact analysis (BIA), where you determine what levels of impact to your system are tolerable, such as the duration of system outage,.

The process of developing the BIA typically involves interviewing key users of the various computer systems (for example, payroll, accounts payable, and accounting) to get a better understanding of how a disaster could impact the ability to continue operations. Some of the key interview questions might include the following:

- How would an information technology failure affect cash flow?
- Would the disaster impact the level of service?
- How long could the outage last before it began to affect your productivity?
- How long could operations continue if data were unavailable?
- Would there be irretrievable loss of data?
- What key resources are required to continue operating?
- At what point would those resources need to be in place?
- If we implement a mitigation strategy, will there be additional risks? If so, are we better off by implementing or not implementing?

The answers should come from or be agreed upon by executive management. Executive management understands such cost trade-offs as mitigation and loss and has individual accountability either way. Lower management might err toward too much (that is, too expensive).

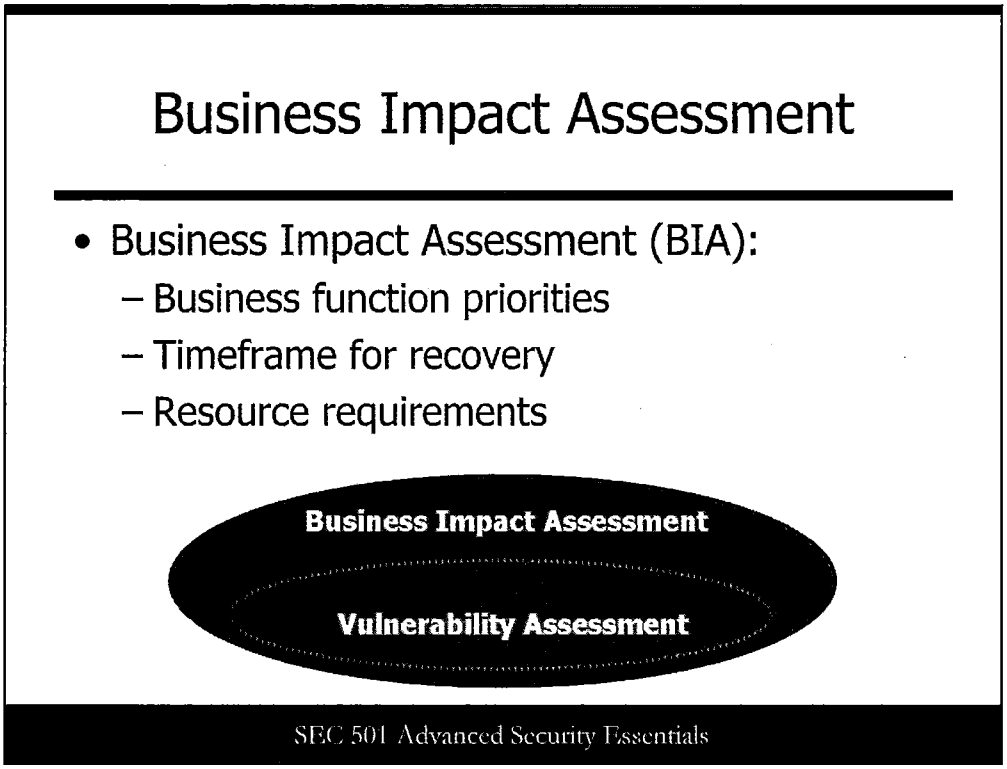

Business Impact Assessment
The business impact assessment (BIA) documents the impact a disruptive event might have on a corporation. The BIA uses the information in the vulnerability assessment to prioritize business functions and calculate business impact.

Obviously, the greater the impact of a business process should it fail, the higher its priority in terms of criticality. A direct relationship might exist between the criticality of a business function and the timeframe for which it must recover. Some business functions, if down for only a few seconds, might dramatically and detrimentally impact the business. Other business functions might be interrupted for days or weeks and have no negative affect on the corporation.

The primary goal of the BIA is to determine the maximum allowable downtime for any given system. A rough guide for timeframes follows:

- **Immediate recovery:** No downtime allowed. Implement a fully staffed, fully equipped alternate site (more on alternate sites later).
- **Quick recovery:** Up to four hours of downtime allowed. Pre-equipped alternate site should be available. Staff can arrive at site within four hours.
- **Same-day recovery:** You can move equipment to another location and set up in an eight-hour period. Same-day recovery can also mean same *business-day* recovery. The alternate site can be anything that affords appropriate power and protection (another office, hotel room, home, and so on).
- **24-hour recovery:** This is self-explanatory.
- **72-hour recovery:** This is self-explanatory.
- **Greater than 72-hour recovery:** This is self-explanatory.

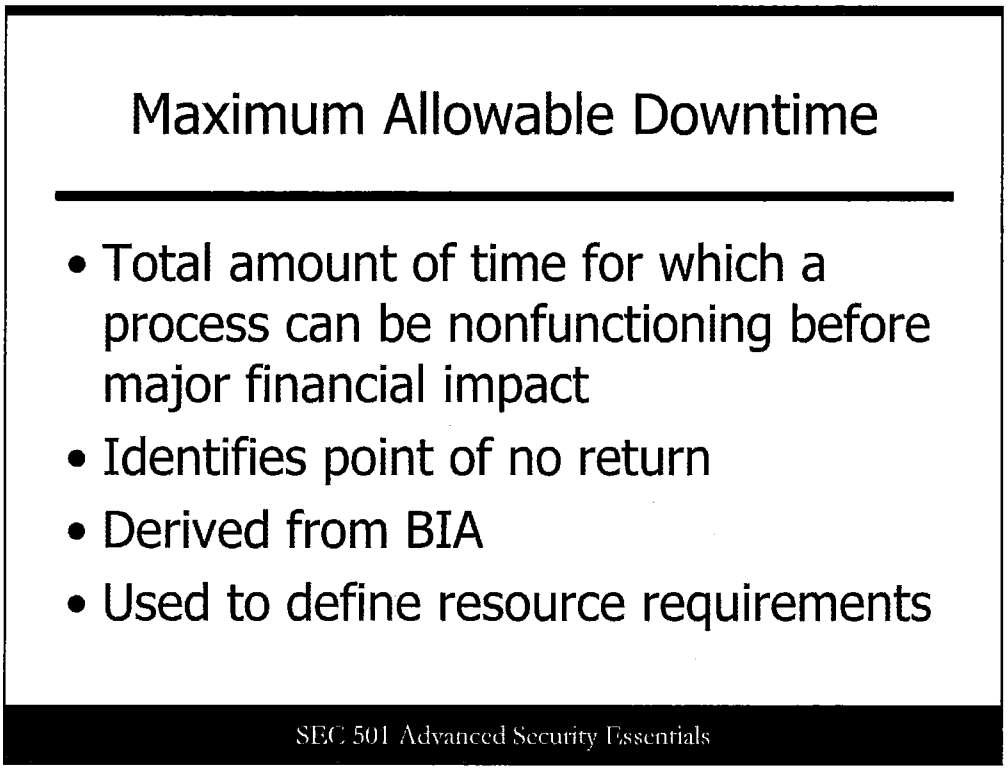

Maximum Allowable Downtime

Maximum allowable downtimes derived from the BIA are the basis for determining recovery resource requirements. The ultimate objective is to define the post-disaster resource requirements upon which a recovery strategy might be based. Ultimately, the need for recovering critical business processes drives the recovery resource requirements (the recovery budget). Often senior managers try to set an arbitrary budget of what the company is willing to spend on recovery and thereby force the business continuity team to work within those parameters. However, the needs of the business should drive the budget, not the other way around.

What is the financial impact you should consider? The following are a few points to consider:

- How much revenue would your company lose if its systems were unable to accept orders?
- What is the cost of lost productivity?
- How much inventory would be lost or spoiled; and how much would it cost to recover the inventory?
- What is the value of IT professionals' productivity while trying to resolve the problem?
- What fines and fees would the company have to pay?
- How much would a public relations campaign cost to regain your company's image?
- Will the company face any legal, health, safety, or liability exposure?
- Can you really afford the cost of implementing a 24X7 operation without any downtime? Do you have the personnel and technical resources to do this? If not, how do you prioritize?

Many companies make a critical mistake by assuming that they should conduct BIAs only once before writing the initial business continuity plan. One BIA is usually never enough; you should conduct BIAs over the lifetime of a corporation, especially if you undertake any major technology upgrades or add, modify, or delete processes. Any alterations to the business affect the recovery-resource requirements. These resource requirements drive your recovery strategies, so it is important to ensure that the business needs are properly reflected.

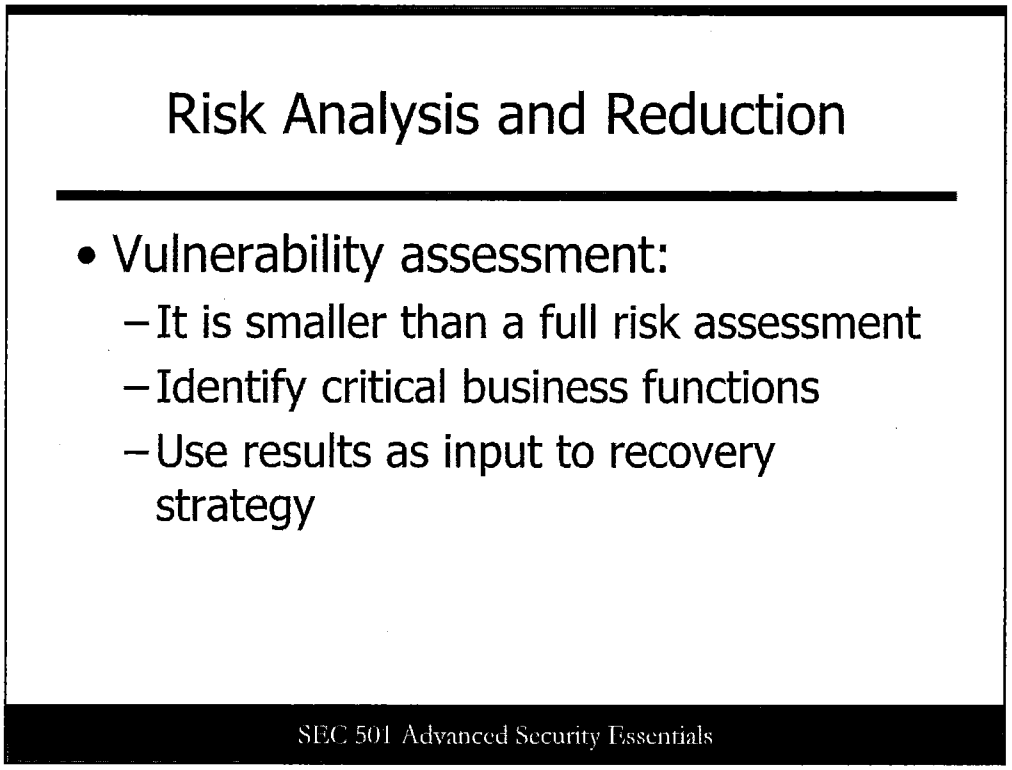

Risk Analysis and Reduction
The vulnerability assessment is typically part of the business impact assessment (BIA) and is smaller in size and scope than a full risk assessment.

The focus of the vulnerability assessment is to provide data that is used solely as input into the recovery strategies and determine the impact of losing a critical business function. You identify and target critical business functions. Any business function that must be present to sustain the continuity of the business or could in any way threaten human life during a failure is a critical business function. In addition, if a business function's failure brings discredit or public embarrassment to a corporation, it is also considered critical. Such a determination is usually up to the corporation, however, because only the business can determine what it deems embarrassing.

Vulnerability assessments provide the bulk of the financial and operational costs of a disruptive event. If anyone in your company is looking for hard numbers, this is the place to find them. You can establish quantitative and qualitative criteria to determine financial and operational costs.

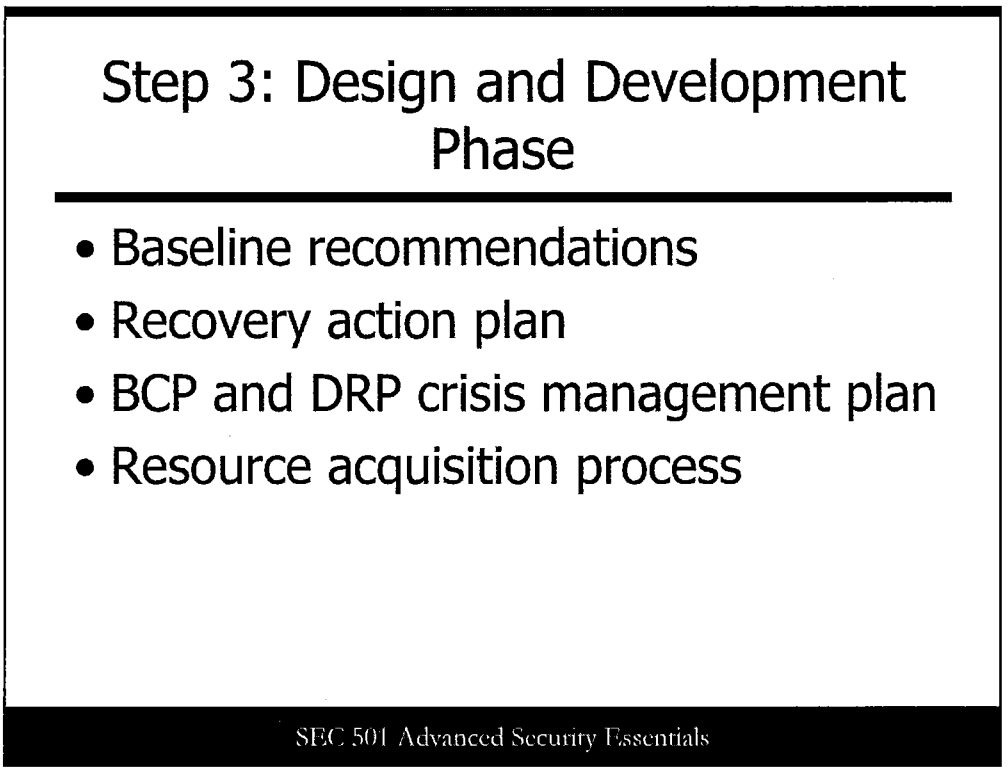

Design and Development
Bases of baseline develop recommendations:
After the baseline is developed, this is used to ensure a documented process to recover to the designed baseline.

Action plans regarding next steps of recovery:
Documented process of procedures ensure timely flow of required steps in the recovery process. Each step of the process must be document.

BCP and DRP crisis management plan:
This allows for defined rules testing and implementation if the need arises.

Begin recover resource acquisition process:
Allows for the identification and acquisition of materials, assets, and locations required to meet the needs of the action plan.

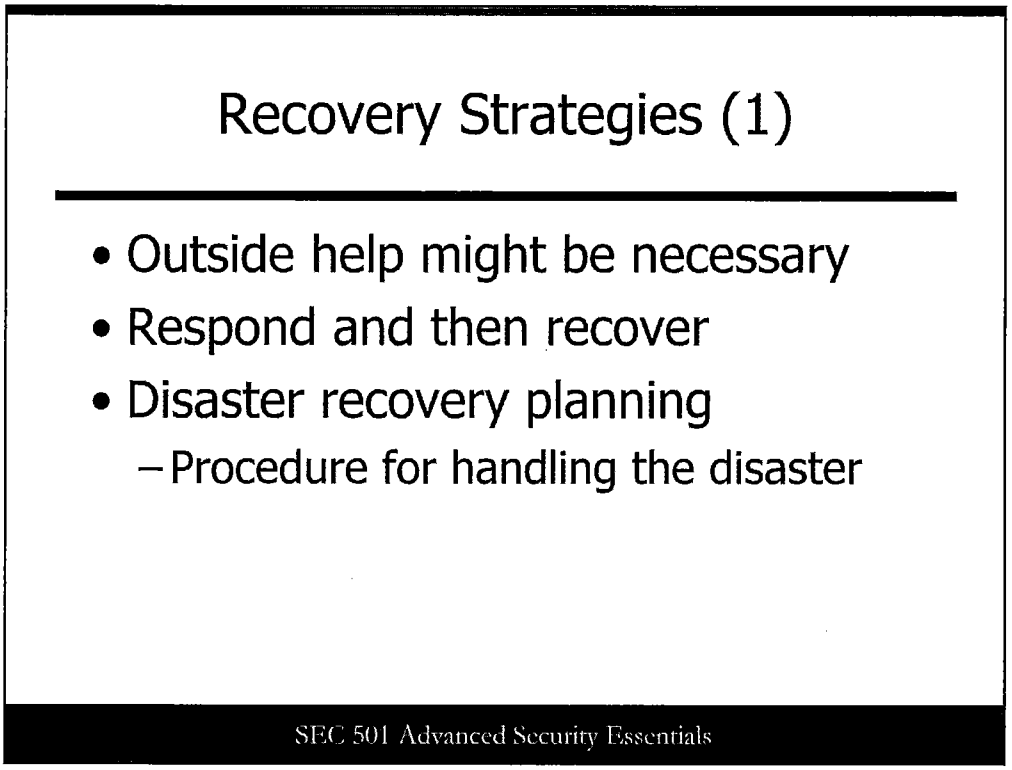

Recovery Strategies (1)
After the BIA is complete and the corporation knows what the impact of a loss might be, it can start planning an appropriate response. The fun—and frustrating—part of recovery strategies is that you are solving problems that do not yet exist. Essential to any top-notch recovery plan are multiple strategies to cover most, if not all, business disruptions. However, actually accomplishing such a feat can be challenging and, in some cases, unrealistic. Most often, companies must look to consultants, vendors, and similar companies for advice and recovery strategy trends. Developing recovery strategies can also be the most frustrating part of the continuity planning process because this is when things get challenging. Costs might loom large, the complexity of replicating, backing up, or mirroring critical business processes all of a sudden seems too prohibitive, and executive and senior management begin to waiver in the face of the beast that is their own business.

The sanity check for the growing hysteria of your senior management comes from the BIA that you, as the team lead, so diligently conducted. Compare the options for attaining recovery, when investigated and priced, against the potential cost of failing to recover. Contrasting these numbers might make senior management wholeheartedly continue with the continuity process. Suppose it costs $1 million to recover a business function, but the company would lose $1 million per day if it were unable to recover in a timely manner. Of course, the numbers are not always that simple, but that is why the BIA is so important: It gives the company some type of objective measurement.

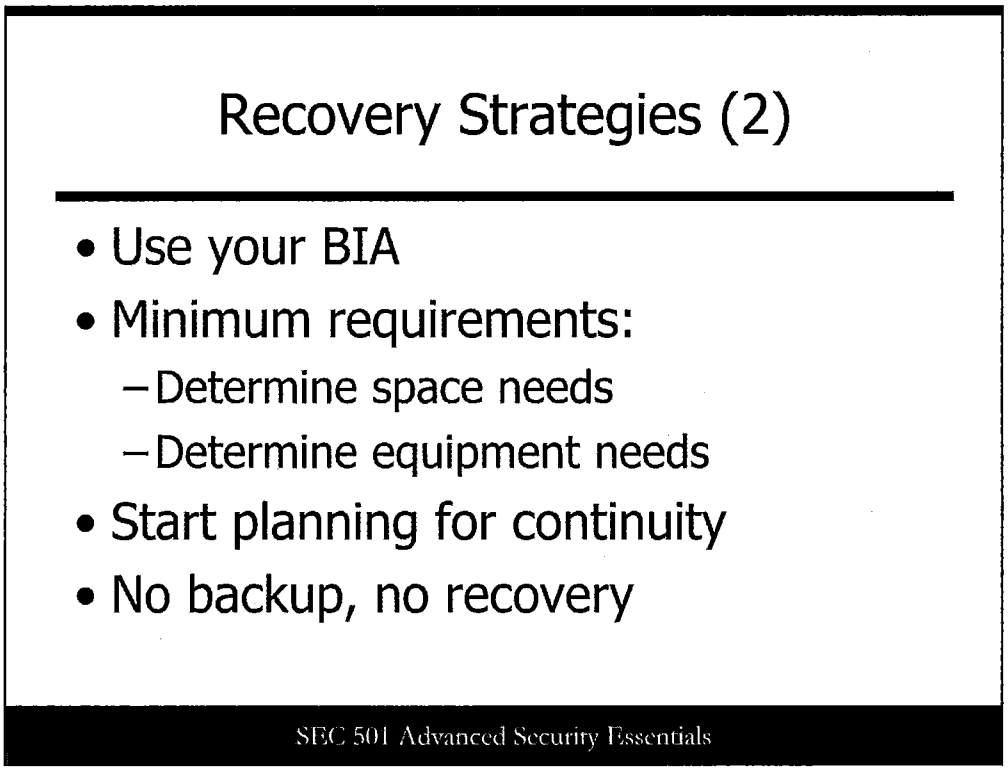

Recovery Strategies (2)
All recovery strategies are driven by the maximum allowable downtime of a given business function and the resources required to continue to perform that function. At a minimum, planners should determine the necessary space and equipment needs for continuing the critical business process and their availability. It might be necessary to put agreements in place with vendors and suppliers to provide equipment, office supplies, services, and even personnel if a disruption occurs.

The disaster recovery plan enumerates all necessary information if a disruption occurs, including but not limited to the location of the emergency operations center (EOC), directions to the EOC, the location of alternative recovery sites (also with driving directions), team members and all contact information, the procedure for handling the disruption, and the declaration and notification procedures.

In all cases, no matter what the recovery strategy is, you must have arrangements in place for the recovery of vital records, whether they exist in hard or soft copy. No backup, no recovery: the mantra of business continuity. Without backups, the business has no way of picking up where it left off.

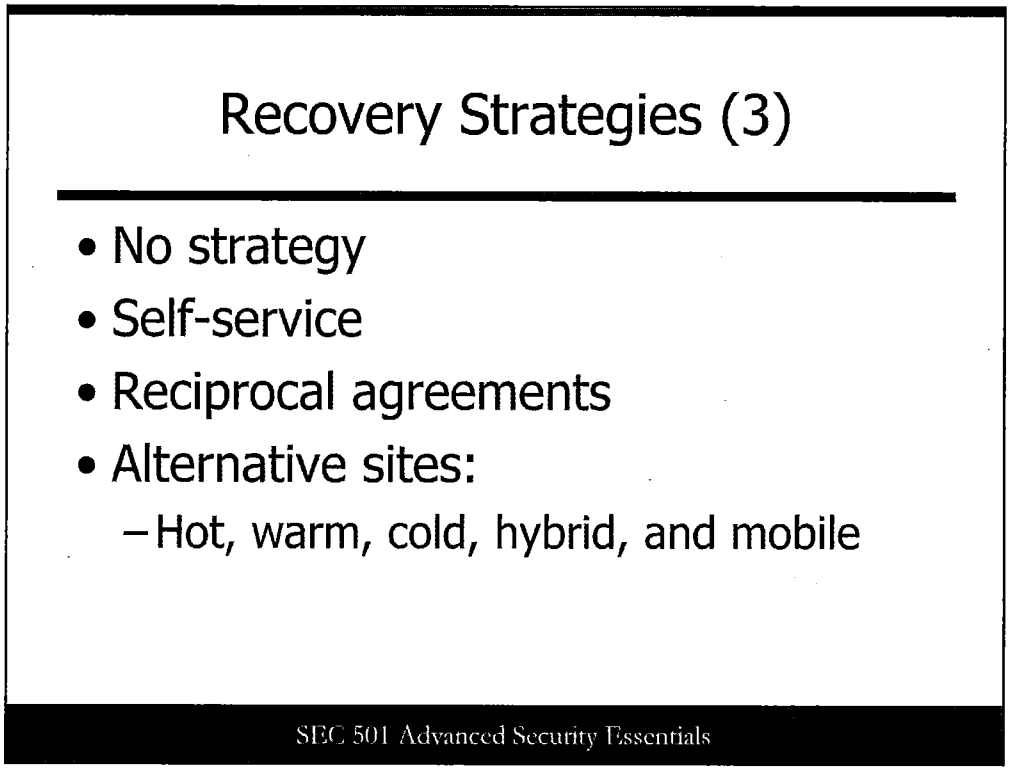

Recovery Strategies (3)
Understandably, for certain business functions the cost of recovery might not be justified. A business function of this type is most likely a low priority because it is not truly critical to the survivability of the corporation. You must use sound judgment, but do not be surprised if you find it reasonable to put no formal response in place.

A self-service strategy uses the corporation's offices to transfer or host disrupted business functions. Conference rooms, cafeterias, satellite offices, training rooms, even employees' homes might be equipped to temporarily support business functions. Given the scope of the disruption, this might or might not be a plausible strategy.

Reciprocal agreements involve making arrangements with other (possibly competing) companies that have similar needs to your own. Depending on your industry's specialization, there might be only a handful of businesses with unique operating needs. These needs might make it too cost-prohibitive to replicate business functions; therefore, by forming a reciprocal agreement, each company agrees to help the other if a disruption occurs.

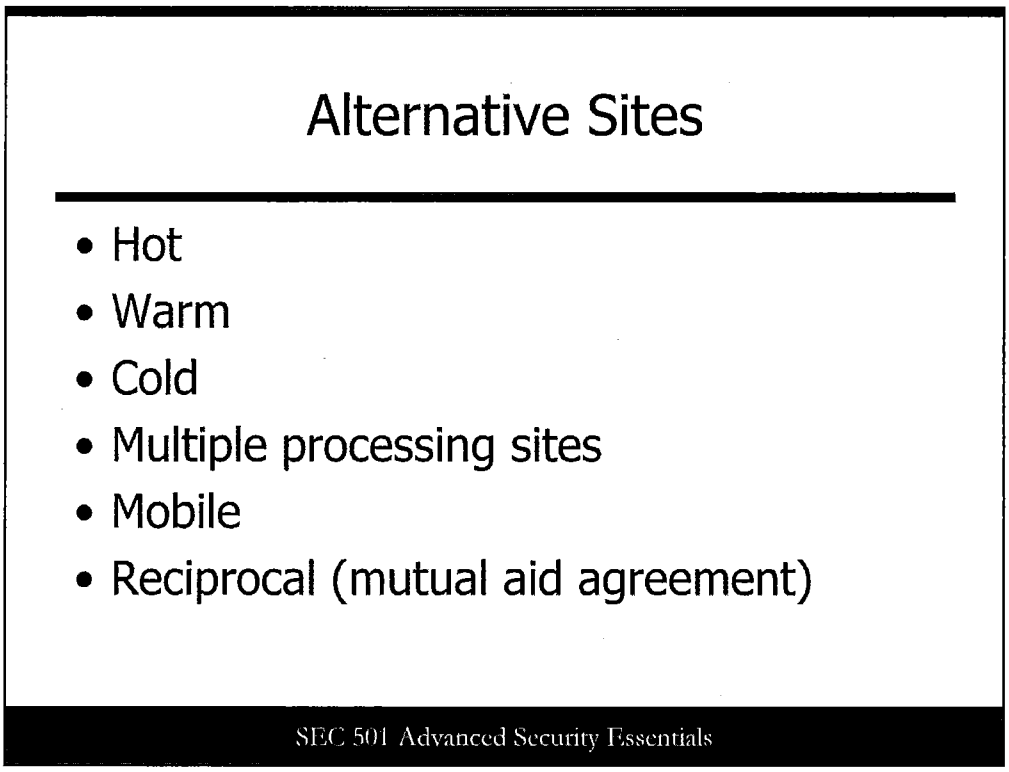

Alternative Sites
Third-party continuity service providers offer many types of alternative sites, ranging from empty shells to full-fledged operation centers:

- **Hot sites**: Fully equipped and staffed facilities running 24/7 that intend to serve an organization that has sustained total disruption either through catastrophic failure or total physical destruction. A hot site is feasible for critical business functions that cannot tolerate any downtime.
- **Warm sites**: Descriptions vary for a warm site, but it is primarily a facility that is pre-equipped but not necessarily ready to go. Business processes that can tolerate a few hours of downtime might be ideal candidates for a warm site, but companies should fully investigate what they are paying for and what they are getting.
- **Cold sites**: The simplest and least responsive of the alternative sites, the cold site is simply an empty facility the company must equip if a disruption occurs. Given that the response time of a cold site is in the order of several hours, if not days, it is debatable whether a cold site will meet the recovery requirements of a business.
- **Multiple processing sites**: Multiple internal processing locations geographically dispersed to assist in the backup and recovery of vital company data, aka (Mirror).
- **Mobile sites**: Mobile sites are routinely identified as the up-and-coming alternative sites because they provide almost the same capabilities as a hot site, but not quite. Depending on the type of disruption, a mobile site is akin to an "office on wheels" that you can locate conveniently near the company, thereby precluding extensive employee travel and personal issues such as daycare and special needs. Unfortunately, mobile sites can realistically meet

only a 12- to 72-hour response time, depending on the proximity of the service provider who will deliver the mobile facility. Obviously, the closer the service provider, the quicker the response. Depending on the business functions' maximum allowable downtime, a mobile site might not be a viable option.

- **Reciprocal**: Formalized agreement between two business entities to facilitate recovery after a disaster. Agreements could include temporary office space and use of company resources to resume operations.

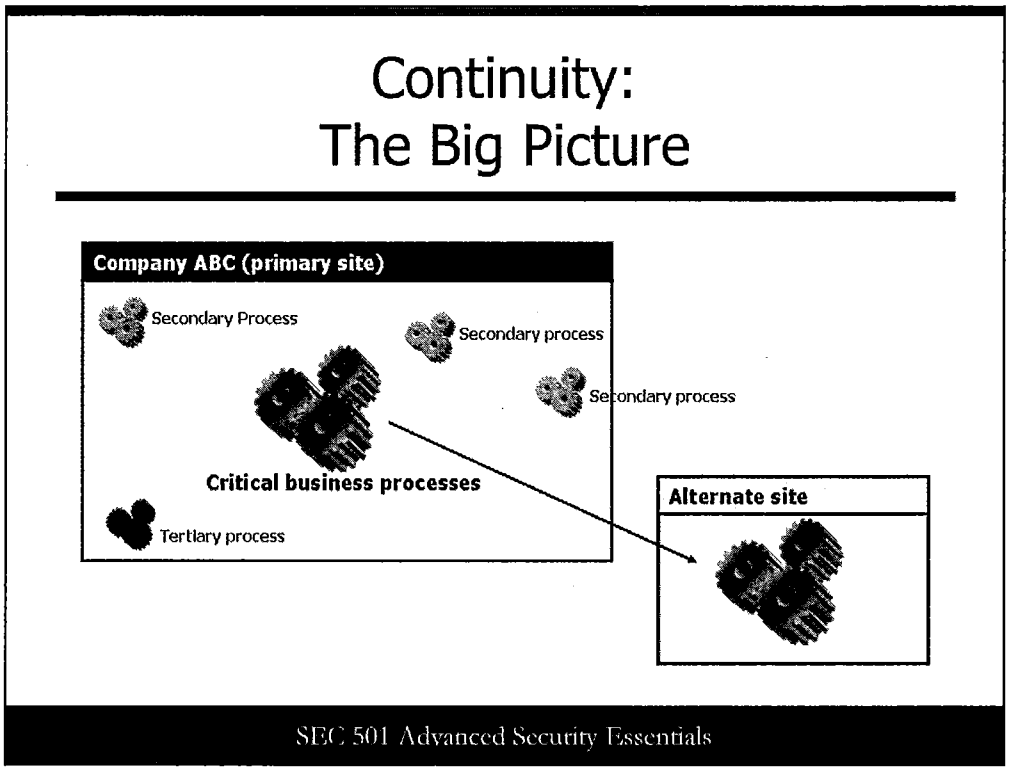

Continuity: The Big Picture
There is a critical business process at the center of the corporation's moneymaking machinery. It might be the production floor of an automobile manufacturer or the web servers of an ecommerce firm.

Whatever the process—and whatever technology supports it—you must identify the process and ensure its continued operation. This is not to say that you should ignore the secondary and even tertiary processes—only that you prioritize in the grand scheme of things.

Some business processes, if disrupted for only a moment, can rob the enterprise of valuable income or market agility, whereas other processes might tolerate disruption for several days without any adverse affect. The big picture of business continuity is making sure your business keeps making money.

In this diagram, Company ABC conducts business at a primary site; that is, the critical business process functions at the critical facility. In addition, secondary and tertiary processes surround the critical business process.

Company ABC constructed an alternative site where the critical business process can be moved or continued if the primary facility is unavailable for any reason. Note that the alternative site does not provide any support for secondary or tertiary business processes; this might not always be the case.

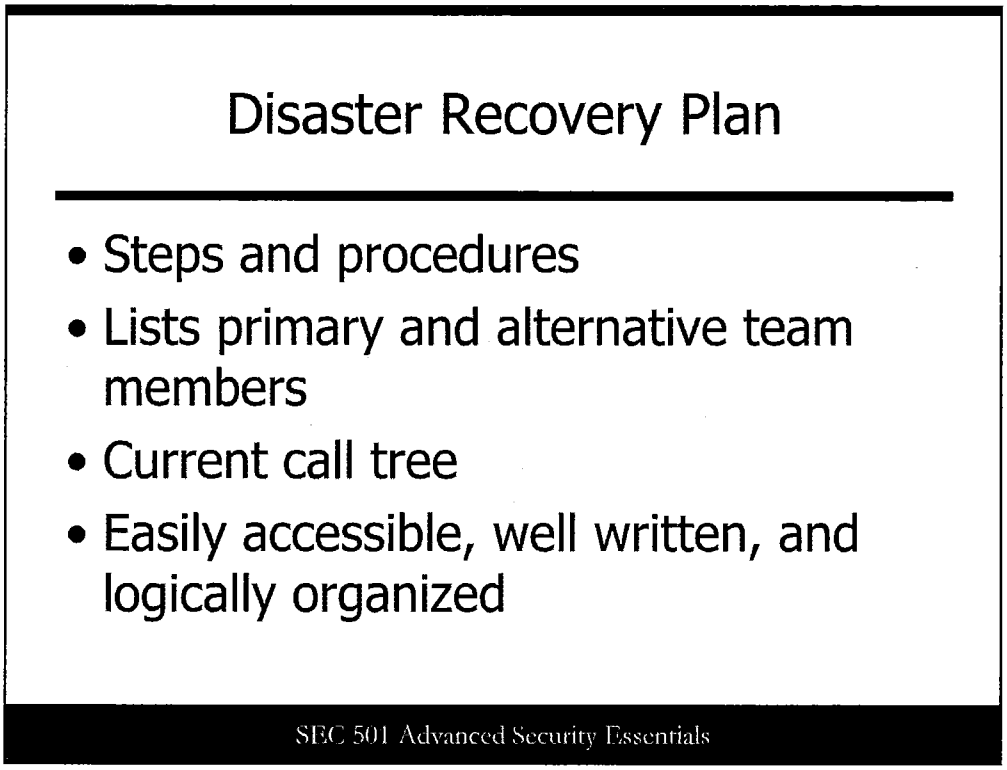

Disaster Recovery Plan
The disaster recovery plan is the corporation's immediate response to disruption. As such, it is a primary recovery strategy and is, itself, a critical business function because it ensures the company's ability to recover. The plan should include a list of all primary and alternative team members who are responsible for handing the crisis, specific and detailed steps that members should execute, and a current contact list (call tree) of all personnel and the functions they are qualified to perform.

Most notably, because the disaster recovery plan will be used by people under duress, the plan should be well structured, clear, concise, and complete. The first draft of a plan might be none of these, which is why testing is so important in business continuity planning. As a purely academic process, business continuity planning is of limited use. Where the plan is worth its weight in platinum is how it is revised during the testing phase.

The structure of a disaster recovery plan might include the following elements:

- **Introduction**: The organization's goals and, in general terms, what it considers an emergency and the potential risk.
- **Emergency management team**: Who is on the emergency response team, contact numbers, roles and responsibilities, and alternative team members.
- **Emergency operations center**: Operational concept, location, driving directions to primary and alternative sites, availability of communications, communications protocols, and physical and logical layout.
- **Emergency notification procedure**: How the organization will be alerted to an event and the communications protocols for notifying internal and external entities, particularly relating to the news media. List contact information for all pertinent personnel and the conditions under which they will be contacted.

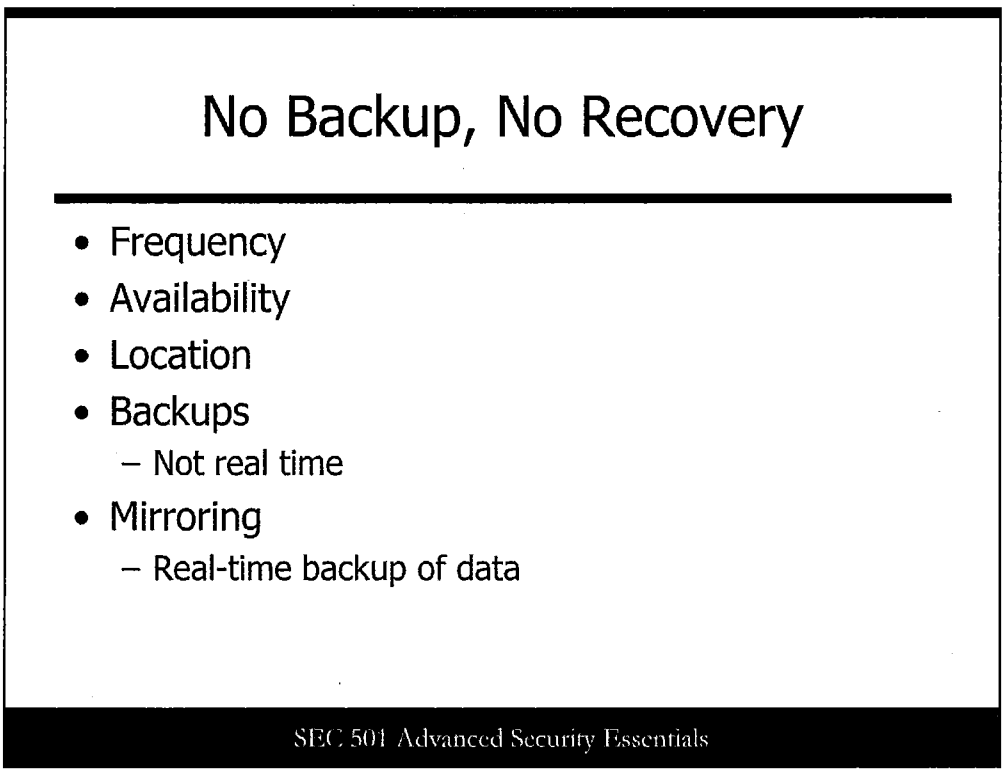

No Backup, No Recovery
Without backups, a company cannot recover...at least, not quickly. You must back up all vital records relating to a corporation and duplicate any hard copy. The archival process of backups allows a company to return to some specific time in the past and rebuild its business functions. The more time-synchronized backups are with data as it is created, the quicker a corporation can recover from a disruption. The pinnacle of time-synchronized backups is called *mirroring*.

As important as time is to backups, location is equally important. If a system's backups are destroyed, so too is the company's capability to recover. Storing backups offsite increases the likelihood that backups will survive most types of disasters or emergencies. Storing backups within the same physical facilities as the company defeats the purpose of trying to protect your data.

The next consideration is the availability of backups when the time comes to recover a system. Will the backups be delivered to the company's primary or alternative site? How long will it take? Is delivery time guaranteed? Will a corporate representative pick up the backups? How long will it take to complete the backup process? Is that time within the maximum allowable downtime timeframe? Who will restore the system?

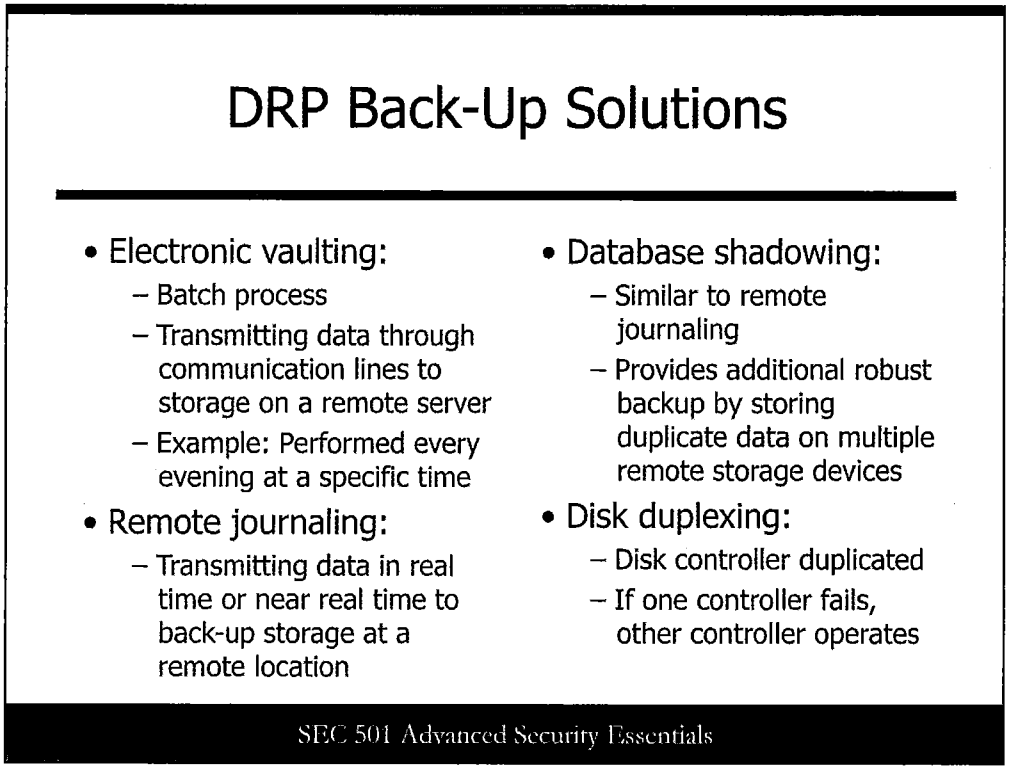

Transaction Redundancy Implementations
There are many approaches that are used to backup databases. One is done in near real time (journaling) and one is done in batch mode (vaulting), usually at the end of the day.

In remote journaling, data is available at the backup at any time and provides a high degree of fault tolerance in the event of a disaster.

Note that disk duplexing does not denote multiple disks as backups, but multiple disk controllers.

Clear Plan: What is expected for how things should go and what are the anticipated outcomes for any planned event?

- **Short term**: This includes both testing and maintenance.
- **Long term**: This includes both testing and maintenance.

Testing and Training Strategies

The area most overlooked is the area of training and testing. Without these areas the BCP is likely to fail. You must educate the workers and management on their roles and responsibilities. Live drills:

- Checklists
- Walk through
- Simulations
- Parallel
- Full Interruption

Enterprise Crisis Management Plan

EOC development planning, and testing. This includes the type of location manning and any other pertinent requirements that are business-specific.

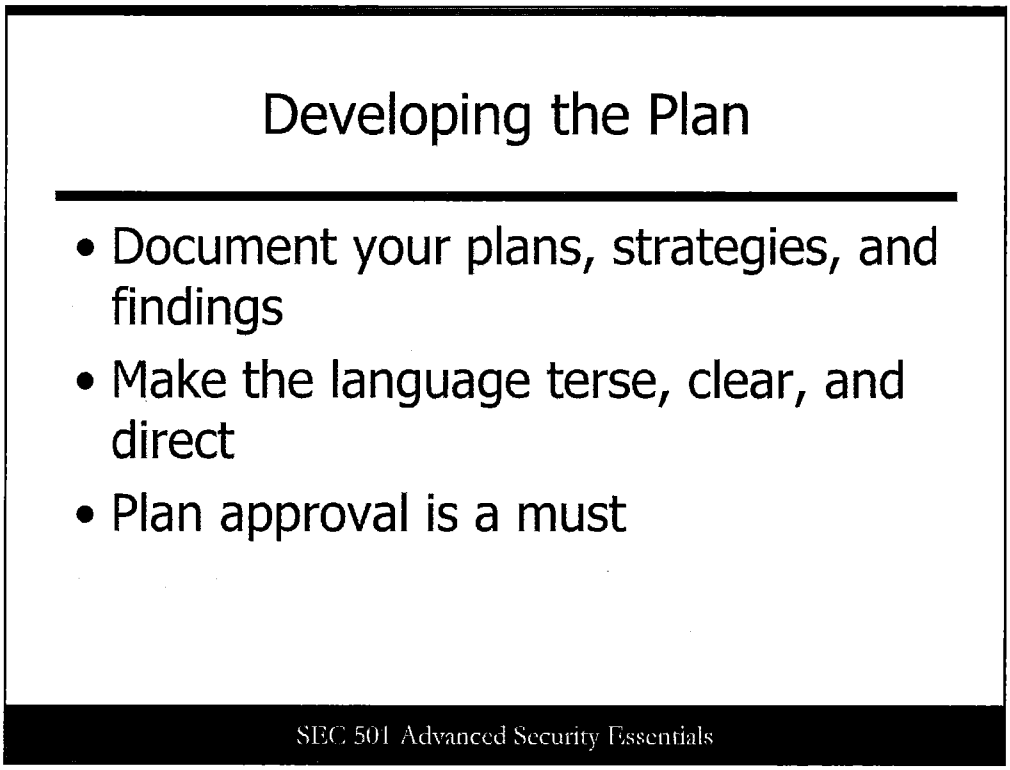

Developing the Plan
Business continuity plan documentation is the end product of the process and is therefore terse and direct. Business continuity plans take on many forms and many different formats, but you can organize one as follows:

- **Introduction**: Explain why the plan is necessary and detail the scope of the plan, including who is part of the response and recovery process and the range of events addressed.
- **Crisis-management structure**: Include details about the roles and responsibilities of everybody involved.
- **Locations**: Document the location of the command center and the procedure for activating the center. Include the location of alternative and backup sites.
- **Procedures**: This section includes the alert procedure when an incident is first discovered, damage assessment, declaration procedures, notification procedures, and team procedures.
- **Exercise log**: Document the calendar date on which you tested the continuity plan, what type of test you conducted, and any shortfalls you encountered during the test (phone numbers that were out of date, a team member who no longer works at the company, and so on).
- **Revision history**: Document the date changes that were made to the document, the person who made the changes, which executive approved the changes, and the details of the change.

The executive management team should sign off on the finalized plan. Although the business continuity plan is finished when the executive(s) signs it, it is far from complete. You must now flesh out the document's usefulness, removing any theoretical remnants and leaving only tested, certified, and accurate procedures for recovery.

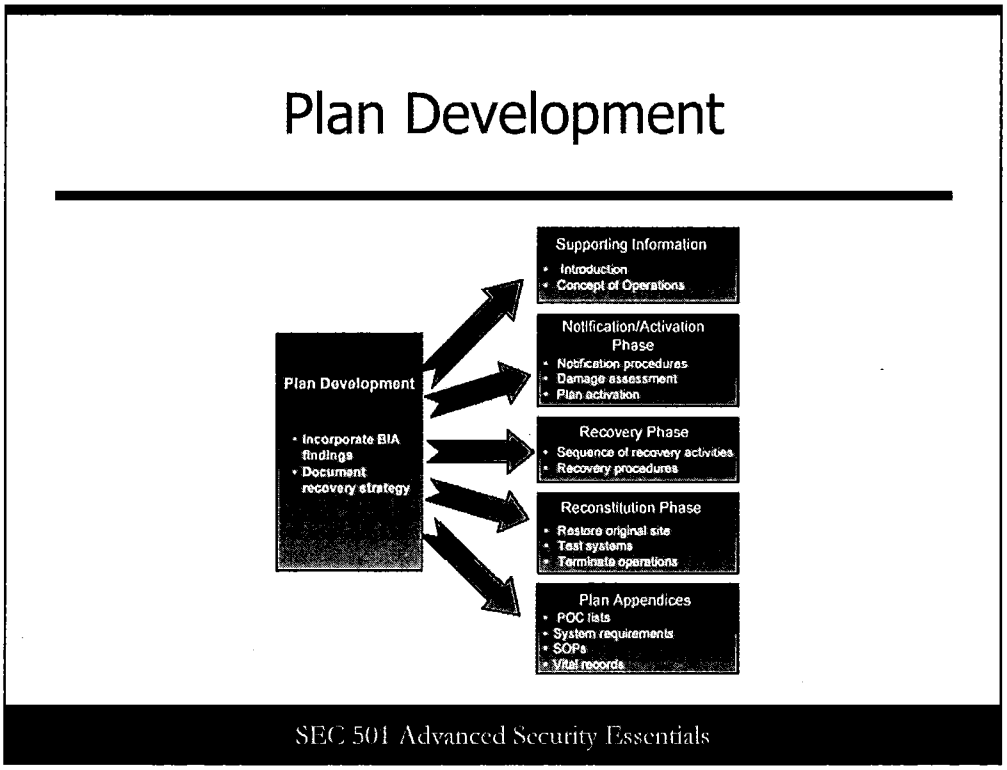

Plan Development
From SP 800-34, "IT contingency plan development is a critical step in the process of implementing a comprehensive contingency planning program. The plan contains detailed roles, responsibilities, teams, and procedures associated with restoring an IT system following a disruption. The contingency plan should document technical capabilities designed to support contingency operations. The contingency plan should be tailored to the organization and its requirements. Plans need to balance detail with flexibility; usually the more detailed the plan is, the less scalable and versatile the approach. The information presented here is meant to be a guide; however, the plan format in this document may be modified as needed to better meet the user's specific system, operational, and organization requirements.

Plans should be formatted to provide quick and clear direction in the event that personnel unfamiliar with the plan or the systems are called on to perform recovery operations. Plans should be clear, concise, and easy to implement in an emergency. Where possible, checklists and step-by-step procedures should be used. A concise and well-formatted plan reduces the likelihood of creating an overly complex or confusing plan."

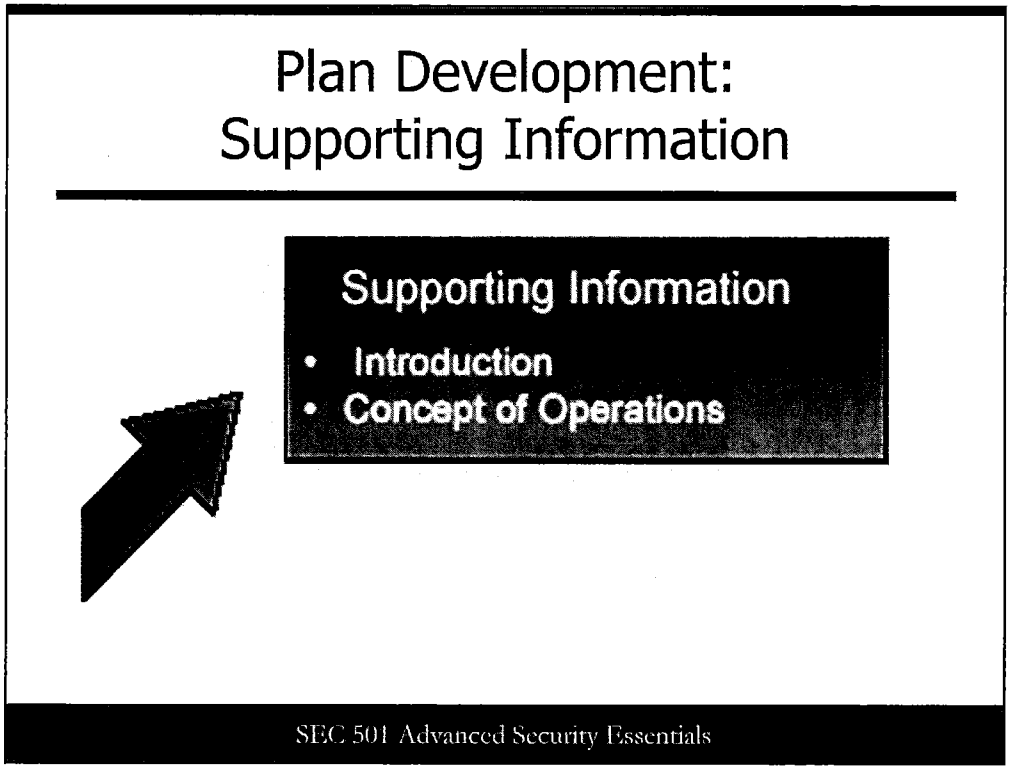

Plan Development: Supporting Information
The planning and development stage includes the purpose, applicability, scope, references and requirements, and record of changes:

- **Purpose**: Establishes the reason for developing the contingency plan and defines the plan objectives.
- **Applicability**: Parties affected by this plan. Any related documents are detailed here.
- **Scope**: Details the assumptions, issues, situations, and conditions discussed in the plan. It also defines the systems that are managed by this plan.
- **References/requirements:** Identifies the legal requirements and governing institutions that impose conditions on the plan.
- **Record of changes**: Contains issues of how change control is managed.
- **Concept of operations**: Provides additional details about the IT system, the contingency planning framework, and response, recovery, and resumption activities.
- **System description**: A general overview of the hardware and software systems.
- **Line of succession**: Identifies personnel to assume that authority in the event others are not available.
- **Responsibilities**: Details the team's functions.

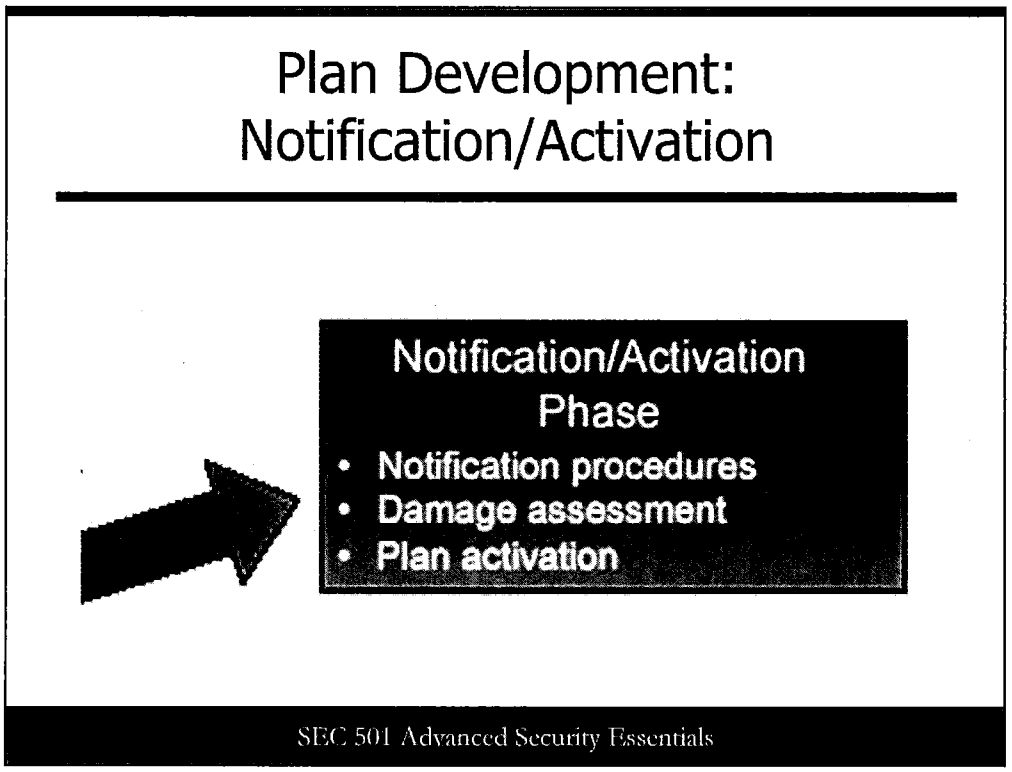

Plan Development: Notification/Activation
The Notification/Activation Phase defines what triggers a declaration of action.

Notification Procedures:
Because we cannot determine the type of notice we will be given for an event, we assume that no notice will be given. At this point, documentation starts with damage assessment. Because some events cause us to use different communication channels, a checklist is appropriate. Each communication channel should be monitored regularly. A phone tree is one of the best tools for little cost. A global organization should consider an automated third-party phone tree. Whatever the technique for notification, it must be clear who to call. That information must be in the plan. The documentation should list team members, any other organizations, and the various points of contact.

Notification information should include the following:
- Nature of the emergency
- Loss of life or injuries

Damage estimates
- Response and recovery details
- Where and when to convene for briefing
- Relocation estimated time period
- Instructions to complete notifications using the call tree

Damage Assessment
This step might be unreliable, depending on the team member's experience. Quick, efficient, and safe are the major concerns in damage assessment. The damage assessment team is the first team that is notified of the incident. Documentation should include: the cause of the emergency or disruption, the

potential for additional disruptions or damage, the area affected by the emergency, the status of physical infrastructure, the inventory and functional status of equipment, the type of damage to equipment or data, the items to be replaced, and estimated time to restore normal services.

Plan Activation

The contingency plan should be activated when the damage assessment triggers have been achieved based upon some or all of the following: the safety of personnel, the extent of damage to system, the criticality of the system to the business objective, and the estimated duration of disruption.

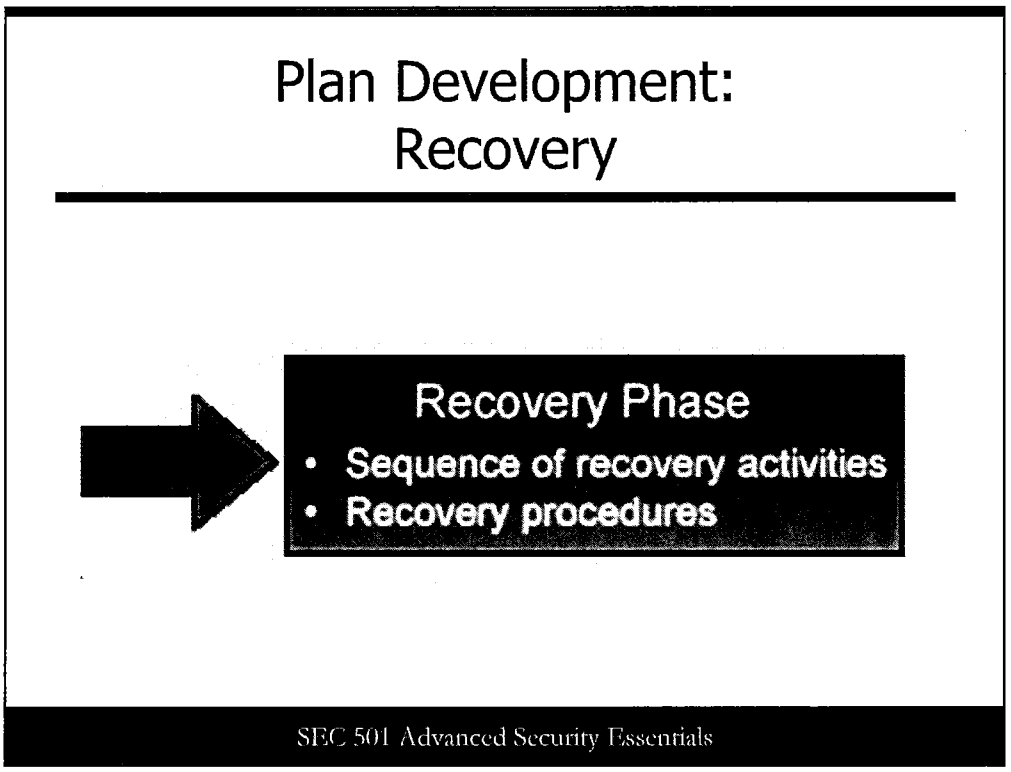

Plan Development: Recovery
The recovery phase has three categories of activities:

1. Execute temporary processing capabilities.
2. Repair or replace.
3. Return to original operational capabilities.

Sequence of Recovery Activities
This order of activities follows business requirements of critical systems to return to production. Any critical, prerequisite supplies should be on hand to achieve the goal. For example, if you were to choose to use a manual system at an alternate location, you would need the supplies to be delivered first and the personnel transported second. It would do you no good to have staff without supplies.

Recovery Procedures
Recovery procedures are the most time-consuming part of the plan to maintain. This section should be tested heavily. Recovery procedures are the step-by-step of how to return to an operational state. If you were a help desk support person, what kind of instructions does the help desk need to support 100 servers in running order with the proper software configured for the business? What would be the order of restore from backup? What configurations would your telephone operator need for the phones? How do you test that the system is working properly? Checklists work well in this area.

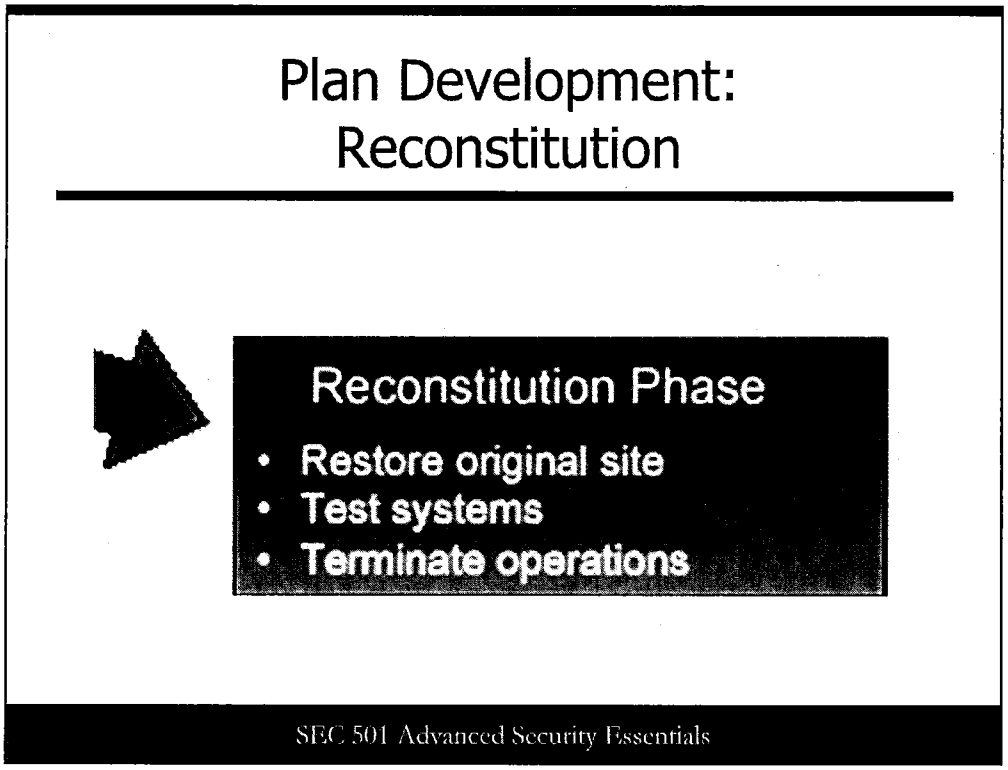

Plan Development: Reconstitution

The reconstitution phase consist of return to permanent facility, testing systems, and shutting down temporary facility:

- **Return to permanent facility**: If the original facility was destroyed, you might be resuming operations at a new location. Hopefully you can resume operations at the original site.
- **Test systems**: Resuming at the permanent facility requires that you are able to cease BCP and go to "business as usual" mode. All systems must be verified, certified, or accredited depending on the regulatory bodies.
- **Shutdown temporary facility**: It might seem simple, but the infrastructure set in place might be very costly to maintain. It will take time to break down the systems. You must also make assurances that the residual data at that site is disposed of properly.

If possible, each section of the reconstitution phase should have a team at each location.

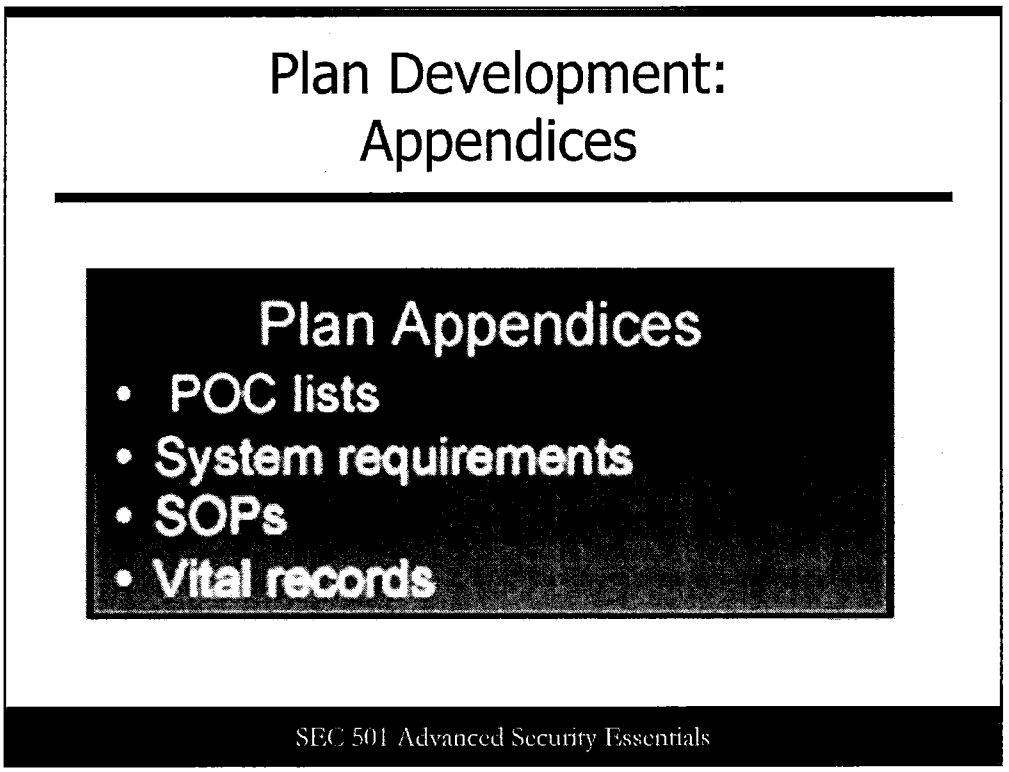

Plan Development: Appendices
Each section of the plan should be complete by itself, but it might not contain the step-by-steps of server restores or the phone list of every customer to call. This is the realm of the appendices. Follow the BIA to guide you regarding the requirements for your organization. Here is a list of common items:

- Contact information
- Vendor contact information
- Standard operating procedures
- Checklists for system recovery or processes
- Detailed equipment and system requirements lists of resources
- Vendor service level agreements
- Reciprocal agreements
- Alternate site information

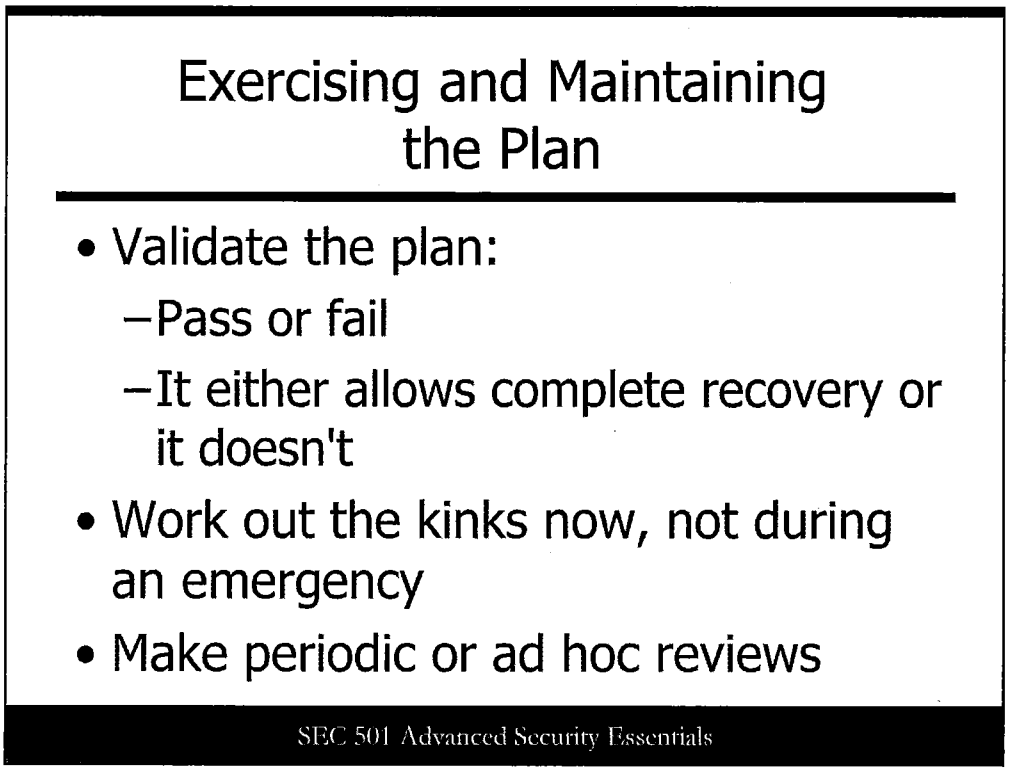

Exercising and Maintaining the Plan
Confidence in the company's continuity plan can only be achieved through testing. Leaving the business continuity plan on a shelf somewhere, forlorn and forgotten, immediately vaporizes any value the plan might have initially possessed.

Testing verifies the accuracy of the recovery procedures and highlights any discrepancies or areas that were unintentionally overlooked during the plan's creation. Also, testing familiarizes personnel with the plan's objectives and provides the necessary preparation for quick, decisive response to disruption. Remember, "Plan the dive. Dive the plan."

In short, testing the business continuity plan provides the following:

- **Consistency:** Testing ensures that there are few, if any, gaps between the plan and the organization's current characteristics. The organization's hierarchy, infrastructure (network), business processes (and maximum allowable downtimes), staffing levels, and vendor and outsourcing relationships should all be current and accurate within the document.
- **Validity:** Testing ensures that the plan is still valid and that the assumptions contained therein are logical and practical. The question to be answered is, "Does the plan truly enable recovery?"

Testing and maintenance of the plan can happen periodically or on an as-needed basis. Periodic review consists of testing and reviewing the plan at specific times within the calendar year, either quarterly, bi-annually, or annually. Ad-hoc review consists of testing the plan as needed or as warranted by executive decision-makers. Opinions differ about which method is better, but it is incumbent upon the company to make the investment in the business continuity plan worthwhile. The company is free to combine the two methods for optimum coverage.

In all cases, any time the company adds new business processes, upgrades, or alters the infrastructure or makes any other modifications, you should review and update the business continuity plan. Preferably, test the plan within a reasonable amount of time to ensure that you can recover the new business processes or infrastructure. Remember that you should conduct BIAs to ensure that the company is responding to the right weaknesses.

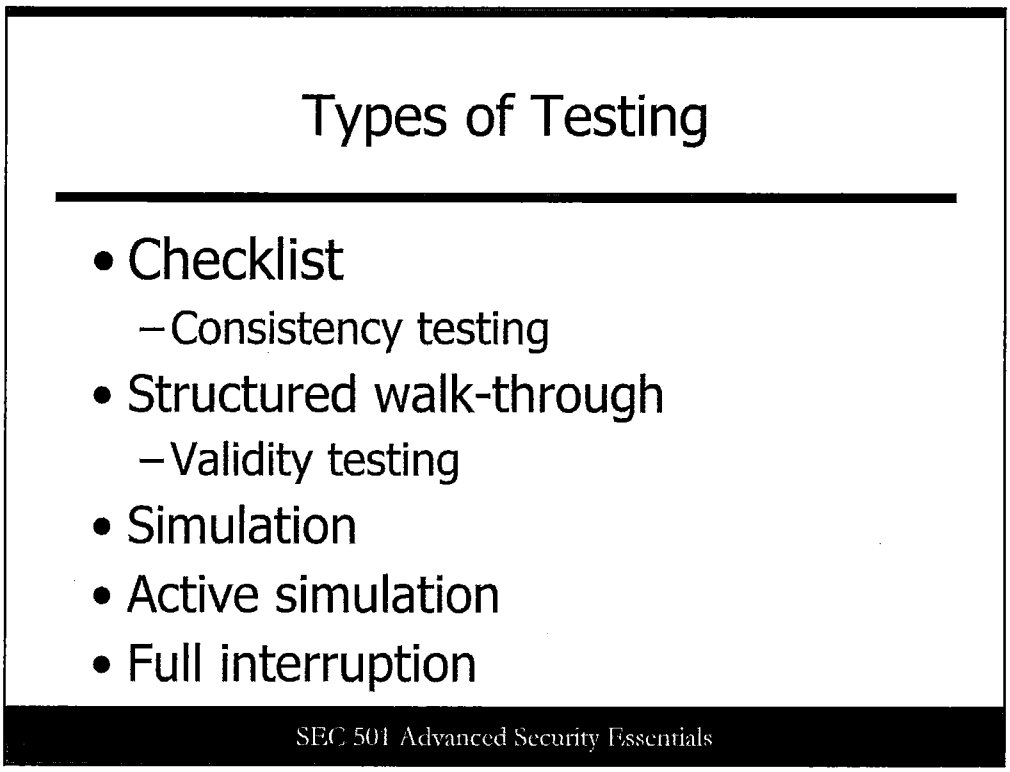

Types of Testing

Checklist testing, also known as *consistency testing*, simply involves reviewing the business continuity plan to ensure that it addresses all critical areas of the enterprise and that the procedures to recover those areas are accurate and consistent. Checklist testing is the least expensive of all the testing methods; however, it is also the least valuable because it does not depict the company's responsiveness to disruption. Checklist testing is for sanity checking and should not be considered a viable testing method in and of itself.

Structured walk-through testing, also known as *validity testing,* ensures that the plan contains no errors, erroneous assumptions, or blind spots, and that it accurately reflects the company's ability to recover from disruption. Team members and other individuals who are responsible for recovery meet and walk through the plan step-by-step.

Simulation testing involves a mock-up of an actual emergency where team members respond as if an emergency is occurring. This test is really a structured walk-through test on steroids or at least some type of amphetamine. You may recover locations (including the emergency operations center and the alternate sites) and enable communications links while team members execute the recovery steps in a walk-through manner. You do not actually perform recovery actions (restore backups). This testing method can be expensive for a company, but in comparison to the following two testing methods, it can prove invaluable for the dollars spent. A simulation test is a satisfactory testing method because it gives the enterprise fairly good insight into its recovery responsiveness.

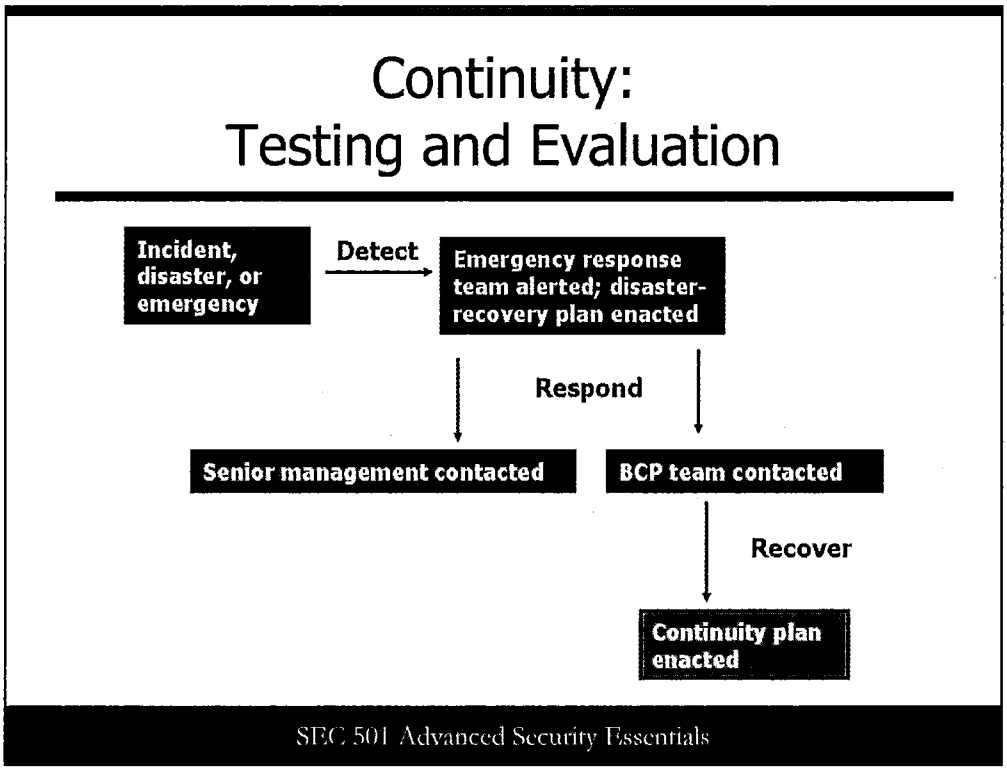

Continuity: Testing and Evaluation
Again, a diagram offers a simplistic view of the response mechanism for a corporation; however, it does give the reader a good idea of the intent of the business continuity process before discussing BCP in detail.

When a disruption occurs, it is detected and the emergency response team (ERT) is contacted in some predetermined manner (documented in the BCP of course).

The emergency response team assesses the situation and determines whether an emergency exists and whether BCP/DRP Activities are required. If so, ERT contacts both senior management and the business continuity team members and apprises them of the situation.

In collaboration with the business continuity team members, the emergency response team determines the affected business functions and enacts relevant sections of the business continuity plan.

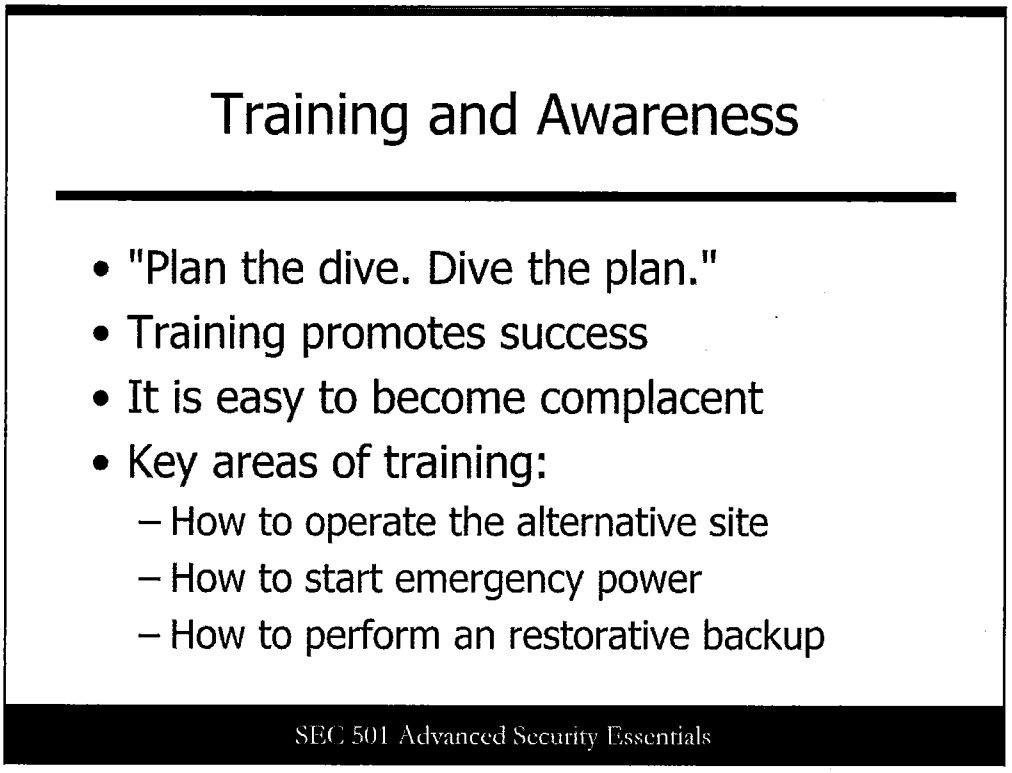

Training and Awareness

A crucial aspect of the business continuity plan, training is often minimized because it takes employees away from their primary responsibilities.

An organization should and must train all staff in the recovery process. Recovery procedures might be significantly different from those pertaining to normal operations, and team members should feel confident in their ability to recover the company. Confidence is the driving factor in the plan's success, especially when employees are under duress. Training might include the following:

- How to operate the alternate site
- How to start emergency power
- How to perform an restorative backup

"Plan the dive. Dive the plan," is a SCUBA diver's saying about conducting a safe dive. A lot of effort and attention goes into a safe dive, but above all, the diver must execute the agreed-upon plan. It requires expertise with equipment, environment, and the diver's own personal limitations.

As a company, you are giving your team members the same expertise by providing training. Training also gives the team valuable feedback on the readiness and preparation of emergency response and recovery team members and the overall recovery process itself.

Most importantly however, a corporation cannot assume that all members of the response and recovery teams will be available or even alive. Training all employees (or a large subsection) increases the likelihood that individuals will be available when loss of life has drastically reduced the number of experienced individuals.

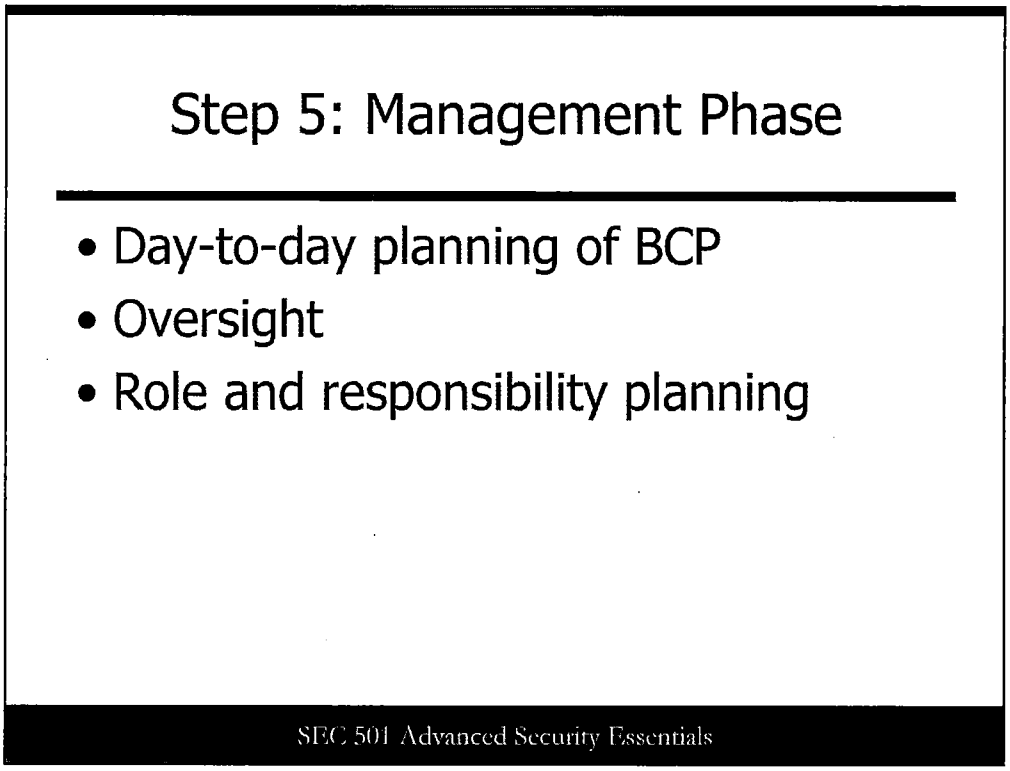

Management Phase
Day-to-day management of the continuity plan: Required to evaluate day to day changes in operation and company function to ensure the current BCP and DRP are applicable and updated.

Oversight: In control of all aspects of the BCP/DRP and have stock in the implementation if the need arises.

Roles and responsibilities: Delegate specific roles and responsibilities for key tasks to the appropriate personnel.

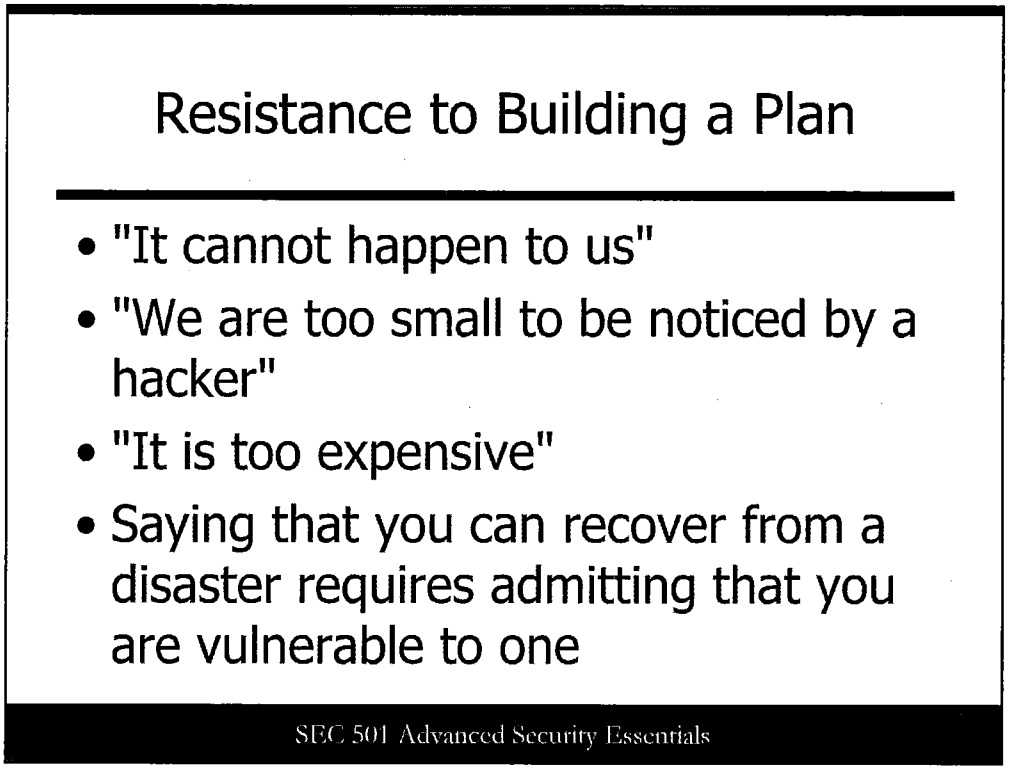

Resistance to Building a Plan

Most people believe that disaster will not happen to them; it will happen to someone else instead. This belief is normal, but it represents perhaps the greatest resistance to building a business continuity plan because, at some fundamental level, we do not want to bet against ourselves. By thinking that bad things *might* happen, we believe we *make* bad things happen. By not thinking about bad things, we preclude Fate's attention. Of course, reality is immune to such superstitions.

Although probabilities do much to assuage our fear regarding the frequency of tragic events, even the most in-depth cognitive analysis never quite overcomes our para-rational uneasiness, no matter how conclusive or definitive statistics might be. Again, it harkens back to not betting against ourselves. However, BCP is not about betting *against* ourselves; it is about betting *for* ourselves.

Unfortunately, uneasiness is a far louder voice than reason or rational thought alone, stating obtuse authoritative directives such as:

- It cannot happen to us.
- We are too small to be noticed by a hacker.
- It is too expensive to implement a solution, and besides, it probably won't happen anyway.
- It has never happened before.
- It will never happen here.

As if obtuse statements invoke a magical talisman immunizing the corporation from harm or hazard. The greatest resistance to building a continuity plan is simply the reluctance of management to see potential disaster. Ironically, many in corporate leadership positions perceive the ability to recover from a disaster as a confession that the business *is vulnerable to a disaster in the first place*, a wholly unacceptable admission indeed.

Top BCP/DRP Planning Mistakes (1)
A number of other mistakes are commonly, almost predictably, made in contingency planning:

- **Lack of BCP testing/over:** Reliance on BCP-Many companies believe that just having the BCP is enough. Without adequate updating and testing, the document is just a lifeless draft. Organizations that test their BCP consistently find areas that need improvement, and they often find critical flaws. The time to discover these is before a real disruption. "Practice makes perfect." If you need less expensive testing and do not want to perform it in house, you can use off-site test facilities. Try simulating a disaster, as in a business simulation game. Pretend that something has happened and that certain resources are no longer available, and have your personnel (who are assumed available) walk through the plan.
- **Too limited in scope:** An incomplete BCP does not address all of an organization's needs for recovery. The BCP plan should cover organizational processes and process dependencies, systems recovery, and the replacement of key personnel, if needed. The organization should continue to function throughout a disruption and beyond.
- **Lack of clear authority and process:** In times of disaster, only partial staff might be available, and their level of empowerment might need to be significantly higher than normal. Definitions for when and under what circumstances change in empowerment and processes begin and end need to be clear and unambiguous.
- **Lack of prioritization:** There is a need to prioritize the key business processes. The risk is to prioritize less-than-critical processes rather than those that are crucial for business survival. This is a time for thoughtful evaluation and decisions.
- **Lack of plan updates:** The BCP should be updated periodically, especially when there are significant system or business process or personnel changes.
- **Lack of plan ownership:** Someone with the power to lead, influence, prioritize, and organize the BCP is instrumental to the success of the program. This is true during planning and during the plan's execution.

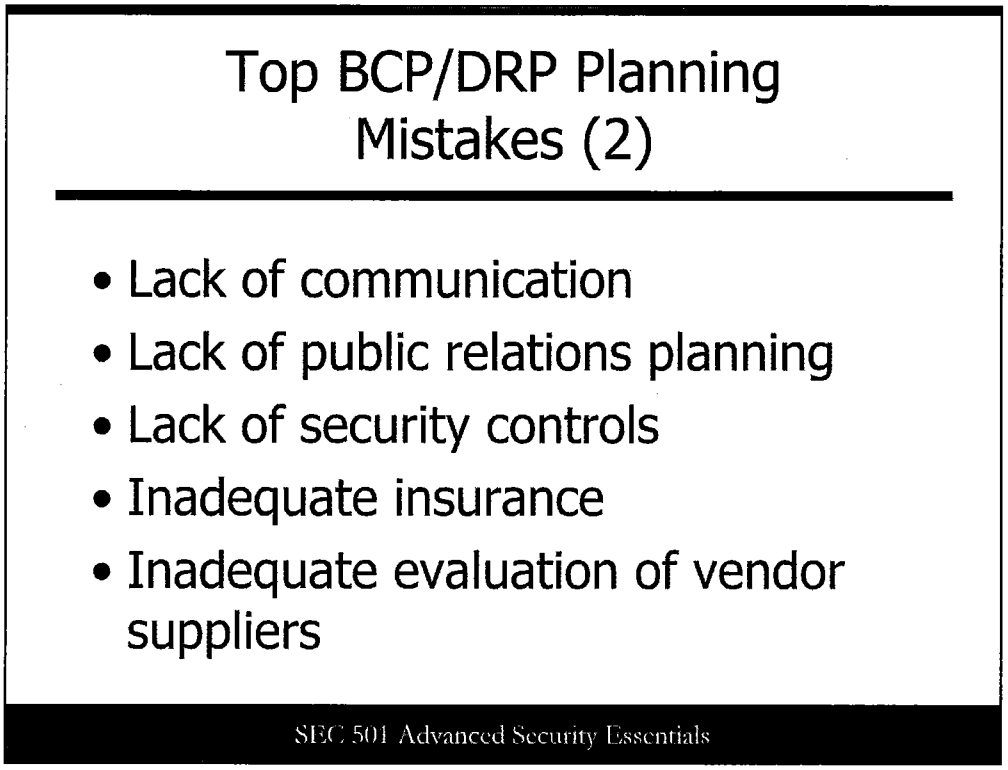

Top BCP/DRP Planning Mistakes (2)

- **Lack of communications:** There is a need for clear and precise communication with all affected stakeholders of the organization - potentially employees, contract employees, vendors, business partners, customers, and shareholders. (This relates to public relations planning, which follows.)
- **Lack of public relations planning:** Organizations often fail to consider public and investor relations to limit the perceived disaster impact. This can literally make or break the organization. Remember the Tylenol tampering scare several years ago and how the strong PR from that company turned a disaster into a marketing opportunity?
- **Lack of security controls:** During the recovery process, security controls are sometimes disregarded; this results in a greater risk of exposure. Security controls might need to be altered and loosened during recovery. However, this should be a matter of conscious decision and empowerment that are built into the plan. Strict adherence to the security controls incorporated into the plan during its execution.
- **Inadequate insurance:** Some organizations lack adequate insurance coverage and fail to support the filing of insurance claims. These inadequacies result in delayed or reduced settlements. The plan might lack appropriate processes for capturing losses and recovery costs, without which the organization might realize a loss greater than otherwise necessary.
- **Inadequate evaluation of vendor suppliers:** Many companies poorly evaluate recovery products (hot site, cold site, and planning software), relying on vendor-supplied information. This often leads to a solution that might not adequately address a company's needs.

Use these examples of common mistakes as a checklist to review your organization's contingency planning - the documentation, testing, integration with organizational processes and personnel, and so on. There are companies and consultants that specialize in designing and implementing business continuity plans. As appropriate for your organization, you may wish to consider their services.

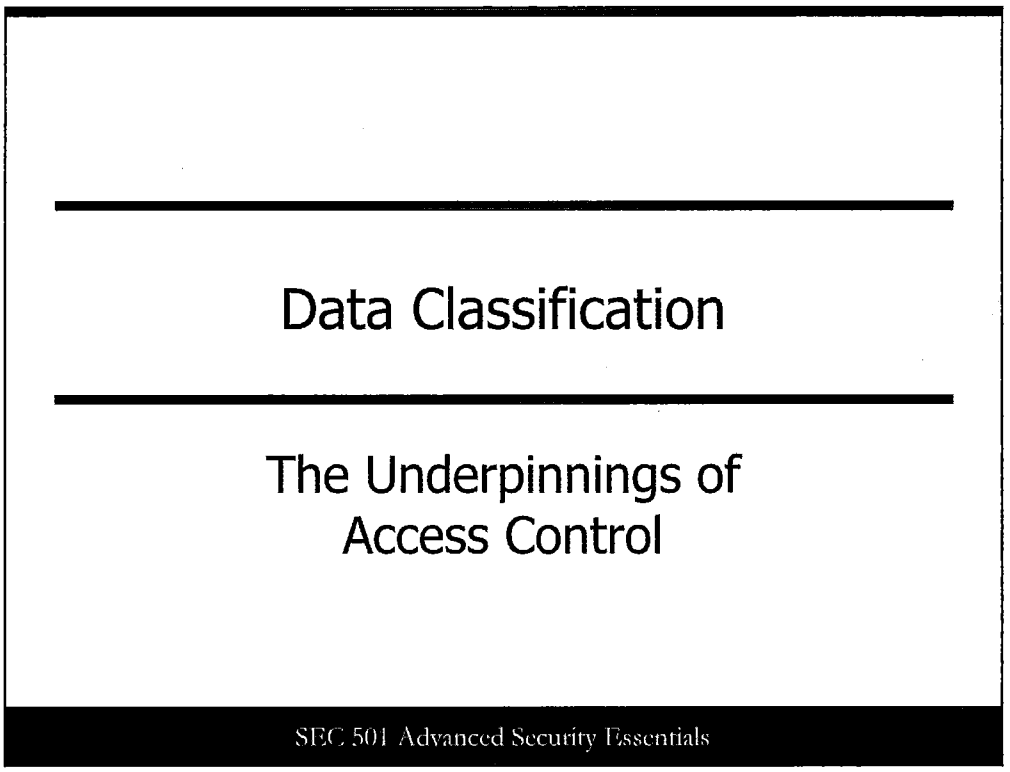

This section will address the underpinnings that make access control to electronic information work. Without an identifiable process to mark data appropriately, access controls cannot be subjective and, therefore, will always fall under manual rules and guidelines of discretionary or more precise, objective guidelines. This may not be appropriate for all data and, therefore, the process of data classification.

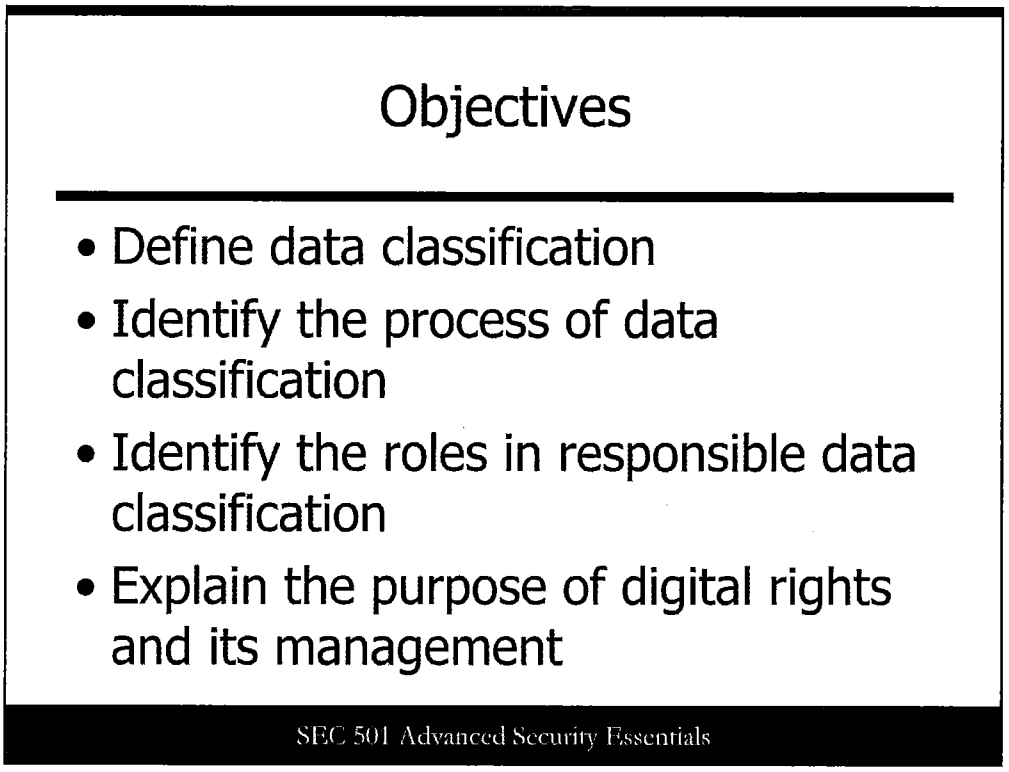

This page intentionally left blank.

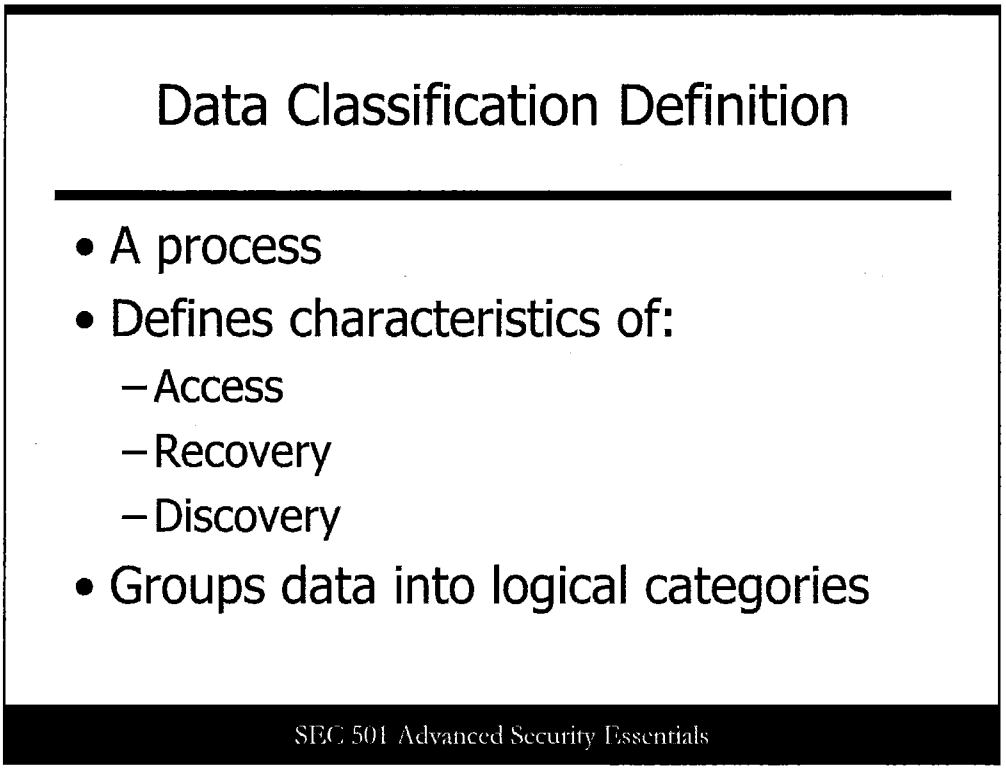

Security is all about managing, mitigating and reducing risk to an organization's critical data. What differentiates one organization from another is the critical intellectual property. It is critical that an organization identifies and classifies information so proper controls and measures should be put in place.

Data Classification Process

- Participants:
  - Identify stakeholders
- Assets:
  - Identify all storage areas
  - Identify all data to be stored
- Assign roles and responsibilities
- Goals:
  - Identify and provide appropriate security levels

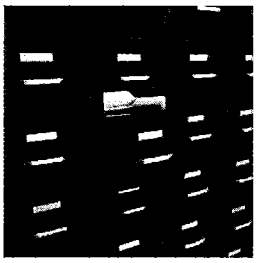

SEC 501 Advanced Security Essentials

In order for a classification process to work all key stakeholders must be actively engaged in the process. Data owners have to engage on a regular basis to make sure that only the people who need access to critical information have it. The key stakeholders can help identify the critical intellectual property and perform data discovery to identify where the information is stored. The ultimate goal of data classification is to make sure that all information is properly protected at the correct level.

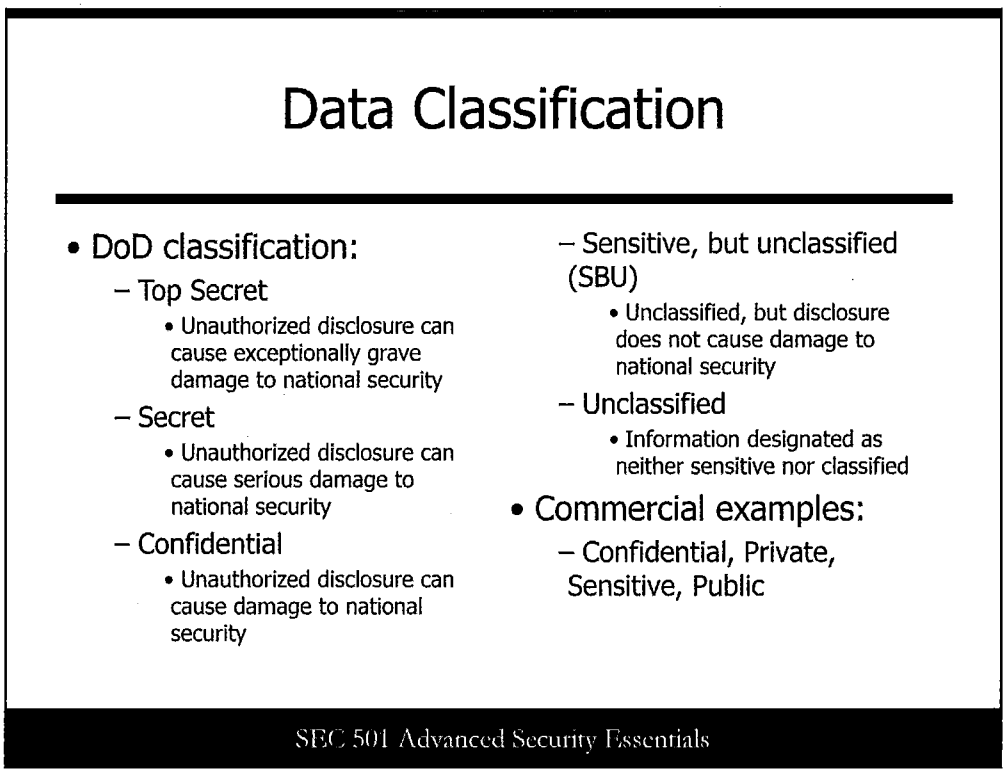

Governments and their militaries, such as the U.S. Department of Defense (DoD), started the phenomenon of labeling data to apply higher levels of protection to data that was so sensitive that if it were leaked it could harm their country's national security. Subsequently, this has become commonplace in the corporate world as well. A quick listing of the DoD and federal levels follows:

- **Top Secret:** The highest levels of protection are given to this data; it is critical to protect.
- **Secret:** This data is important, and its release could harm national security.
- **Confidential:** This is important, and it could be detrimental to national security if released.
- **Sensitive But Unclassified (SBU):** This generally is information that is sensitive and should not be released (like SSNs).
- **Unclassified:** They prefer to keep it from being released but the nation would not be harmed if it were.

Data classification is key to asset management. Assigning risk, deterring threat, and applying the correct protection is key to data classification. All businesses need to assign a value to the data. Based on that value, you can then apply correct protection means. The protection measures may by dictated by policy or by law.

Due to cost it would be impossible to assign maximum protection to all data, Therefore you need to understand the different classes of data according to the assessed risk and value. This will give you the ability to apply maximum resources to critical data that would ensure continuity of business operations.

Data Classification Criteria

- Value
- Age
- Legal requirements
- Reputation risk
- Ability to perform

SEC 501 Advanced Security Essentials

How is data classified?

- Value:
  - What is the information worth to the company?
  - What if it is lost or compromised?
- Age:
  - How current is the information?
  - Is real-time information more important to your organization than information received last week?
- Legal Requirement
  - Examples include medical records, case files, and personnel files.
  - HIPAA and FERPA
- Reputation risk
  - What is the dollar cost to business reputation if data is lost?
- Ability to perform
  - If there is data is loss, what is the down time to business or system?

These are just a few examples of criteria that may be used. Some other examples: Victim' s Welfare, Financial Remediation, Data usage, and service agreements.

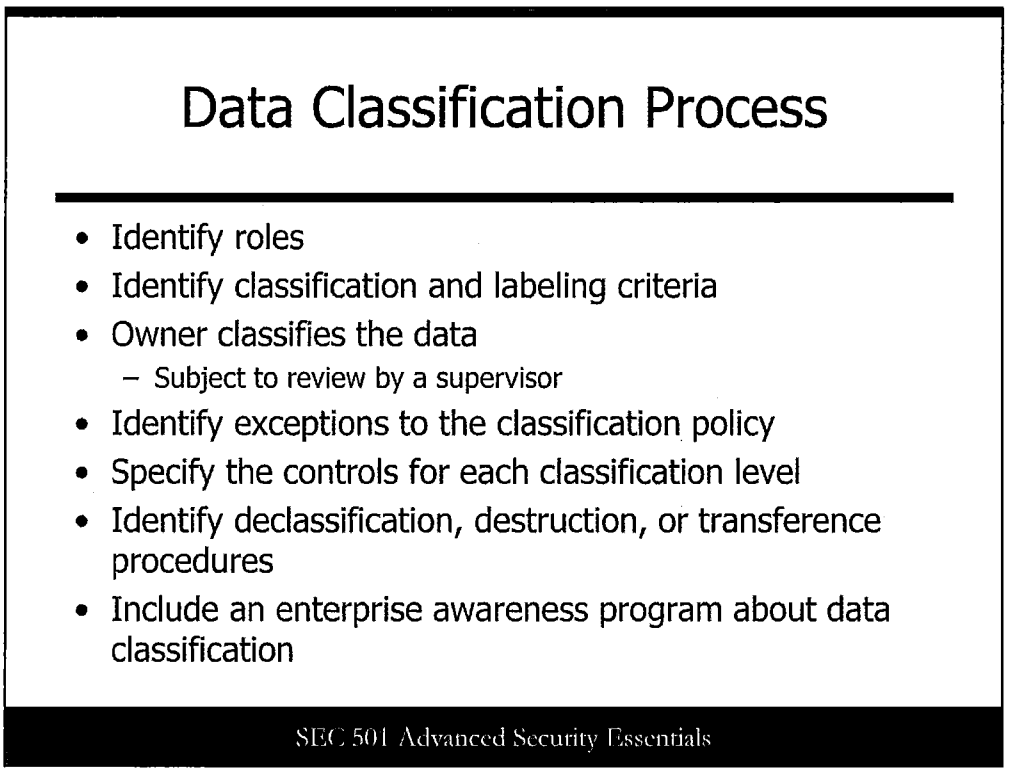

To allow for a complete program, the steps in classifying data, identifying its controls, stating its disposition, and training needs to be addressed and documented.

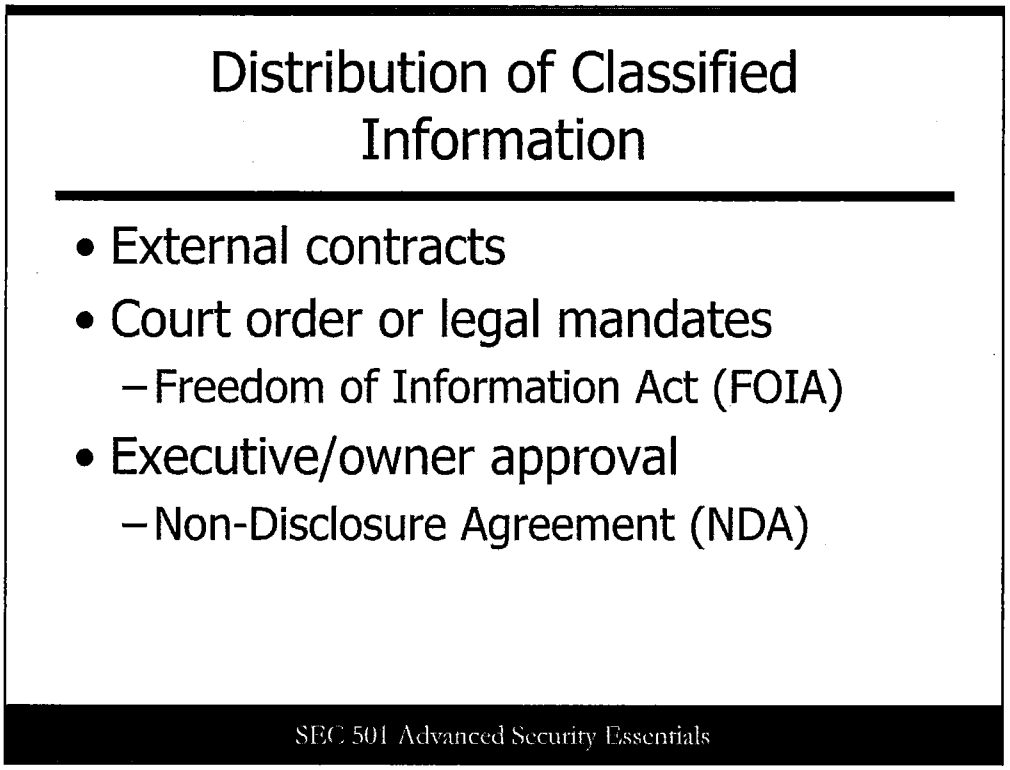

External contracts with either customers or vendors may require the sharing of classified data and is usually accompanied by a Non-Disclosure Agreement (NDA).

Court orders or other legal mandates, such as FOIA requests (Freedom of Information Act), can require release of information that would otherwise remain protected.

If a court says you need to present specific documents, you need to present those specific documents, regardless of their classification.

Executive/owner approval may be done for a number of reasons, including but not limited to corporate strategy and tactics such as mergers and partnerships.

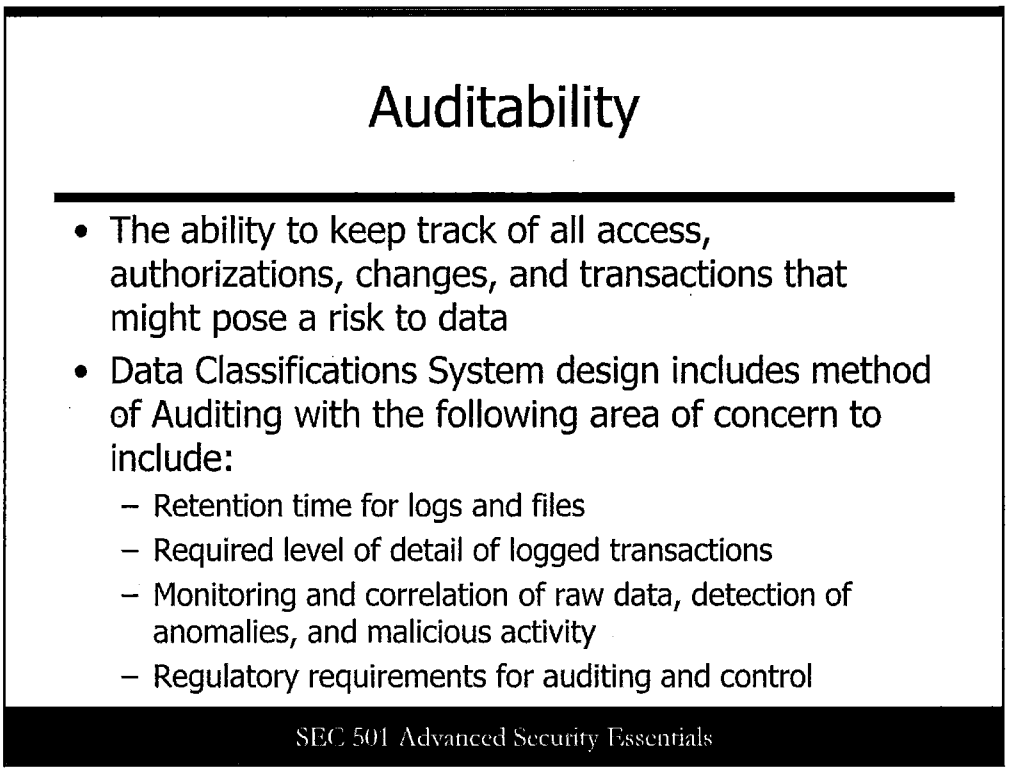

Auditability is the ability to keep track of all access, authorizations, changes, and transactions that might pose a risk to data.

Areas of concern include:

- Retention time for logs and files
- Required level of detail of logged transactions
- Monitoring and correlation of raw data, detection of anomalies, and malicious activity
- Regulatory requirements for auditing and control

Data Classification Roles

- Data owner
- Custodian
- Application owner
- Data manager
- Data user
- Auditors

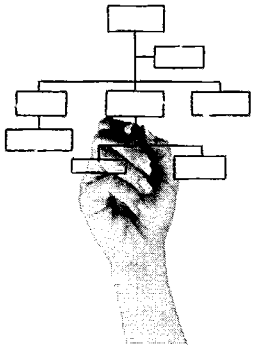

SEC 501 Advanced Security Essentials

The following are key data classification roles:

- Data owner
- Custodian
- Application owner
- Data Manager
- Data User
- Auditors

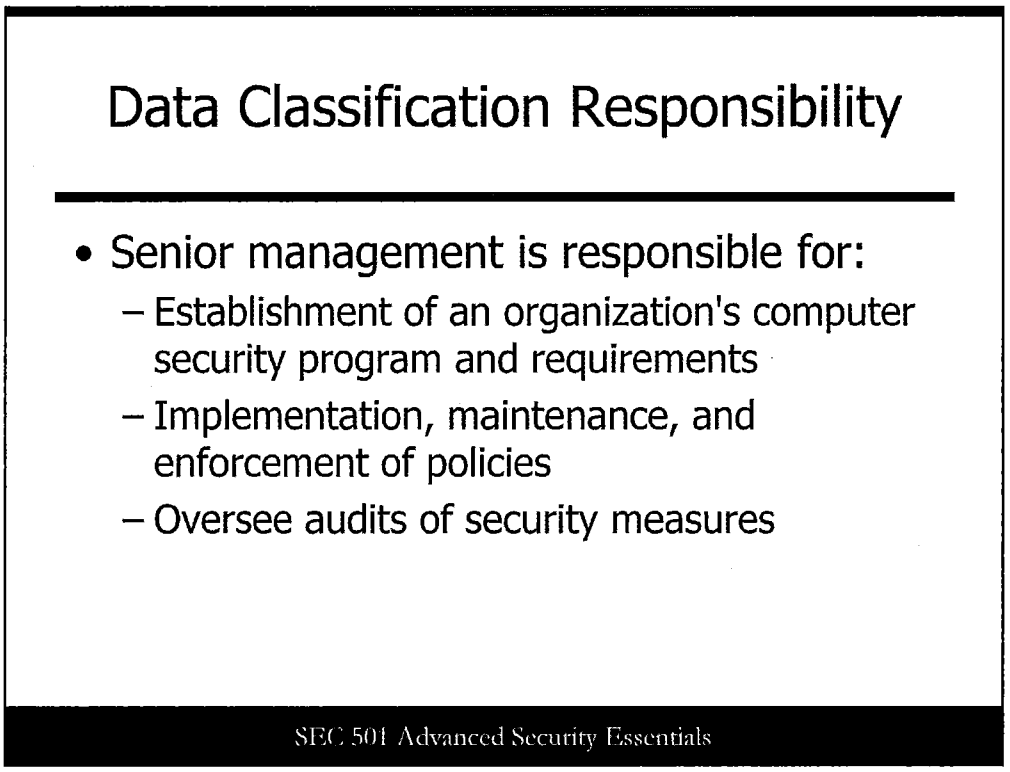

The person responsible in an organization is not necessarily the one who implements it or manages it; they are the ones who mandate the requirements and monitor its use.

Senior management is responsible for implementing an effective and appropriate data classification program. They are responsible for ensuring funding and manpower to implement, maintain, and enforce the program policy when needed. Management oversees audits of security measures in support of all programs.

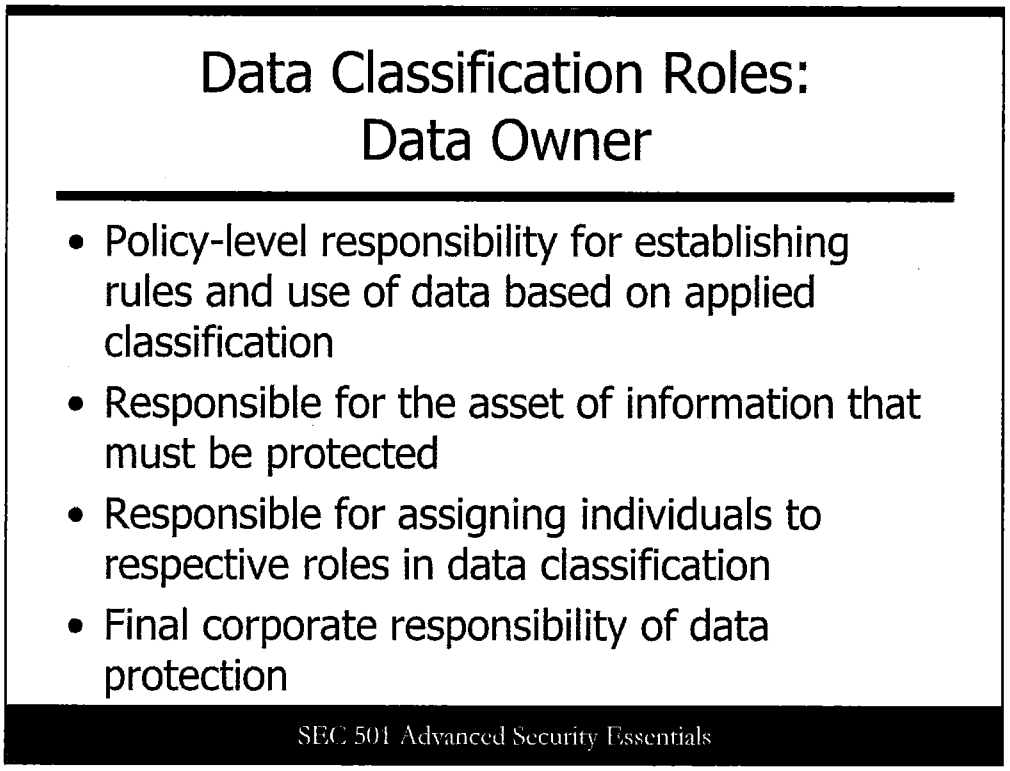

The Data Owner has policy-level responsibility for establishing rules and use of data based on applied classification.

The head of any agency is the overall Data Owner and is still responsible for:

- Assigning and approving final classifications
- Ensuring the protection of agency data
- Approving appropriate use of agency data

Employees of an agency may be delegated some portion of this responsibility on behalf of the agency.

The Data Owner assigns individuals to their respective roles in data classification.

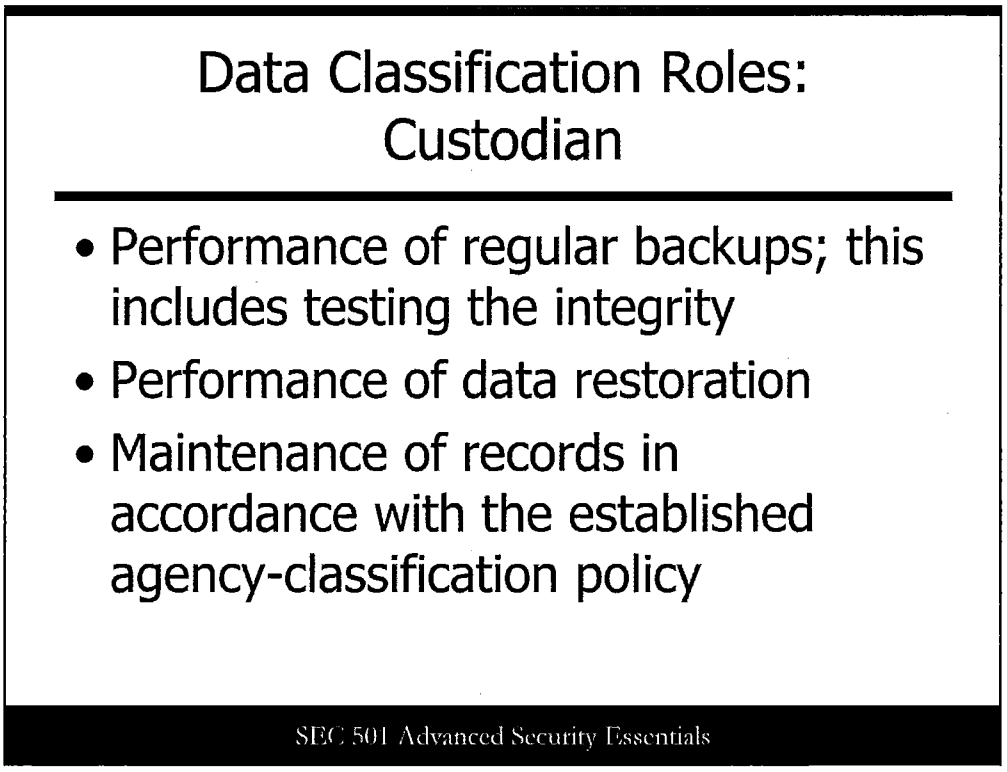

The custodian is responsible for the maintenance of the data. Custodians are responsible for backups and restoration of data, and may be required to administer a classification scheme. The custodian performs the management of data as indicated by data owners.

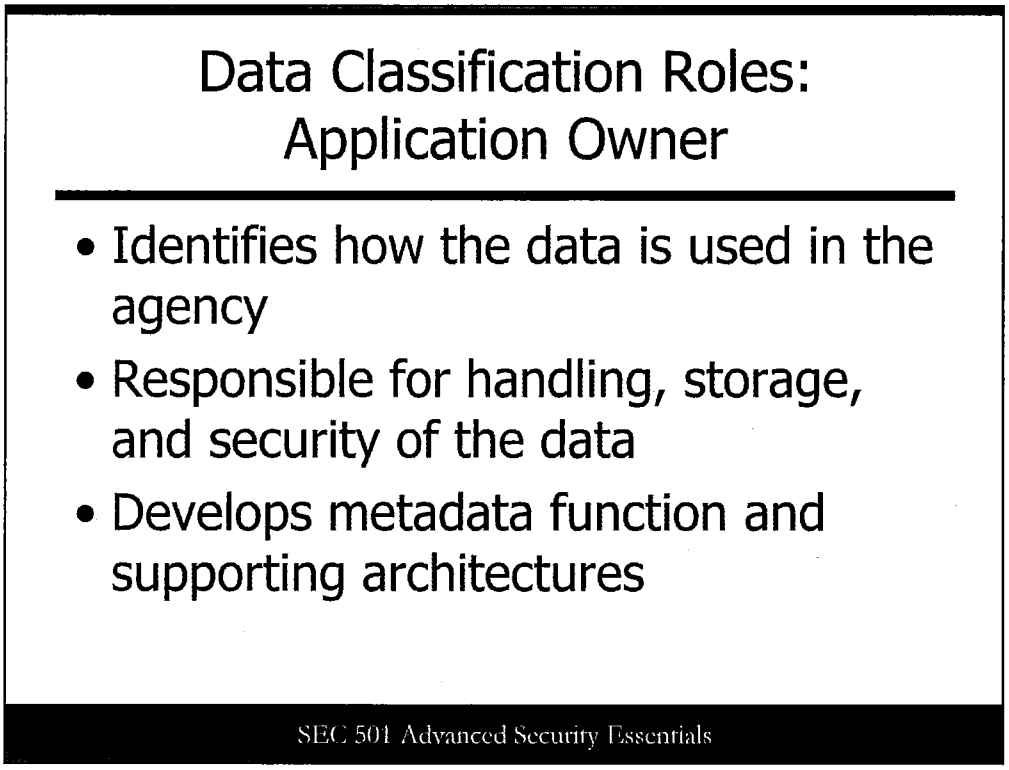

Application owners (developers) design and address the functions that handle data within an agency.

Application owners, due to the programming nature, develop meta-data functions and support architectures to manage and implement data classification. Data classification implementation must remain simple to allow end-user understanding and easy implementation to allow for an effective program.

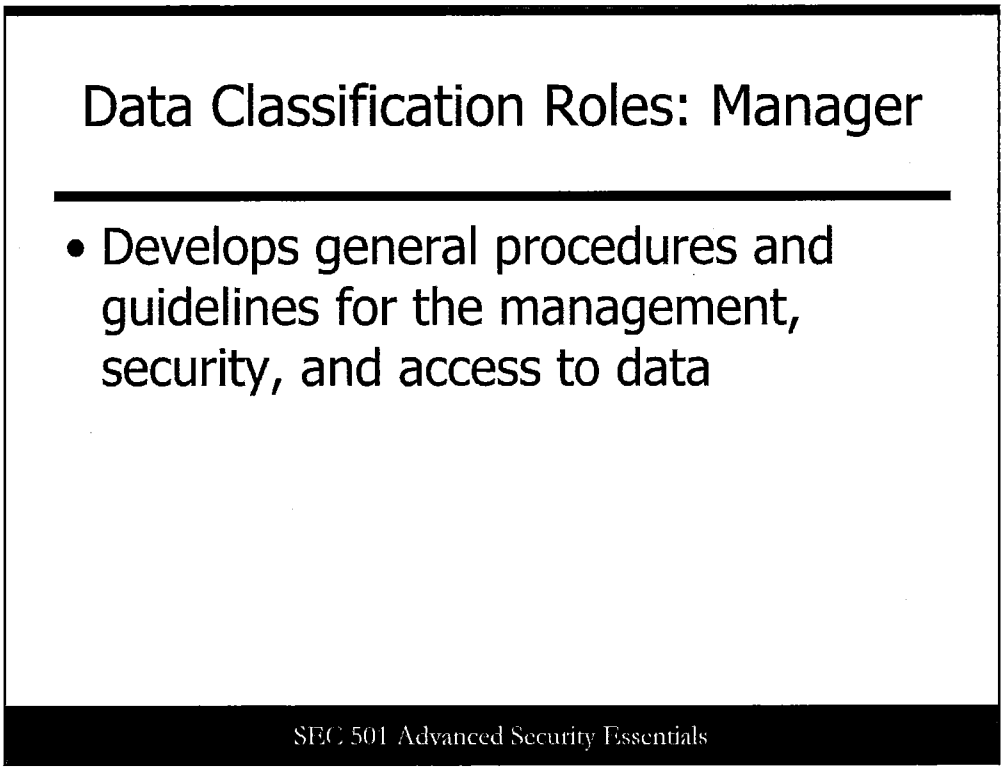

The agency senior management identified the overall, broad scheme security classification guidelines. The Data Manager further develops the procedures and guidelines to ensure proper management, security, and access to data is maintained. The Data Manager's roles fall between data owners and custodians and ensure smooth operation and access to data.

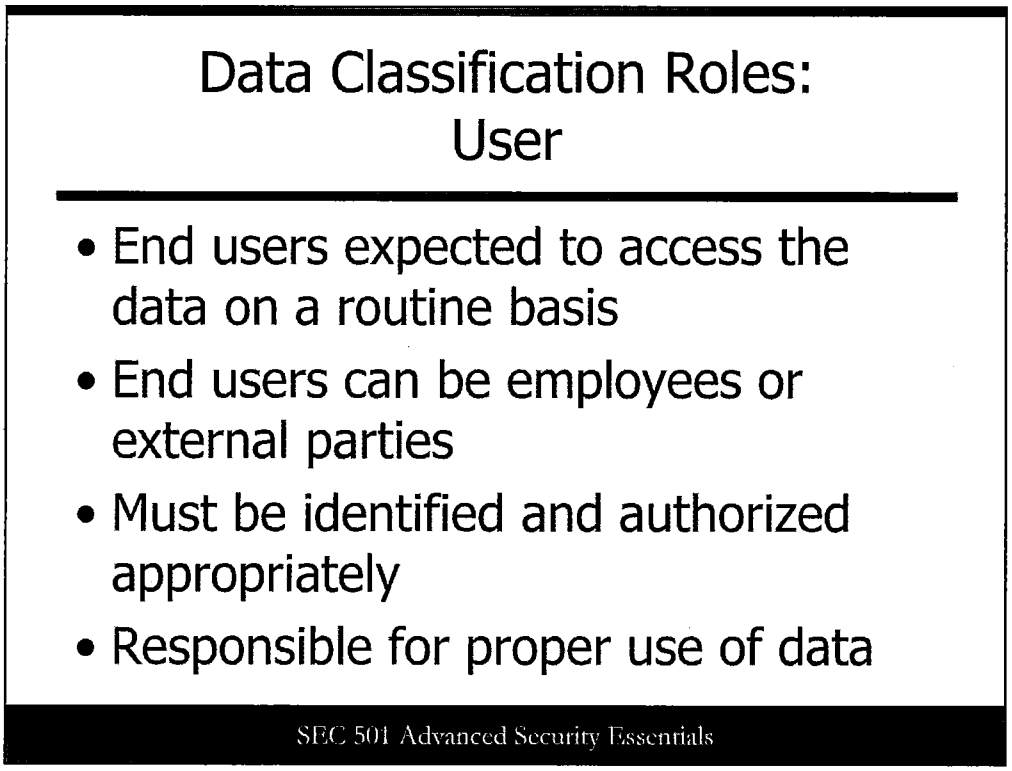

Users are generally first and last access points to data and, therefore, are responsible for its proper use. Users are responsible for ensuring adequate protection and appropriate classifications are afforded the data. Reading, modification, and deletion of data may fall under their responsibility based on the type and purpose of the data.

Security awareness training involving appropriate handling and classification of data should be provided to users on a regular basis. Strict adherence to regulations, policies, procedures, and guidelines should be enforced.

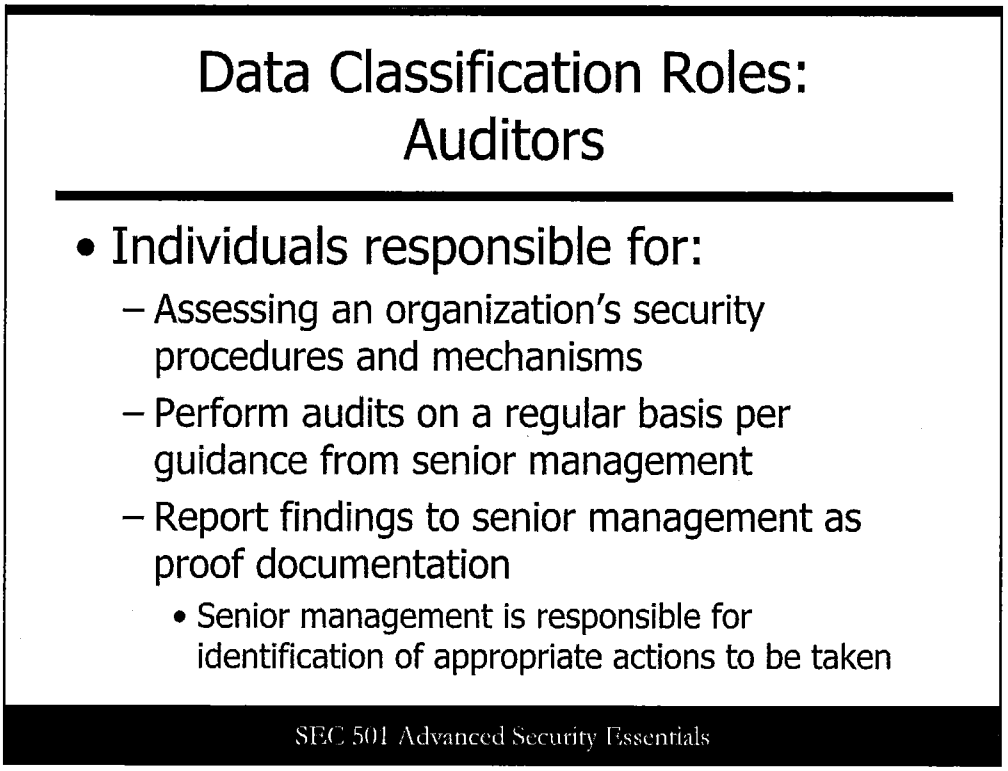

Auditors are the patrolmen and insurance for an organization. Their functions allow them to test, analyze, and evaluate the security mechanisms in place at an agency and report on them. They receive their guidance from senior management and are responsible only for the identification, and sometimes the recommendation of security measures.

Examples:

- The health care industry is governed by HIPAA and states that audits must be performed yearly.
- The Sarbanes-Oxley Act and the rules issued by the US Securities and Exchange Commission (SEC) require auditors to maintain, for seven years after the conclusion of the audit, all "records relevant to the audit or review."

Regardless of the industry, senior management should document and approve the audit process.

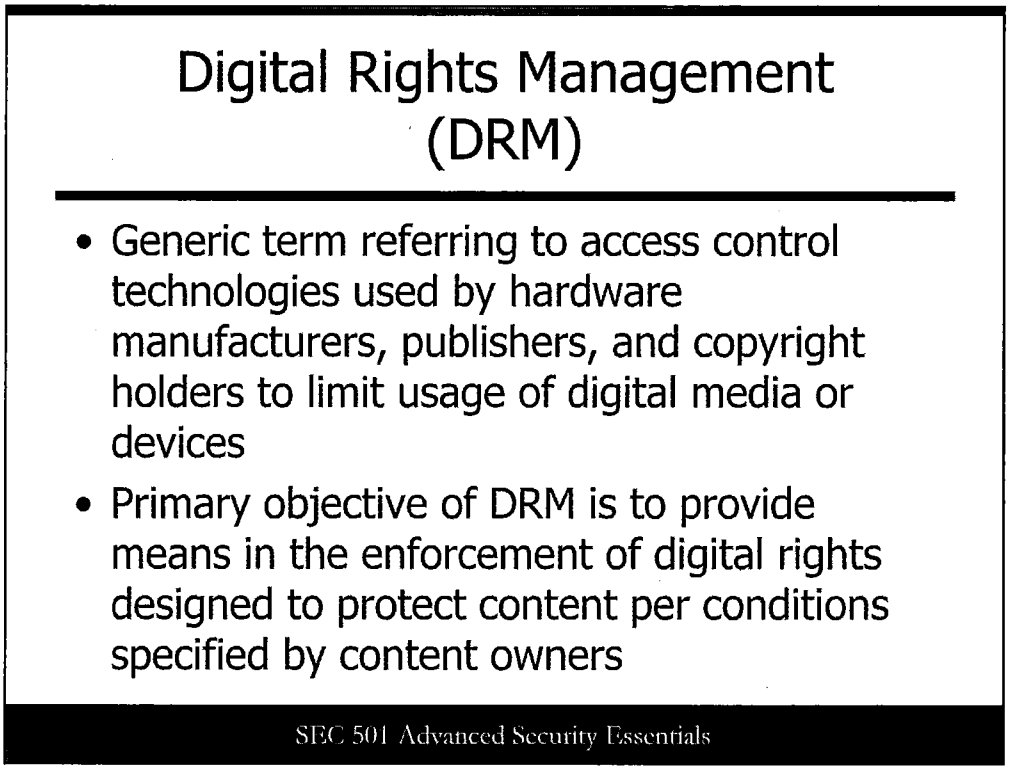

The primary objective of DRM is to build a DRE (digital right enforcement) environment. These environments may entail access control, data tracking, or simply data marking methods to ensure original ownership of data.

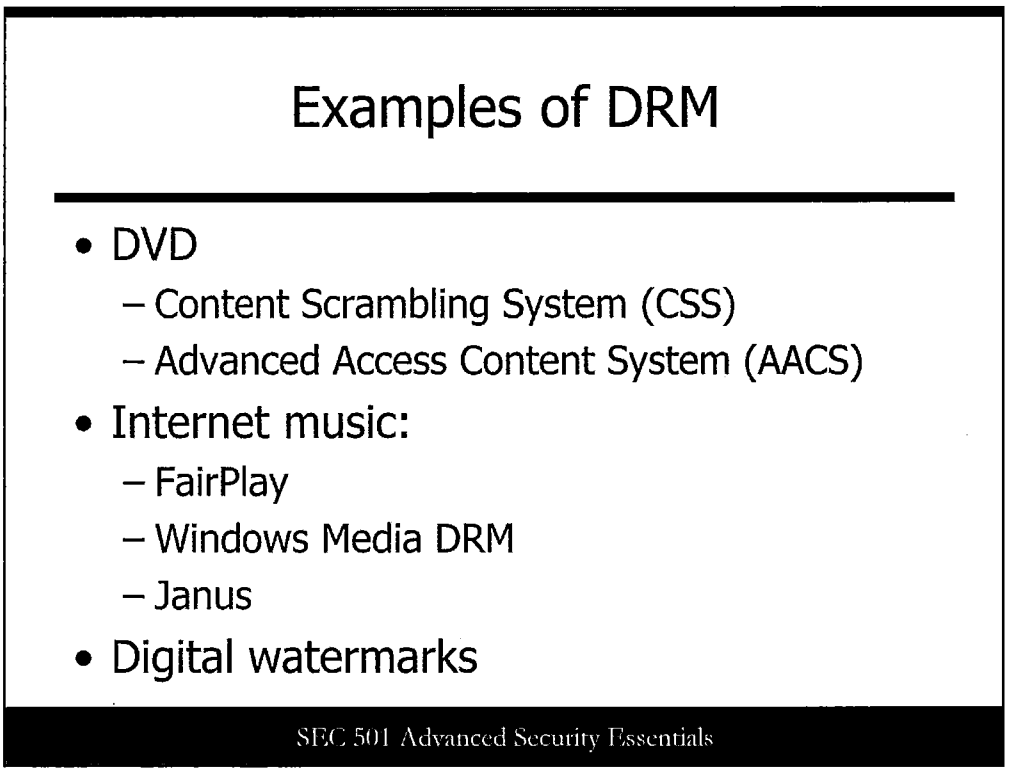

Examples of DRM methods include:

- **Content-scrambling system (CSS)**: Utilizes a weak, 40-bit stream cipher to actively encrypt DVD-Video.
- **Advanced Access Content System (AACS) HD-DVD, Blu-ray**: Encrypts content under one or more title keys using the Advanced Encryption Standard (AES). Title keys are derived from a combination of a media key (encoded in a Media Key Block) and the Volume ID of the media.
- **FairPlay**: Designed for The iTunes Library, iPod. Purchased music files are encoded as Advanced Audio Coding (AAC), and then they are encrypted with an additional format that renders the file exclusively compatible with iTunes and the iPod.
- **Windows Media DRM:** Provides secure and controlled delivery of audio and/or video content over an IP network to a PC or other playback device.
- **Janus**: Codename for portable version of Windows Media DRM for portable devices.
- **Digital watermarking:** Process of embedding information into a digital signal. This is the most widely used form of digital copy protection.

Digital rights help you handle and mark data appropriately. Data classification is used to mark and control access to data, whereas digital rights are designed to identify ownership and authorized use.

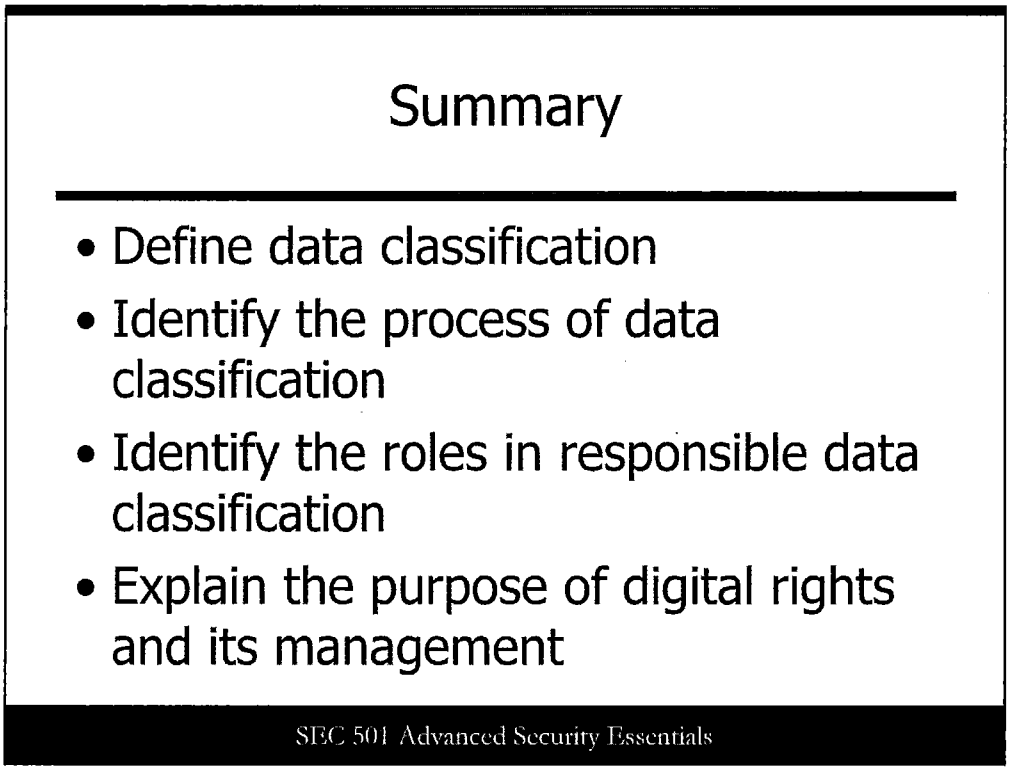

The following are the key areas that are needed for a robust data classification process:

- Defining data classification
- Identifying the process of data classification
- Identifying the roles in responsible data classification
- Explaining the purpose of digital rights and its management

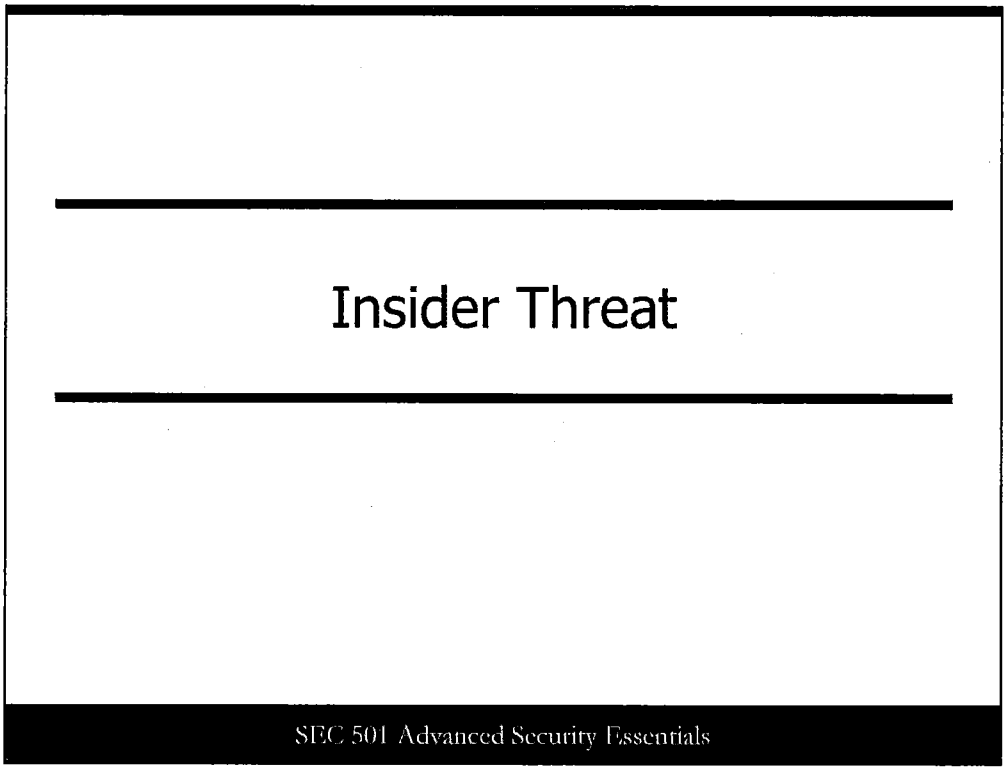

In this section, we discuss the controversial topic of the insider threat. Every organization has its skeletons in the closet, so to speak, but when we discuss security, it is imperative these skeletons not become an issue. We will attempt to make you aware of the types of threat, methods useful in identifying them, traits, and how to detect and mitigate them.

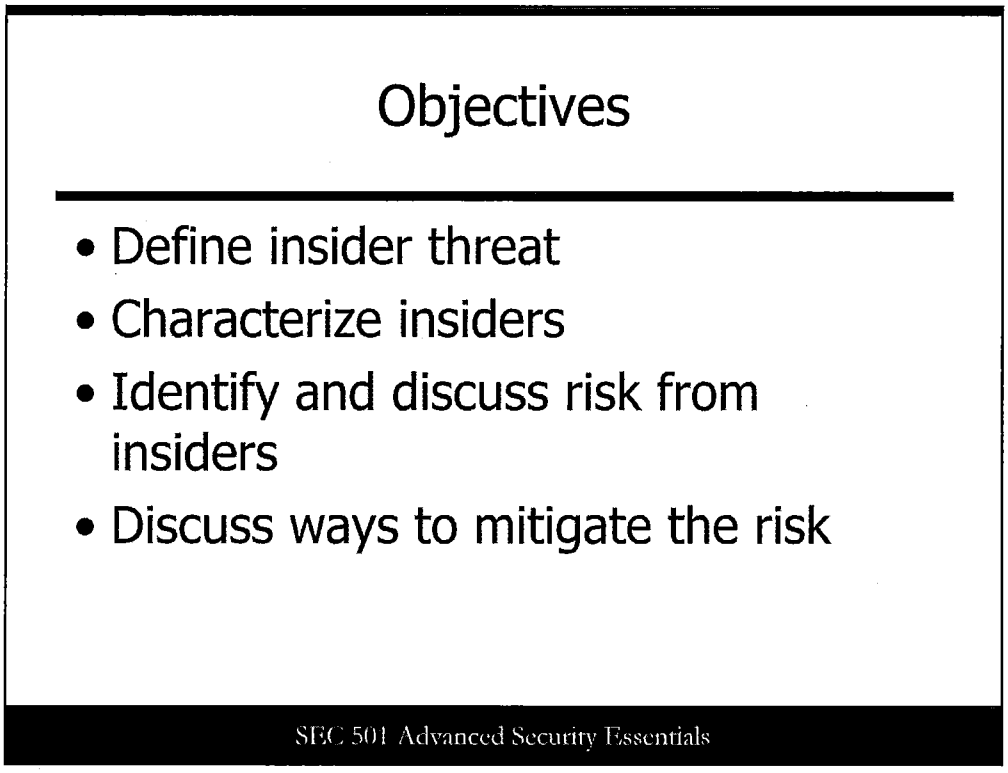

This page intentionally left blank.

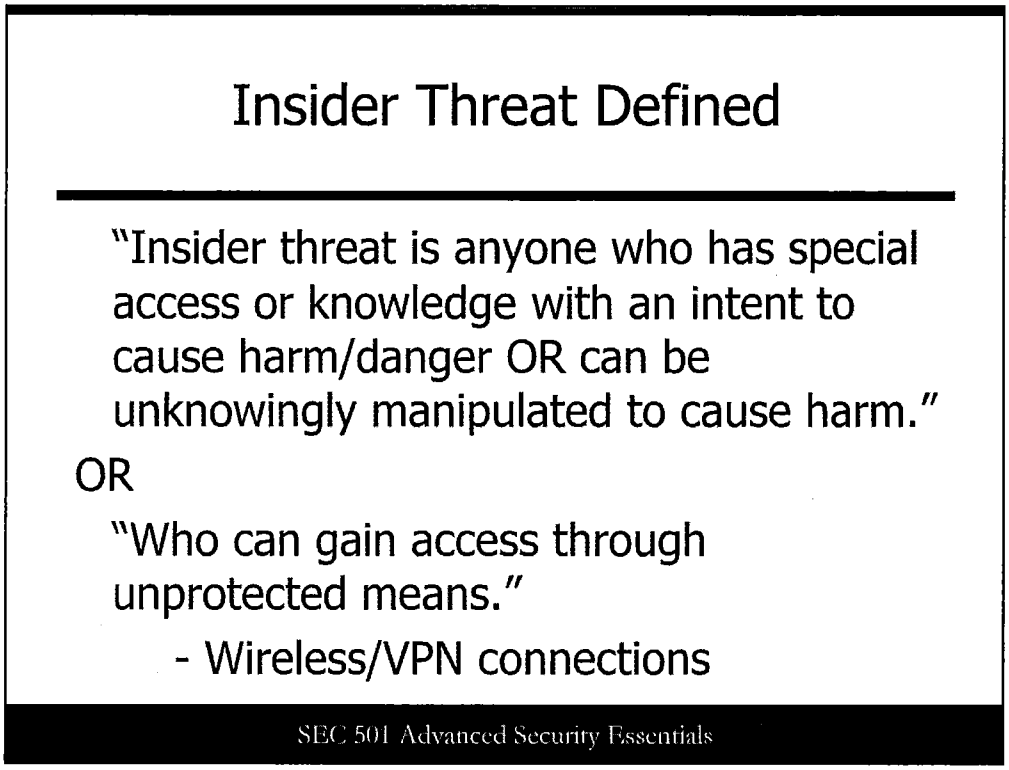

To protect against a threat, we must identify, define, and measure a threat. When I say measure, I mean to calculate or assign a risk to a particular threat or the mitigation process assigned to it.

An insider threat is not easily defined. Oh sure, you can say "well he works for us and is a company paid asset, so he's an insider". What about the temps we hire for part-time work or that holiday excess work load, or the contractor / sub-contractor we used on the last program. Don't forget about past employees. They are insider threats as well, even though they may not work for the company anymore.

Here's one for you, have you given any thought to the access given to our cleaning crew to areas no one else wants to clean up, yet houses our top dollar servers, or backup devices. Hmm, now there's a thought, the backups. I wonder where we store them, on-site, or off-site like everyone recommends. And if it is off-site, are we using a trusted source and have we identified the risks associated with them?

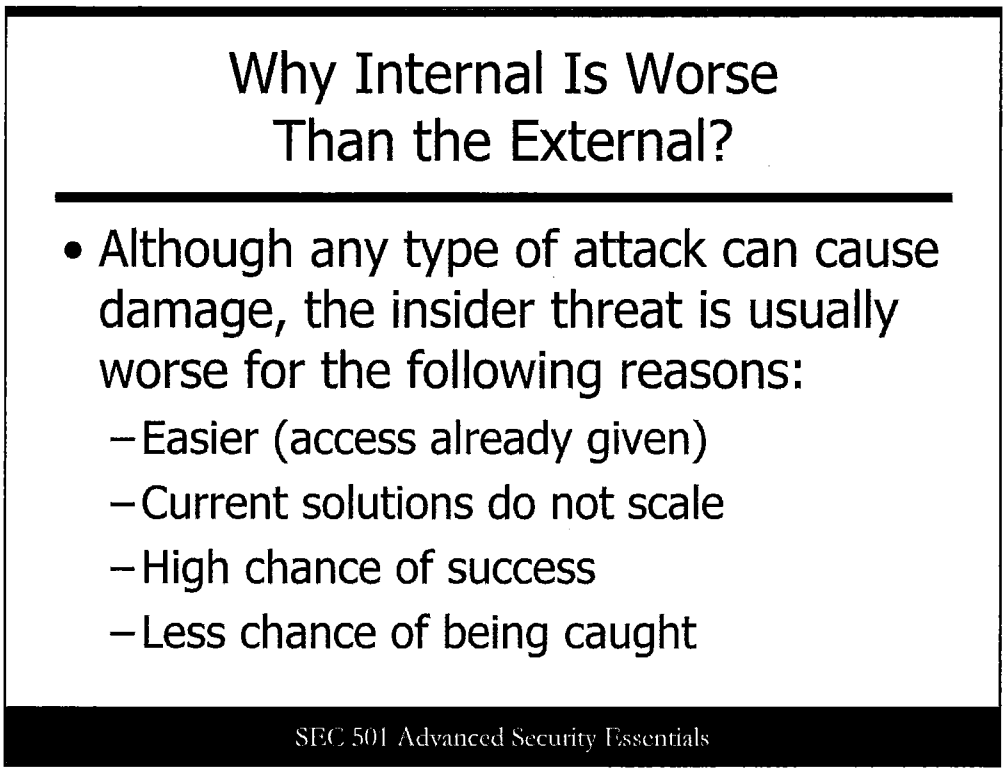

Internal threats offer one the biggest threats to any organization. Due to the nature of the threat from within, it makes it difficult to decide on the best protective measures.

Levels of access matter. The key to all security can be answered with a few simple questions: Is access required, and if so, what is the least amount required to get the job done? Initial access of individuals may be established and assigned properly, but given time, this access generally grows either through additional roles and responsibilities of the individuals or complacency and misplaced trust in the administration of accounts.

For example, John was a new employee several years ago. At that time he was assigned to the research department and allocated rights to a personal account (read/write/modify), a shared space (read/write/modify) and public space (read). Over time, John was assigned to various projects within the company, some of which have been mothballed (closed & archived), others that involve outside sub-contractors. Due to these rotating project assignments, John was given access to numerous areas on shared storage systems, but may not have been removed once the projects were completed or closed. This type of "access over time" without periodic review lends to complacency and weaknesses in our security controls.

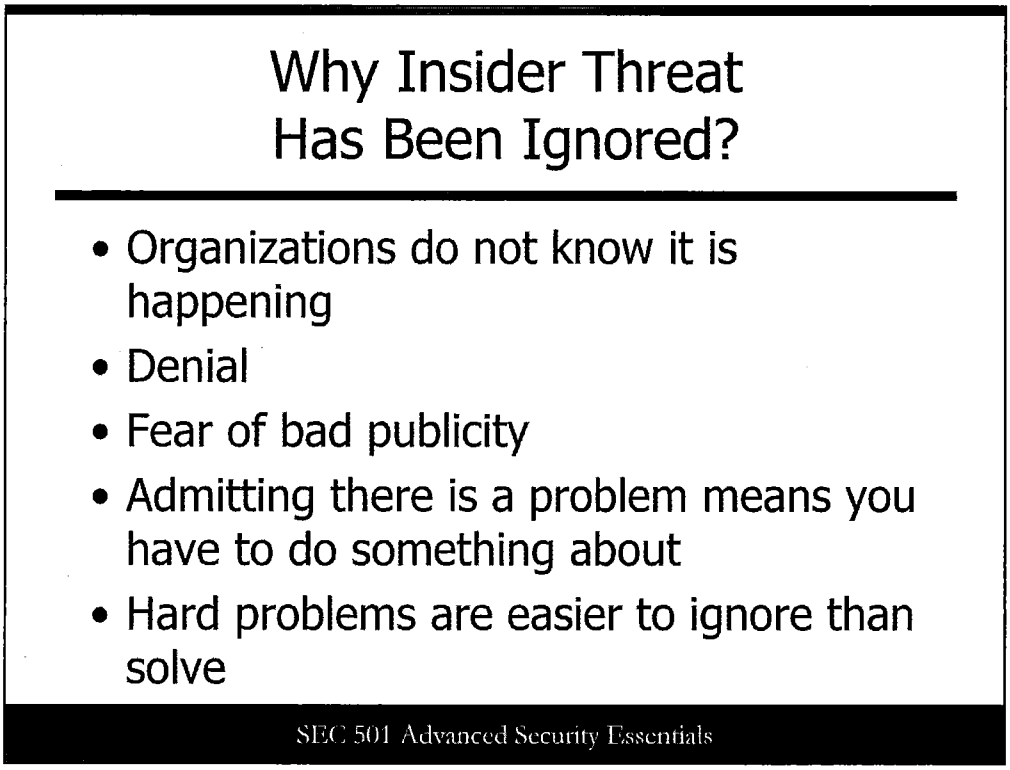

Insider threats are right in front of us, afford us the easiest (or so it seems) to control since we own the direct assets, and yet are the most overlooked and forgotten. Without on-going, actively participated in auditing and analysis (emphasis on analysis) of access activities, insiders can get away with almost anything.

Organizations, through fear of losing public support and trust and damaging their reputation, will refuse to announce or totally deny internal actions caused harm. In some cases, openly announcing the identification of insider threat actions has strengthened a company's reputation when put in the right context.

For example: a financial institution discovered an employee had been skimming some of a customer's funds and placing them in the employee's account. The institution's audit led to the discovery but was unable to determine the full extent of the loss. The institution contacts the customer and fully discloses what was found and agrees to work with the customer to fully compensate the loss and remedy the problem. In this instance the customer could choose to withdraw and expose the institutional flaws, but in most cases, may feel appreciative in the identification, restitution, and commitment of the organization towards its customer's best interests.

Types of Insider Threat

- Authorized versus unauthorized
  - The more people with access, the greater the risk
- Categories of insider:
  - Direct employee
  - Contractors
  - Consultants
  - Temps
  - Former employees

SEC 501 Advanced Security Essentials

This slide attempts to narrow down the types of insider threats that may pose or are a risk to the company. A few of these I mentioned earlier, like Temps and Former employees.

We begin with the most important protection level, yet sometimes the most loosely managed, and that is one of authorization and levels thereof. How many administrators, or super-users, or back-up operators, or auditors, with elevated privileges exist in your organization? Do they support off-site or alternate site admin roles too, or do you have another group of administrators there as well? Trust relationships are set-up between organizational sites in efforts to enable good company information flow; but we fail to realize that an administrator on one site is not necessarily the one who should have that level of access at our site.

Hence, the threats authorized versus unauthorized. Let's remember a key phrase: Least Privilege.

Characterizations of insider threats. Remember the old adage, "You can't judge a book by its cover." As far as looks go, this may be true and we could quite possibly implement it. But what about the "within?" Mannerisms, actions, verbal innuendos, pictures (or lack thereof) of one's friends and family, and other items we use to judge one's "character"—do they still apply?

This is how we try to "size up" the human threats, because frankly, it's all we have to go on.

Can we rely on individual background checks? Ask yourself these questions:

- Who conducted the background checks?
- How thorough are they?
- What type of checks and issues are included that can help identify someone who may have tendencies to commit fraud or who may work for an outside agency and steal trade secrets?

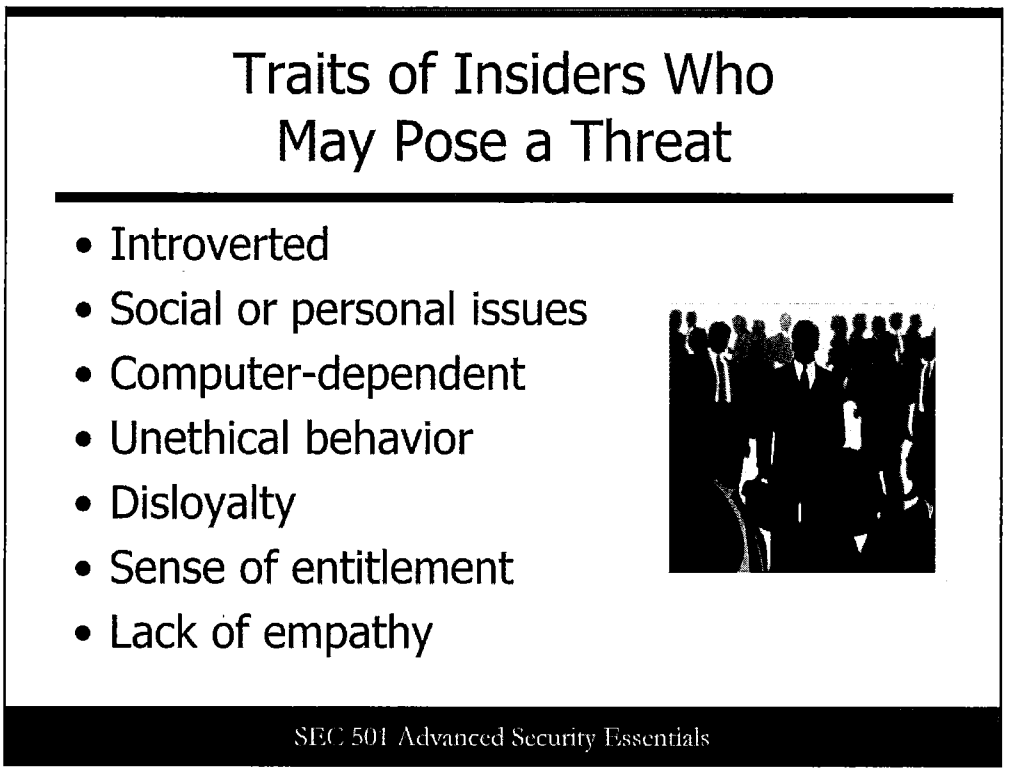

No one, not even the most experienced psychologists and criminal investigators, can always predict who is a definite threat, or who may go to the "dark side" and for what reason. We can train ourselves and our co-workers to be aware of the following traits that have been known or seen in past insider case personnel:

Characterization of insiders who pose a threat:

- Introverted and / or have social or personal issues
    - Keep to themselves
    - Experiencing financial problems
    - Going through a bad relationship (especially if with a co-worker)
- Computer-dependent
    - Exhibit a need to maintain control of data or systems
    - Express long term frustration with existing equipment
- Unethical behavior
    - Excessive communications with outside agencies
    - Non-contractual dealings with outside agencies
- Disloyalty
- Sense of entitlement
- Lack of empathy

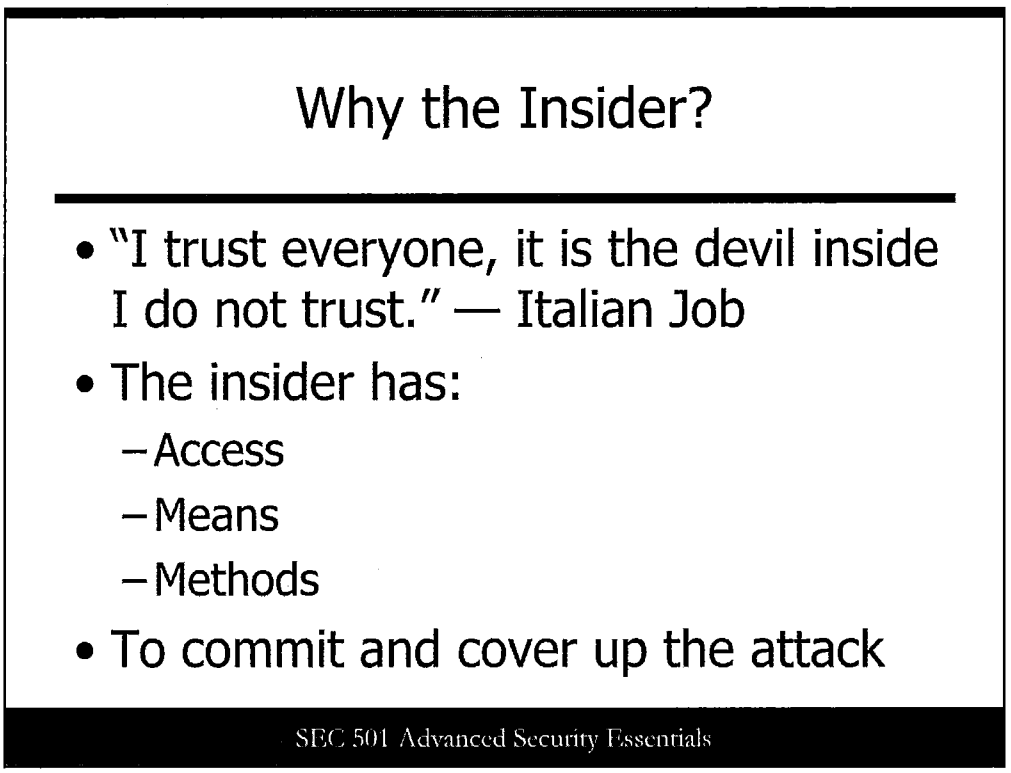

Insiders, at all levels and "categories" are and will remain our biggest threat. We are our own worst enemy. As the slide says, "The insider has: Access; Means; & Methods". Given the right accounts and with the right privileges or the right tools, anyone can be a threat to any type of system or information. These are why our insiders are also targeted by outsiders. Social engineering targeted towards the right people with the right access is as good an attack weapon or tool as any other attack. Actually, it may be better, since attacks from the inside may raise fewer flags of nefarious activity.

The last statement, "To commit and cover up the attack," identifies the ability to take all that we've mentioned (access; means; methods) and we can elevate ourselves, steal data, disrupt business, ruin reputations, and destroy valuable assets, all without touching our perimeter defenses.

Hmm, kind of like putting a 10 foot fence around our house to protect from outside threats, and a giant sink hole swallows the house.

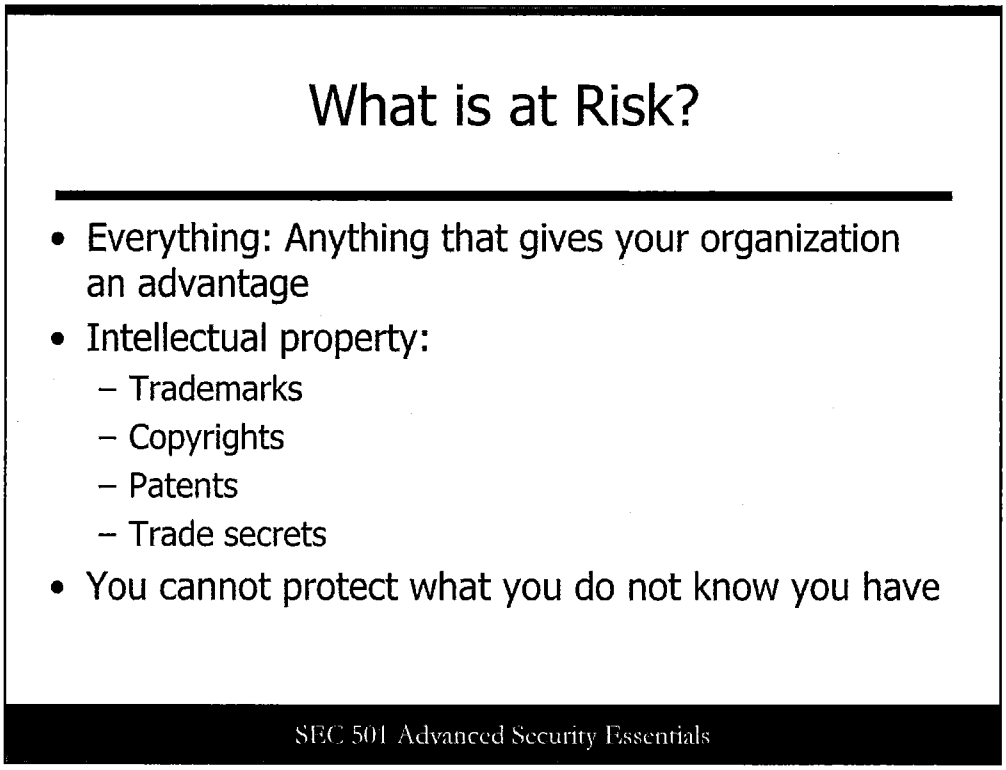

What is at risk? What is important to you and your company? Can you identify assets and information that have value, or better yet, can you assign a weighted value to what is the most important to you, or what you can do without should it disappear? If it is taken away or destroyed, can you recover? These things must be taken into account and then you must develop the protective means to defend them.

Assigning Value, not just monetary, is important. Identify data and resources required to continue operations, make profits, support our customers, ensure our job security, etc.

In order to defend and survive, one must know the risks and assign values to these risks so we can reduce or eliminate them. If all we do is reduce them, how much can we live with before they too become a greater threat?

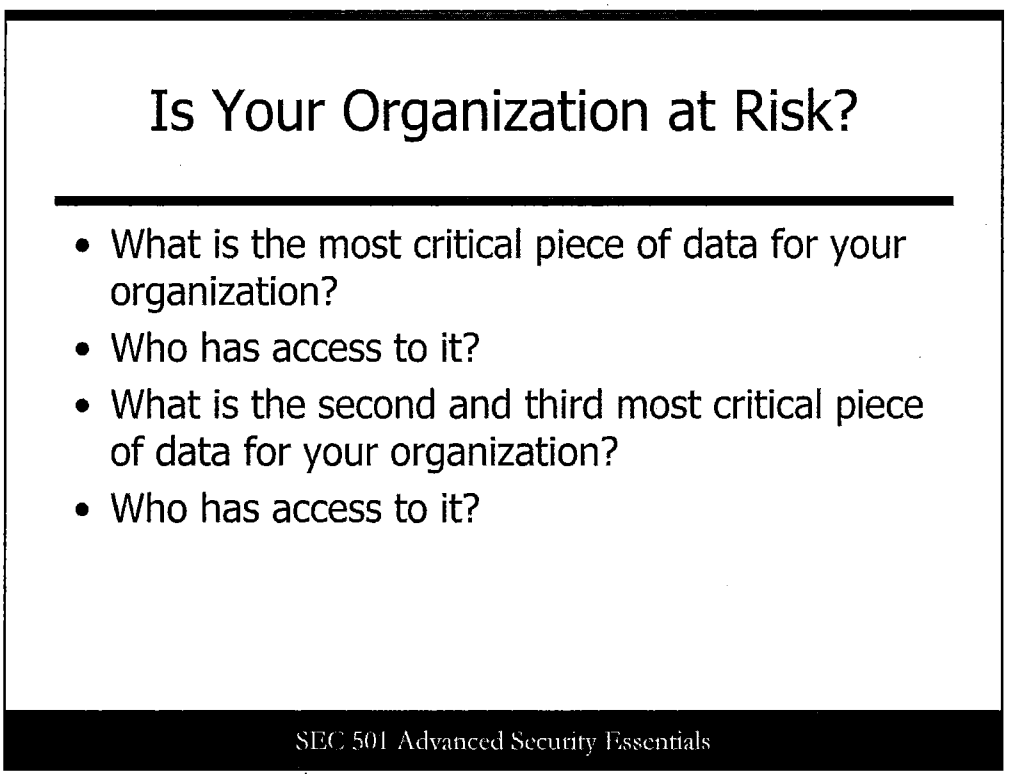

Okay, so you can identify the most critical piece of data for your organization. Can you identify all the individuals with access to it, and if not, what are the limiting factors that may help you to identify them?

Now continue on the path and identify and categorize all pieces of data as to their respective position within the company's needs and reliance. This path is one which will help identify risks, the level of protection, and what to include in recovery plans should they be compromised or stolen.

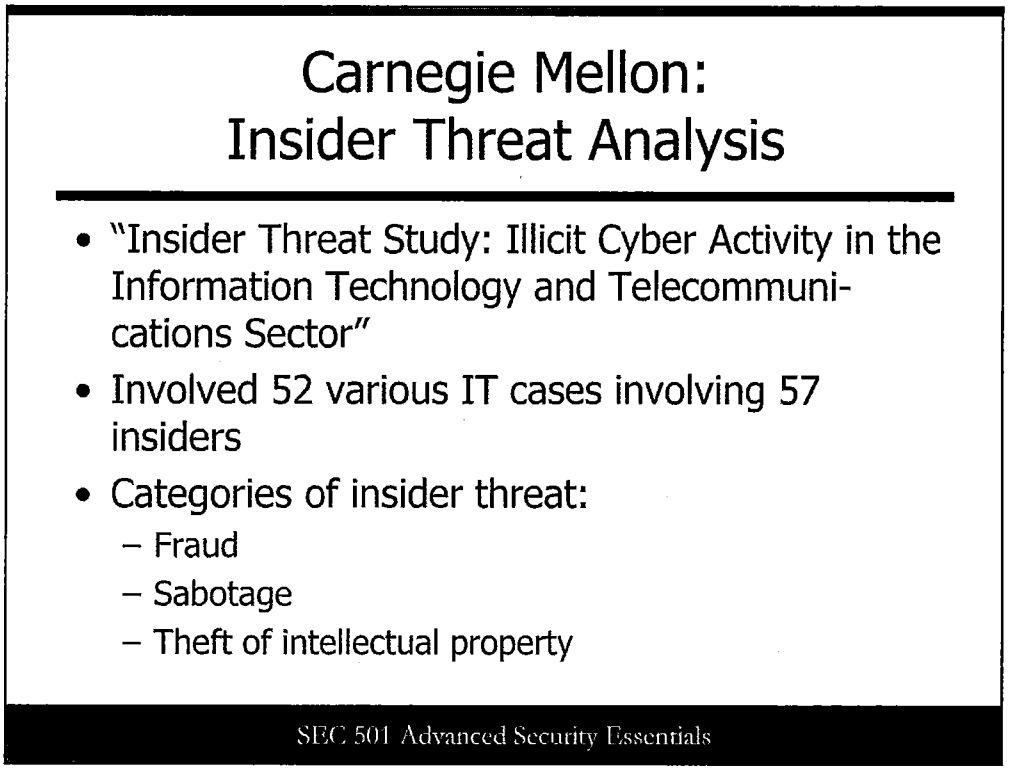

This report is a product of a study that included 52 cases involving various types of incidents involving 57 insiders that occurred within the IT sector.

It identified the attacks and involvement as:

24 – Solely sabotage
11 – Solely theft of intellectual property
8 – Solely fraud
6 – Sabotage and theft of intellectual property
3 – Fraud and theft of intellectual property

Ref: U.S. Secret Service and CERT/SEI, "Insider Threat Study: Illicit Cyber Activity in the Information Technology and Telecommunications Sector," dated January 2008.

Acknowledgements:

U.S. Secret Service, National Threat Assessment Center
Carnegie Mellon University, Software Engineering Institute, CERT Program
Former NTAC Research Coordinator, Dr. Marisa Reddy Randazzo
Former NTAC Special Agent In Charge, Matt Doherty

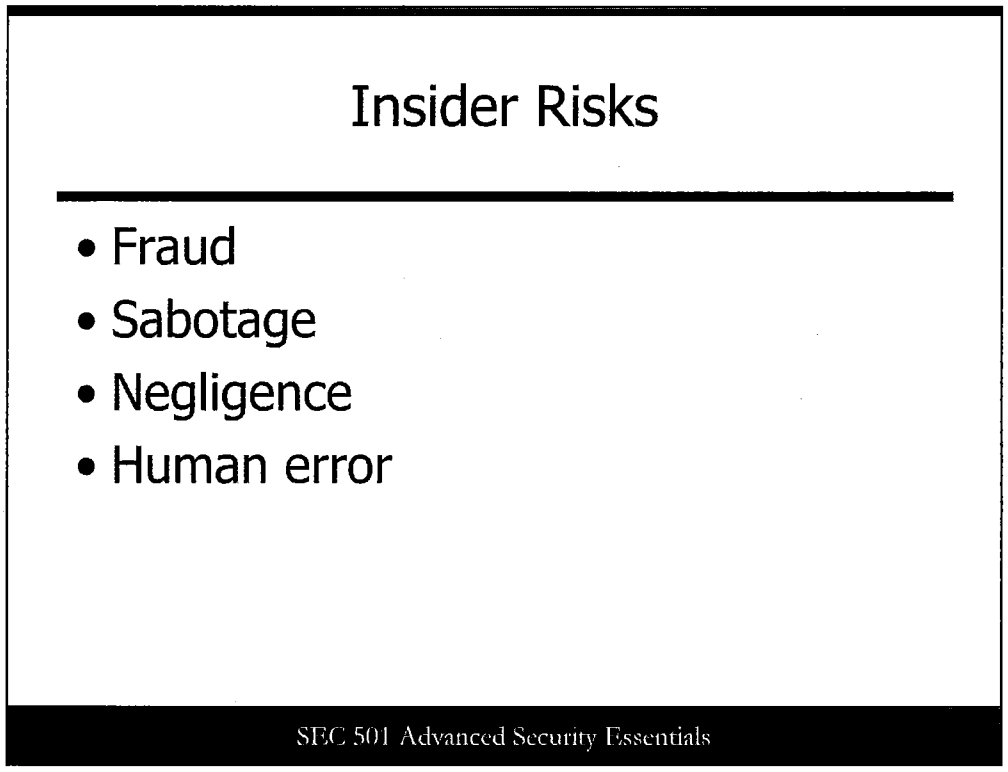

Continuing on from the previous slide, some basic categories of insider risk to become familiar with are:

- **Fraud**: Use of deception or trickery in an attempt to secure unfair or unlawful gain
- **Sabotage**: A treacherous action to defeat, hinder, or destroy property, or an operation in a deliberate act of subversion
- **Negligence**: Failure to exercise due care or diligence resulting in the injury of another party, device, or data
- **Human error**: The outcome of uncontrolled on unknowing actions of a human being lending towards the propensity of common mistakes

I bet the last one surprised you, and you found yourself saying, "I never thought of categorizing that as a risk." The others (fraud, sabotage, and negligence) get attention because they give us targets to blame and allow us to accept the reasons. They are even on the forefront of our minds when making threat determinations and continuity plans. But, human error, now that's one we have to review every day, because the potential for it happening to our data is far more often that than the rest. Human error can occur anytime, by anyone, anywhere on our systems.

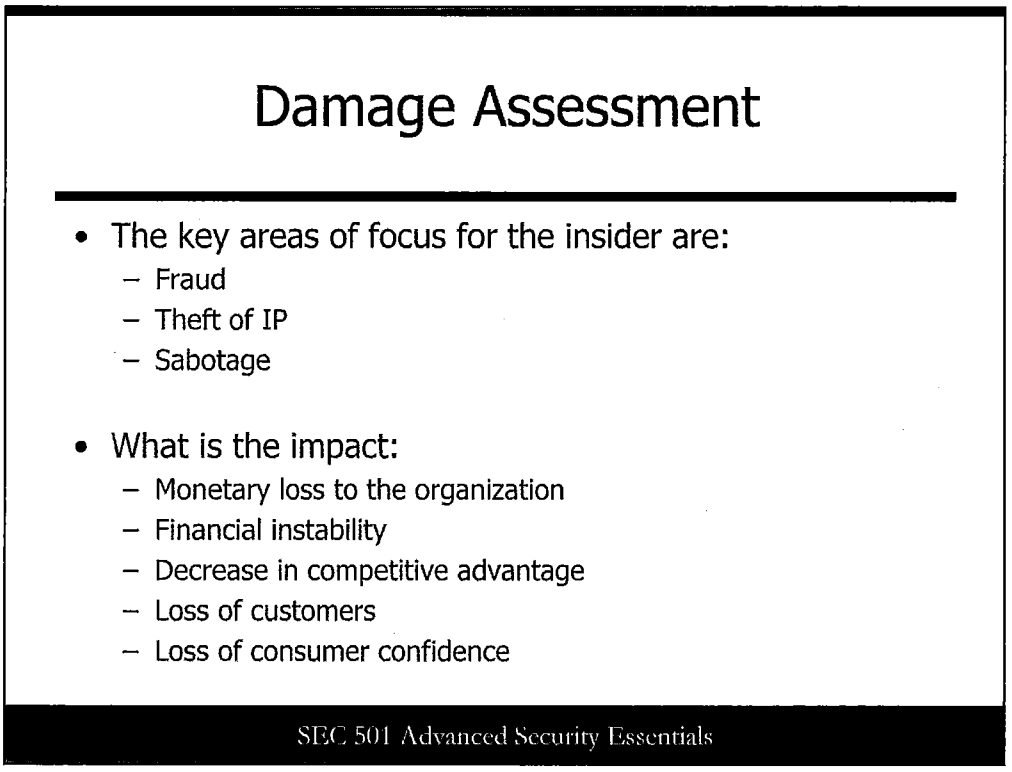

Damage is what we strive to avoid or reduce. Damage assessments pinpoint our flaws or weaknesses, our greatest or insignificant losses, and provide us goals when designing our prevention methods. Damage assessments are not just conducted after an attack, but through identification and analysis before an attack.

In terms of insider damage to an organization, items such as theft of intellectual property or trade secrets, fraud and misrepresentation, and sabotage are in the forefront of our mind.

Should any of these occur, what is the impact? How deep and long could it hurt the company?

Items such as monetary loss, though devastating, are recoverable, unless it was due to another impact, such as loss of trade secrets to a competitor who capitalizes on the gain and you are left with the question: how do we regain that knowledge and competitive edge.

Loss of Reputation is high on all our lists. This can lead to loss of current customers, loss of potential customers due to loss of confidence, and even loss of employees from feelings of betrayal and fear for loss of jobs.

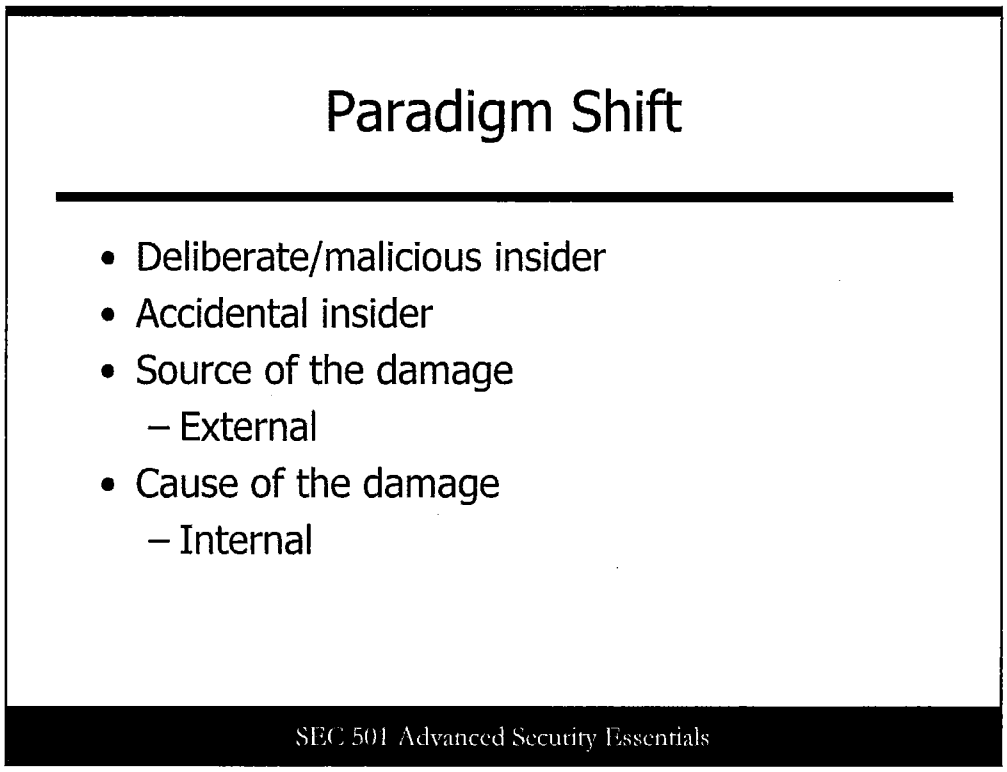

This page intentionally left blank.

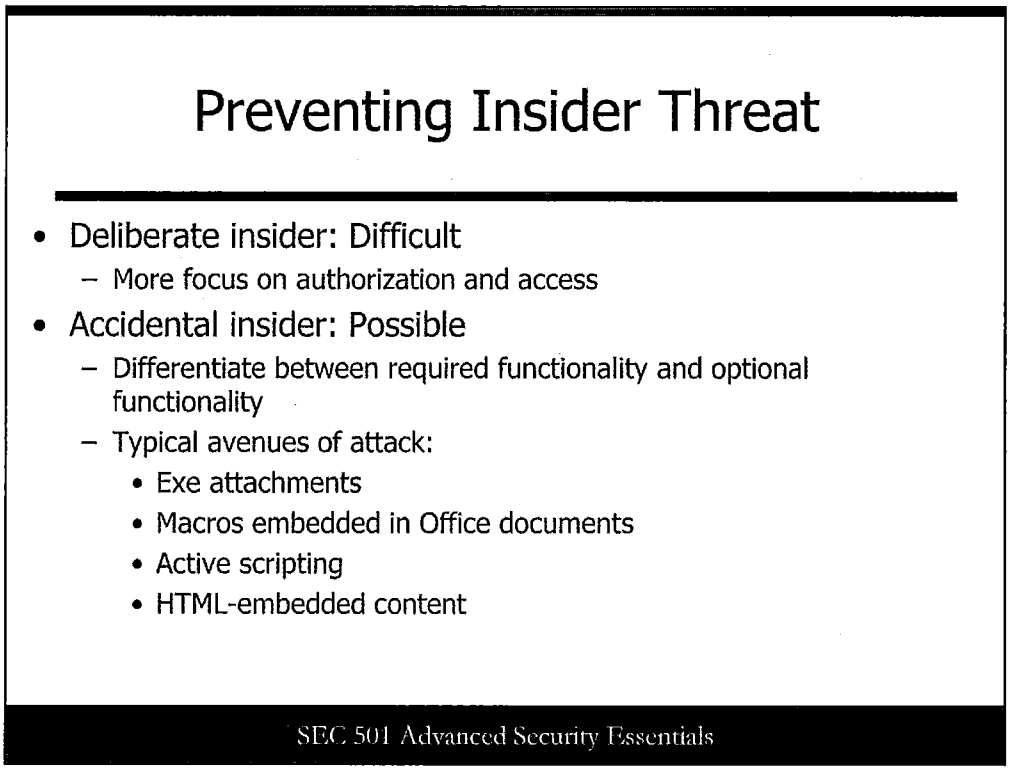

This page intentionally left blank.

Detecting Insider Threat

There are differences in activity between a normal user and an insider threat.

Activity patterns focused on data:

- Amount of data accessed
- Failed access attempts
- Data copied or sent to external sources

This page intentionally left blank.

Detecting Accidental Insider
Focus on Command and Control Channel

- Accidental insider is targeted by external entity
- Almost all external attackers set up C2
- Focus on outbound traffic:
  - Number of connections
  - Length of the connections
  - Amount of data
  - Percentage that is encrypted
  - Destination IP address

SEC 501 Advanced Security Essentials

This page intentionally left blank.

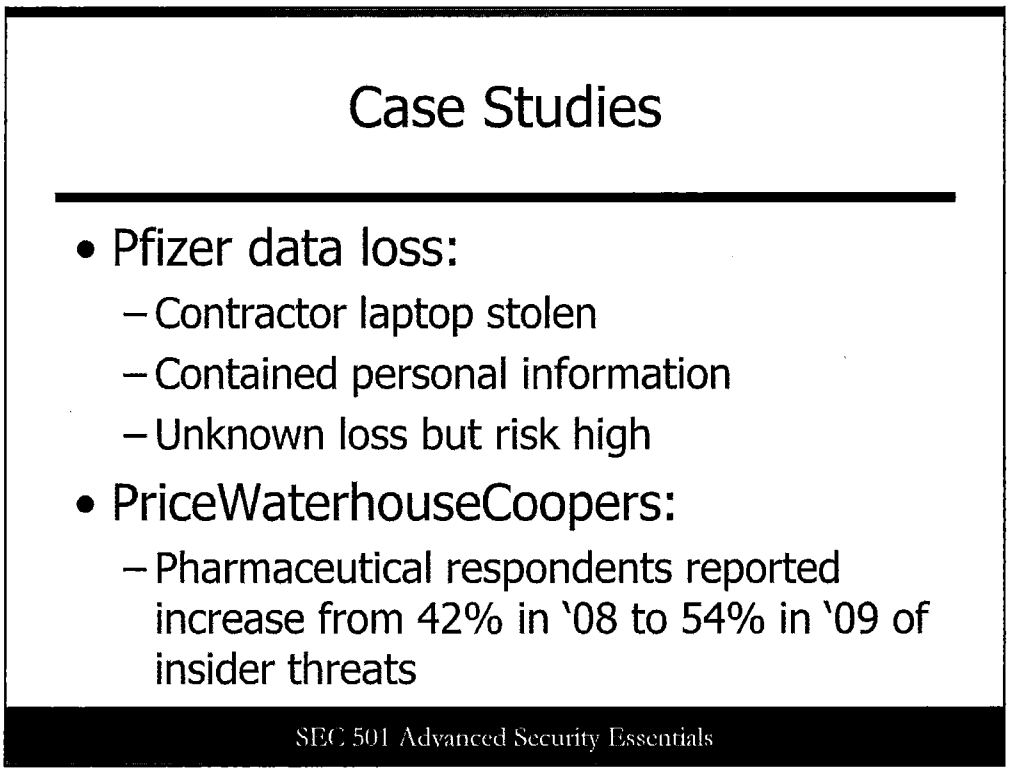

In order to drive home the key points, several case studies will be reviewed.

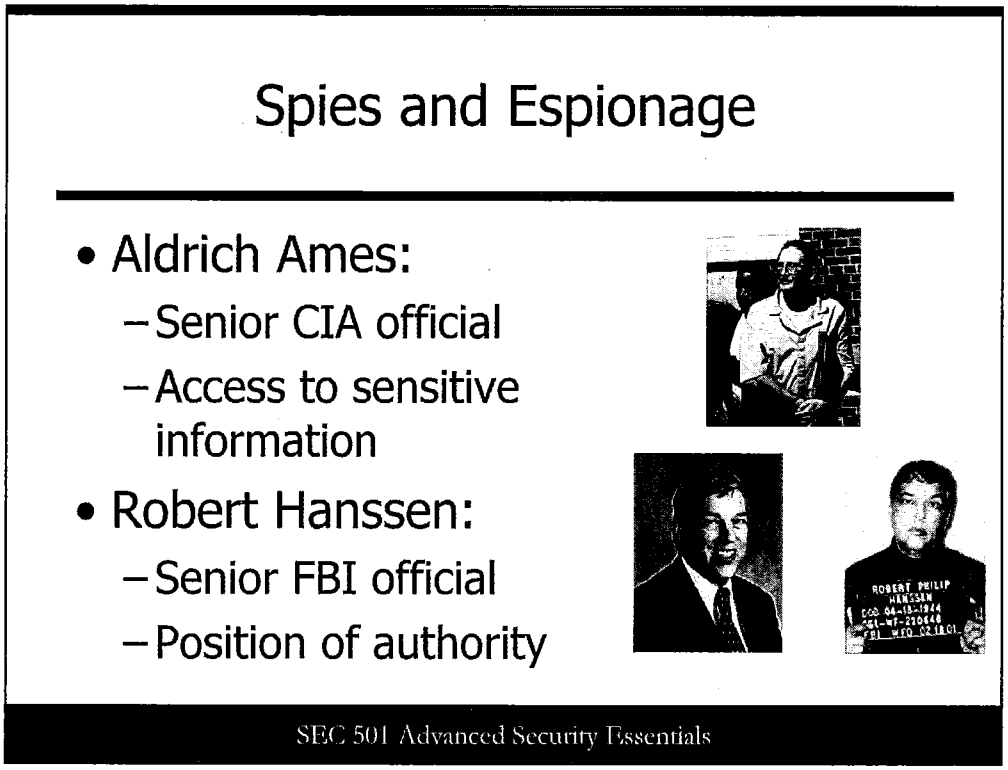

Why has it been over 14 years since we have arrested and imprisoned insiders? Can you say Complacency! Or, I get it; these are our own people responsible for tracking and investigating these things!

This is why our society and organizations NEED to focus on "Checks and Balances," Roles and Duty Rotation.

There are many reasons people go over to "the dark side" and what draws or leads them there. Money, disgruntled, a sense of belonging, all these can be managed and manipulated. Yes I said manipulated. With the proper social engineering and guidance, anyone can be lead to do something most people think is common sense to understand, or is illegal, but can be convinced to believe they are on the right side doing the right thing.

Hanssen

- Supposedly devout Catholic and active in Opus Dei
- Voyeur
- Old school dead-drops yet clever (selected his own drops and had pick-up add a "6" to published drop times)

Ames

- Lavish lifestyle gave him away
- Money

Don't forget the ones who have no reason, they just do it. No sense of remorse, total insensitivity to the outcome or whom they may have hurt. You cannot protect yourself against these types; we place them in a special place of "lost reality." Some may present characteristics as we previously mentioned, but not all of them do.

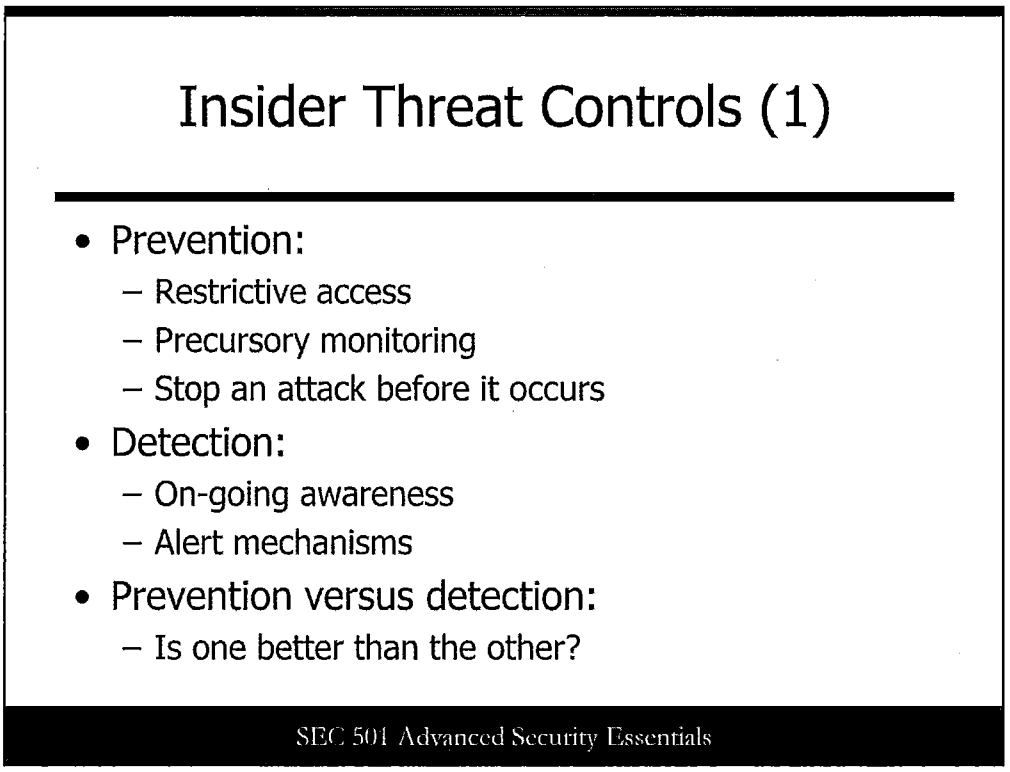

Insider Threat Controls (1)

When dealing with insider threats, what methods or controls are available to us? Two classifications of effort allow us to determine appropriate actions.

Prevention is the action of stopping something before it can start or spread. It's under this category that selecting the right security controls offers the most benefit. Actions towards personnel backgrounds and access fall under this category.

Detection is the ability to identify actions as they occur. Detection is a reactive method, because an event must occur before it comes into play. Detection does not afford protective abilities, only awareness and alerts.

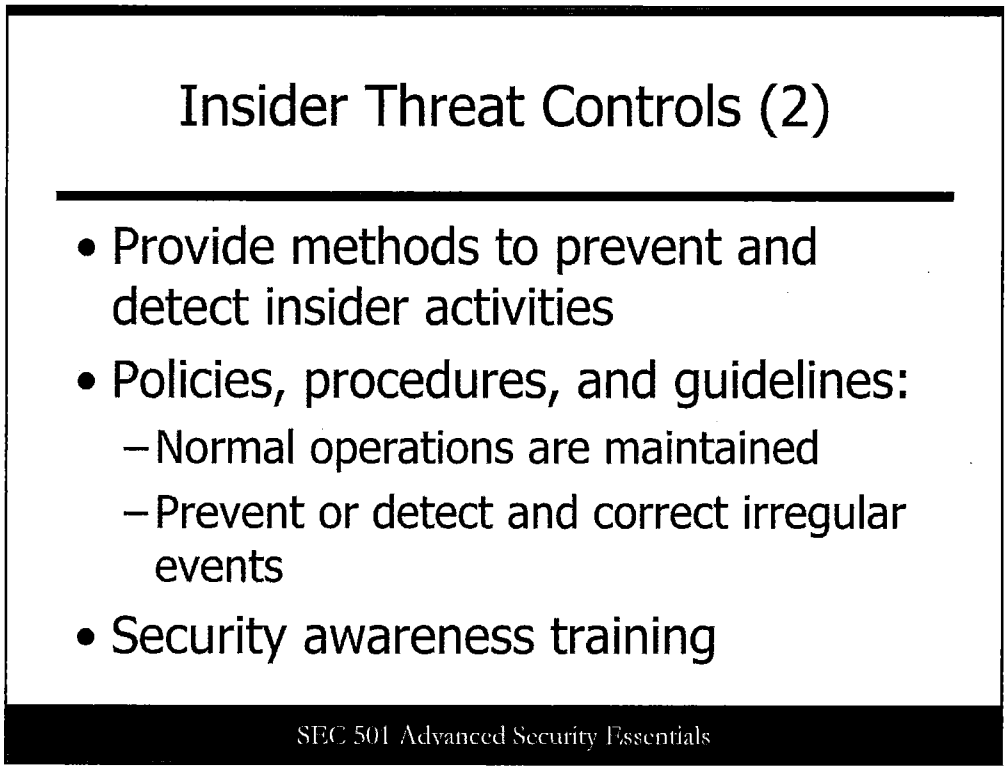

Insider Threat Controls (2)

Controlling insider threats can be done in any number of ways. Once a risk assessment and identification of critical assets is conducted, plans, procedures, protective actions, and devices can be implemented.

Plans, procedures, guidelines, policies, and organizational structure (roles and responsibilities) are the managerial methods available to identify and control insider threats. These measures must be granular enough to offer strict guidance and compliance, yet flexible enough to allow for incorporation of additional measures should the need arise.

Security awareness training is a proactive approach that enhances any security program. It offers protection through the education of employees to heighten their awareness of activities around them.

Insider Threat Controls (3)

Separation of duties allows for the individual and more finite control and monitoring of particular access levels. The threat of individual capabilities still exists. It is now segregated and monitored.

Rotation of duties offers several benefits to an organization:

- Reduces or eliminates "empire building"
- Spreads knowledge and control across the organization
- Ensures continuity of operations should an event occur such as loss of life or departure of key personnel
- Enhances internal growth by sharing knowledge within
- Gives employees a shared feeling of ownership in the organization

Not all positions can utilize rotation of duties, especially those requiring specialization skills. For those duties we move on to other methods of control:

- **Least privilege**: Allow only sufficient access to conduct required duties. This mechanism requires the detailed breakout of required duties and actions required to conduct them. Audit controls can be implemented to monitor restricted or specialized duties.
- **Classification of assets**: This involves identifying all assets, categorizing them on levels of importance or criticality, assigning weighted values, implementing control mechanisms to maintain restricted access controls to each.

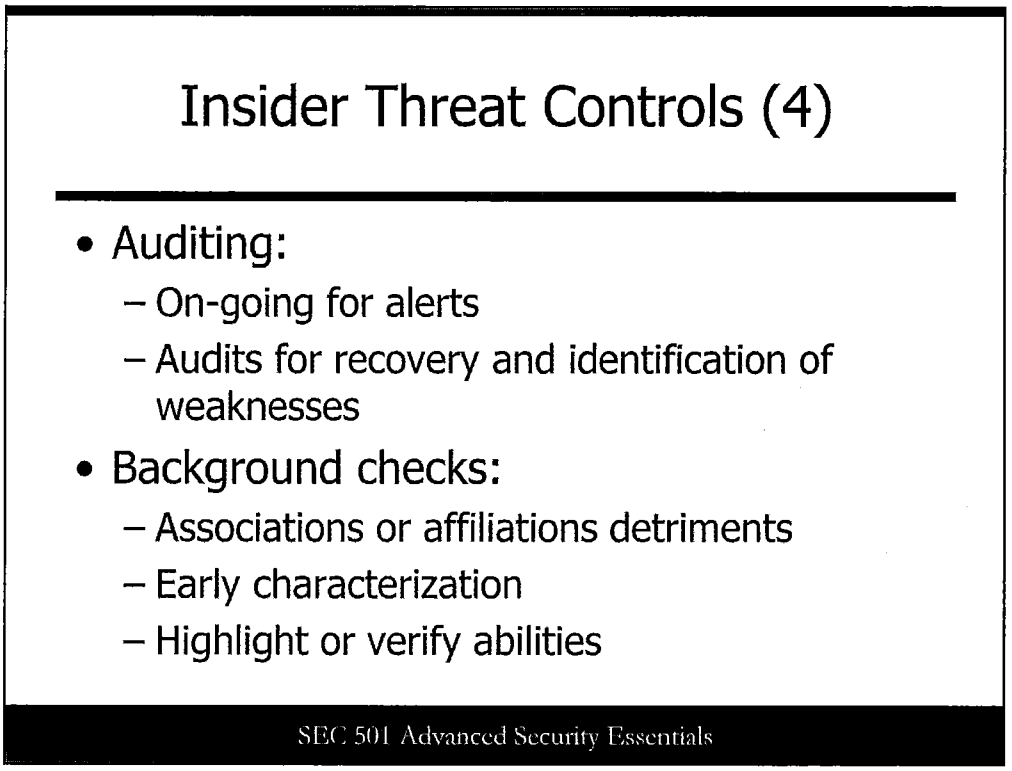

Insider Threat Controls (4)

Auditing controls are instrumental throughout a security program. Auditing is only as good as those who view and analyze it though. To be an effective tool, it must be used. Reviewing logs regularly, and reviewing audit levels and items being audited, can enhance security and offer a good deterrent to insiders.

Background checks during the hiring process are useful tools. They afford us the ability to get to know an individual before placing them into the company work environment. However, a background check is only as good as the one conducting it, and the data to be scrutinized. The higher the classification or sensitivity of data to be accessed or protected, the further back or extent a background check should be. For example: Development of an intelligence collecting software or appliance to be used against a country an employee may be from or still have ties to may warrant a background check that ensures a good representation of the individual's loyalty or on-going affiliations to their home country.

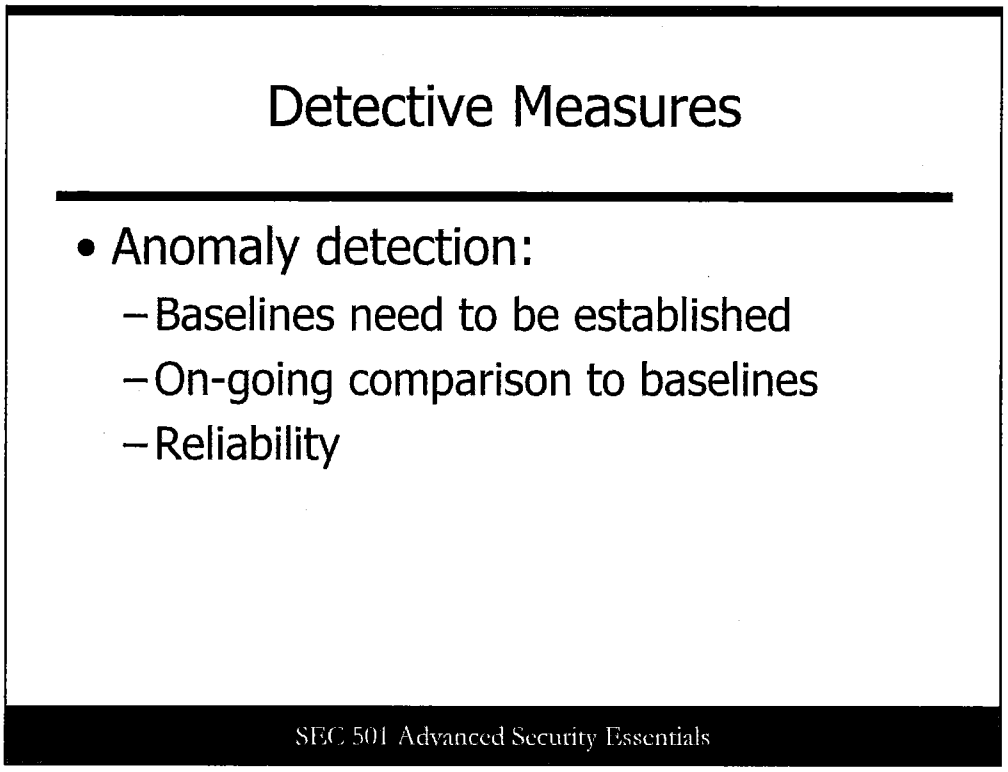

Anomaly detections work off the premise that a good baseline known to be supportive of operations is created, a mechanism is put in place to monitor activity, and activity is compared to the baseline in an effort to determine anomalies. False positives are possible if baselines are not maintained. Anomaly detection applications are more than just data traffic; they include log comparisons of access rights and biometric differences in individual accesses, and may not always be accurate.

- Signature analysis
- Thin Clients
- Background checks
- Defense-in-Depth
- Looking beyond technology
- Archive critical data
- Complete solutions

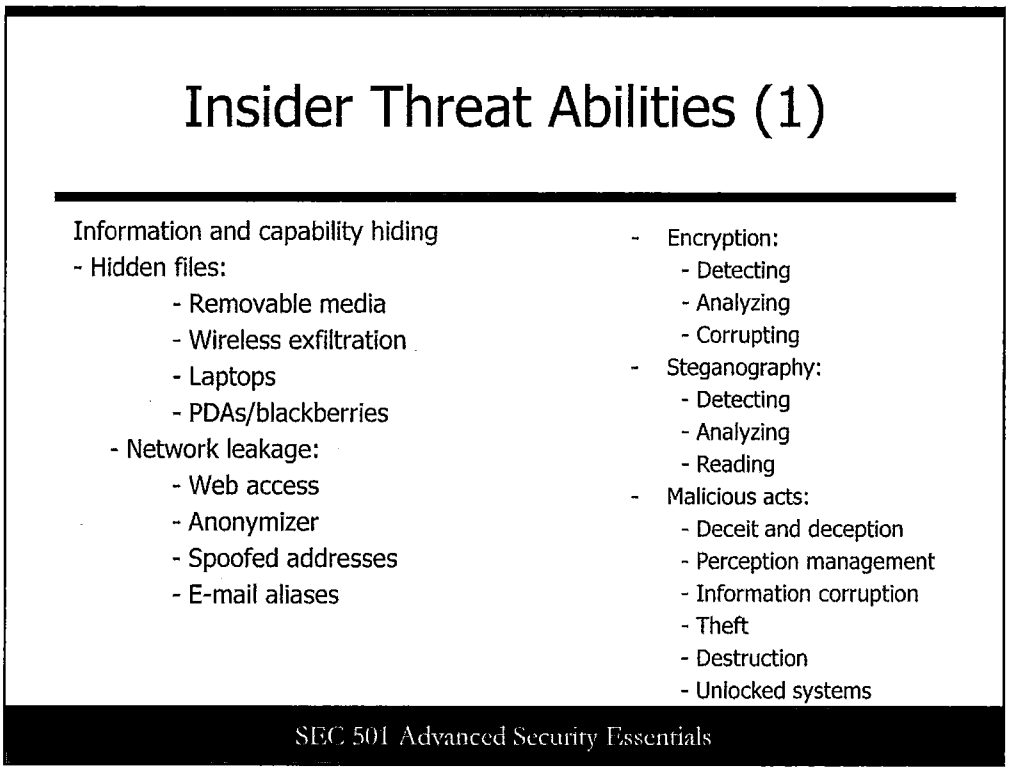

Insider Threat Abilities (1)
Encryption and steganography software and devices, while affording protection of data, also offer an insider threat, which is the ability to circumvent protective and detective methods. If an insider cannot access data he desires, he can still destroy or corrupt the information by simple file manipulation.

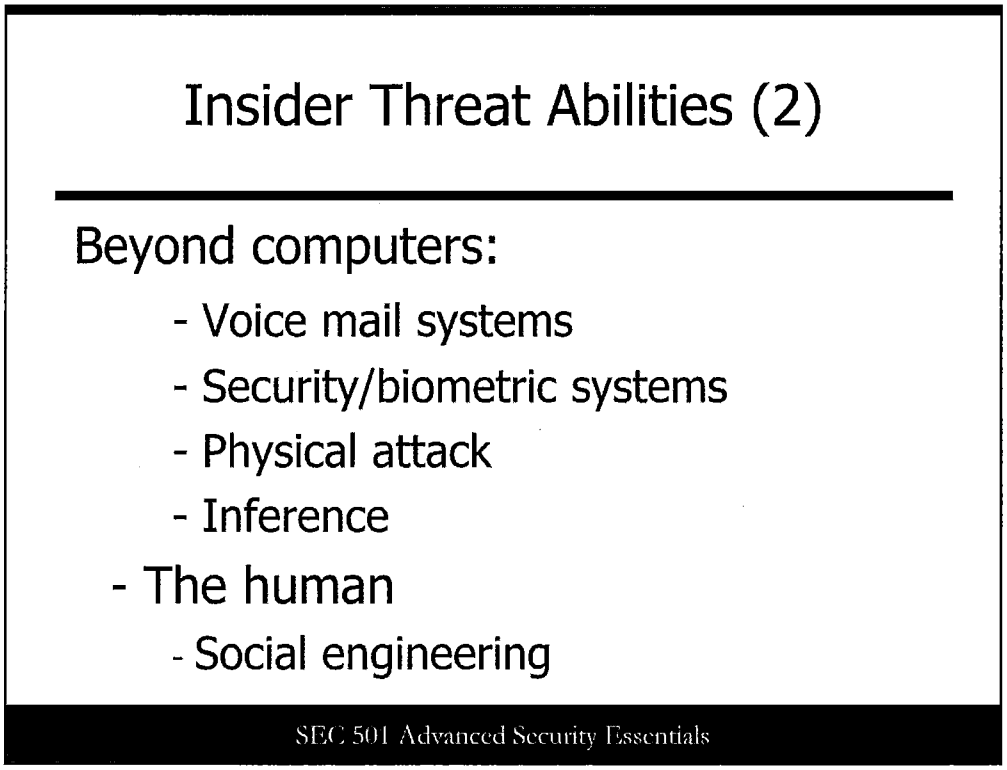

Insider Threat Abilities (2)
While technology has changed how organizations operate, individuals caused harm to an organization even before there were computers. In addition, networks have made the job of the insider threat easier, though traditional methods still work very well. In evaluating the abilities of the insider, it is important to account for non-technical means of causing harm to an organization.

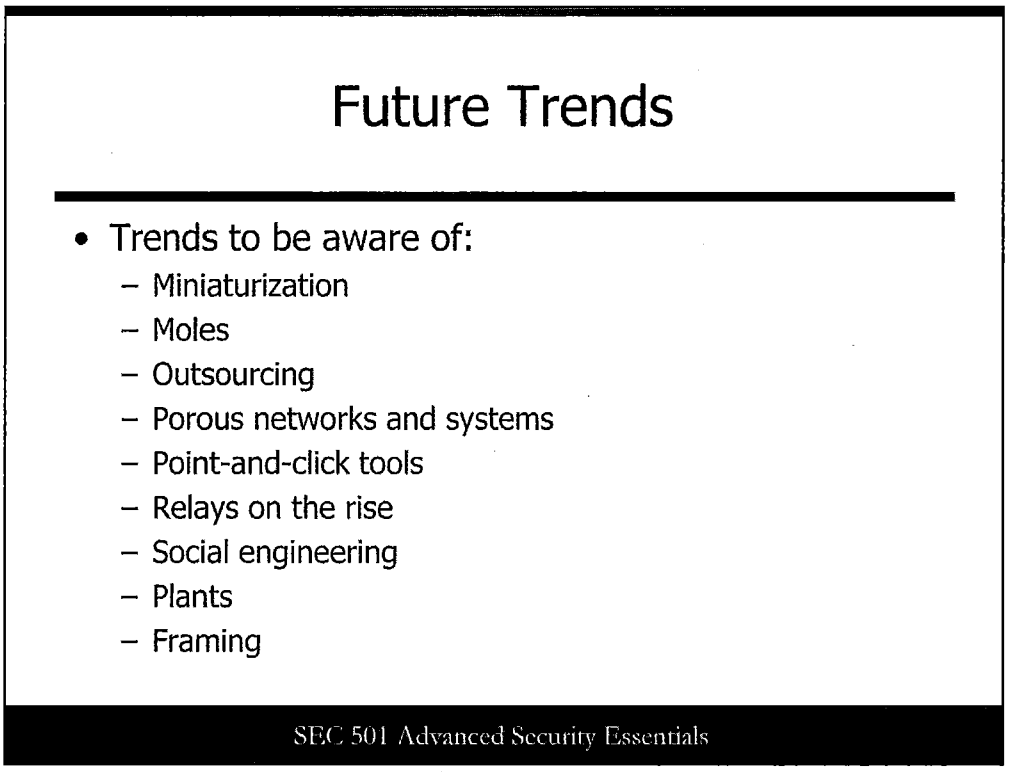

Technology is always changing and it affords good and bad opportunities.

Summary

- Define insider threat
- Characterize insiders
- Identify and discuss risk from insiders
- Discuss ways to mitigate the risk

SEC 501 Advanced Security Essentials

This page intentionally left blank.

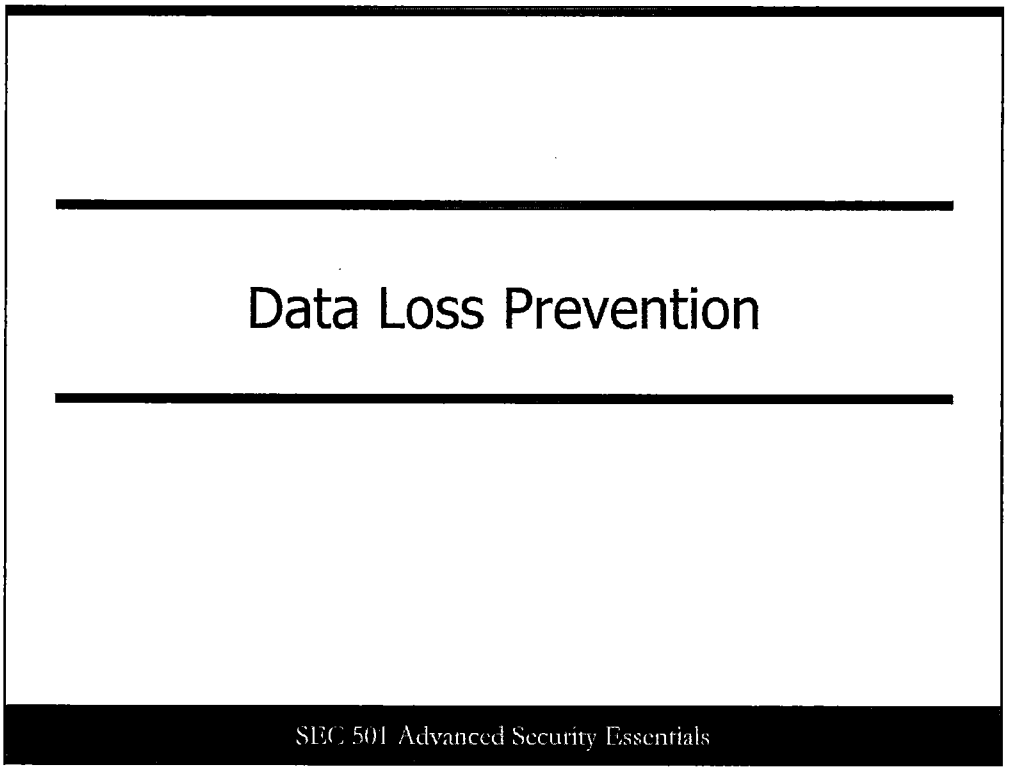

This page intentionally left blank.

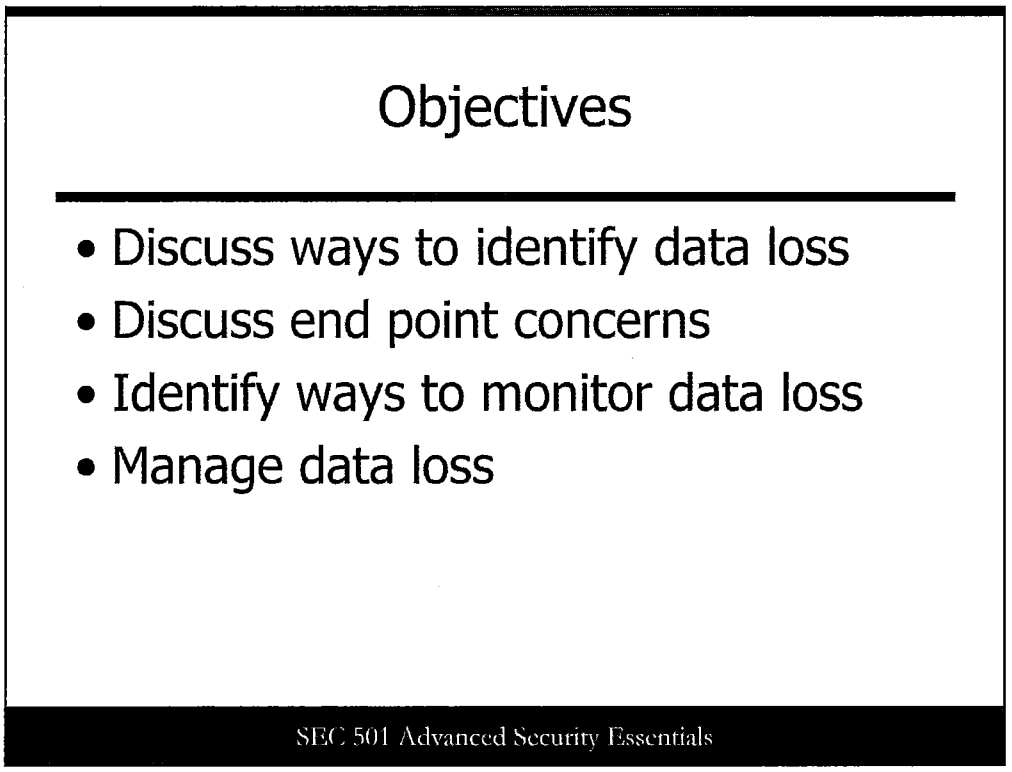

This page intentionally left blank.

Steps to Data Loss Prevention (DLP)

- Identify or discover
- Monitor
- Protect
- Manage or mitigate

SEC 501 Advanced Security Essentials

In this section, we discuss methods we can use in the protection and prevention of data loss.

- Identify or discover
- Monitor or audit
- Protect or prevent
- Manage or mitigate

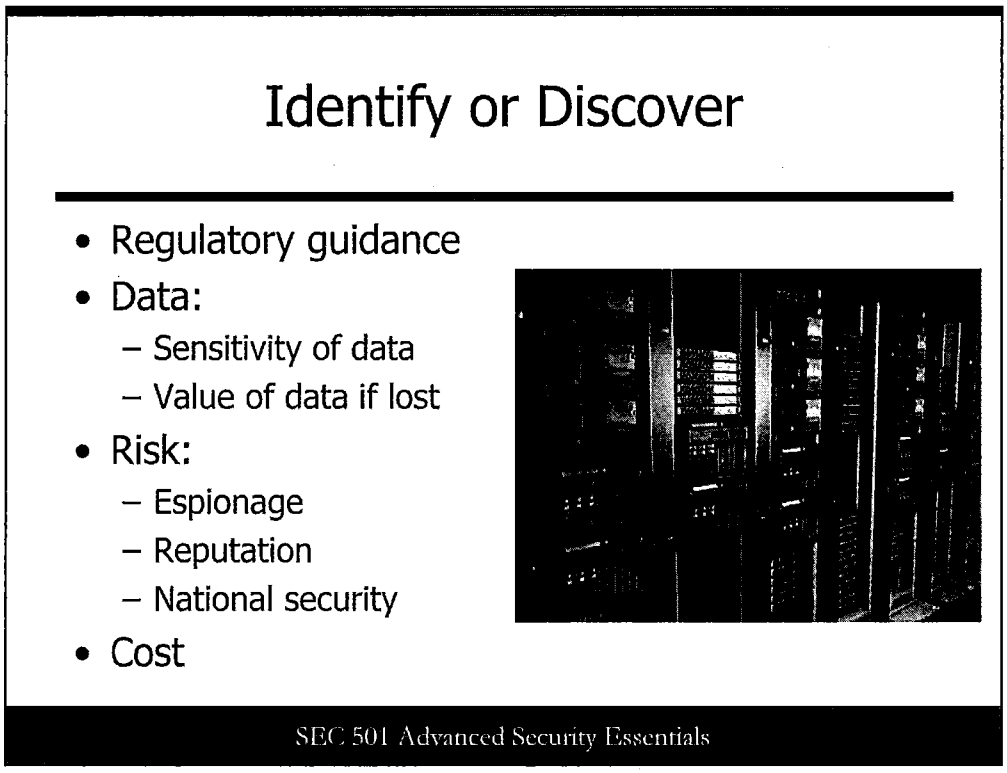

Identifying existing regulatory guidance will help identify existing procedures in protection, as well as afford guidelines in the internal or more finite development of policies and procedures tailored to your organization. Some examples of existing governmental guidance are:

- Health Insurance Portability and Accountability Act (HIPAA)
- Gramm-Leach-Bliley Act (GLB)
- Sarbanes-Oxley (SOX)
- Payment Card Industry Data Security Standard (PCI DSS)

Categorizing data assists in the focusing of procedures:

- Identifying the sensitivity of data, or better yet, developing a table to help place weighted values on information, is one method in prioritizing expenditures and protective measures.
- Placing or labeling values to data if lost is useful in: itemizing protective measures necessary to protect it; extent of recovery emphasis to be placed on it; and restitution issues if legal matters ensue and a value is asked for compensatory damages.

Risk Identification
Associating specific risks and vulnerabilities to data will assist in the identification of the levels of protection required. Information that can cause grave or serious risk to our national security, or that which may cause loss of life, would have a greater protection requirement than that of the reputation of the company.

Cost
This is addressed and useful in many ways: The assignment of values to data loss and to the cost of protective measures required.

Some End Point Concerns

- External USB hard drives
- Laptops, phones, PDAs
- Website postings
- E-mails
- Bluetooth devices
- Peer-to-Peer (P2P)
- Residual data transfers (data spills)
- Wireless access

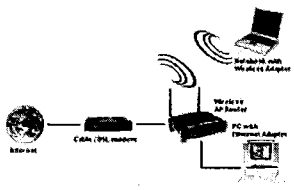

SEC 501 Advanced Security Essentials

Data is stored and handled in many forms. In the following slides, we attempt to break these down into sub-categories to allow for the tailoring of recommended protection methods. Keep in mind that some of these concerns can be found across multiple sub-categories depending on their uses, so one type of protection may not be sufficient.

At-rest Data

- Thumb drives
  - Thumb sucking
- Removable drives
- External drives
- iPods, MP3 players
  - Pod slurping
- Laptops
- Other storage devices (CD/DVDs, and so on)
- Residual data (data spills)

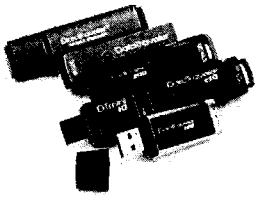

SEC 501 Advanced Security Essentials

At-rest data is a term that refers to information or files that are maintained on a storage device, whether it is permanent or portable storage. This includes your operating system, programs, and actual user data.

Technology has provided us a plethora of storage mediums, large data storage, and various sizes of devices. With sizes reduced and storage capacities increased, the threat of portable devices has increased.

Organizations that allow employees the use of portable MP3 players inherit greater risks because they openly allow storage devices in / out of their sites.

Laptops and Ultra-Portables already are known for internal storage, yet sometimes are overlooked for how they are protected (i.e., lack of encryption) and if they have built-in wireless / Bluetooth capability, which may circumvent networked security measures.

Residual data (data spill) security concerns center around the ability to inadvertently or intentionally copy or access data that was previously deleted, yet not overwritten. Most operating systems copy data in blocks or chunks of information. If a data block is reused by an application or file and it uses less space than the original data that had been stored in that region, a copy of that segment can include "residual data."

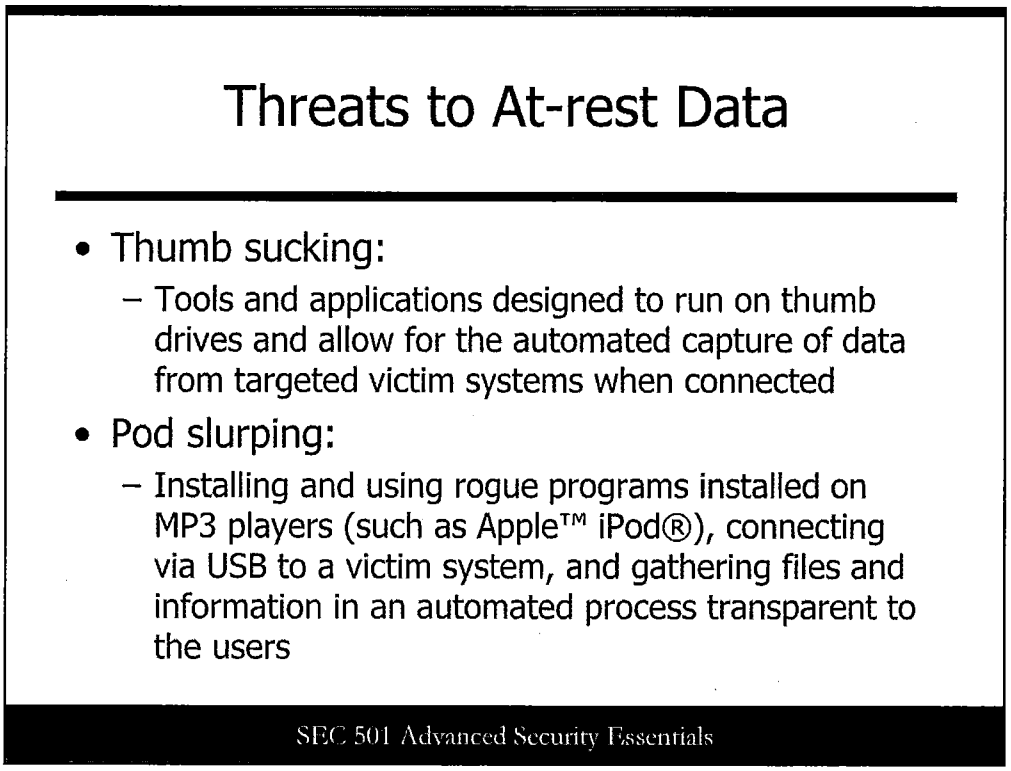

In this section, we discuss just two of the many ways to access and steal "at-rest data."

Thumb sucking, as its name implies, is a term that refers to the use of thumb drives to target and compromise systems. These methods revolve around automating an attack that uses the system's default "autorun" feature to mount and run the first autorun application the victim system sees on the inserted thumb drive. The tools then proceed to carry out specially programmed payloads. These payloads can attempt to elevate privileges of the current user, access and copy all files or specific files onto the thumb drive or compress them into a single file and e-mail them to a drop off point on the Internet.

Pod slurping is a term identifying the use of "pods" (original forms used Apple™ iPods®, hence the name) normally used for music or video files, to access systems and steal information.

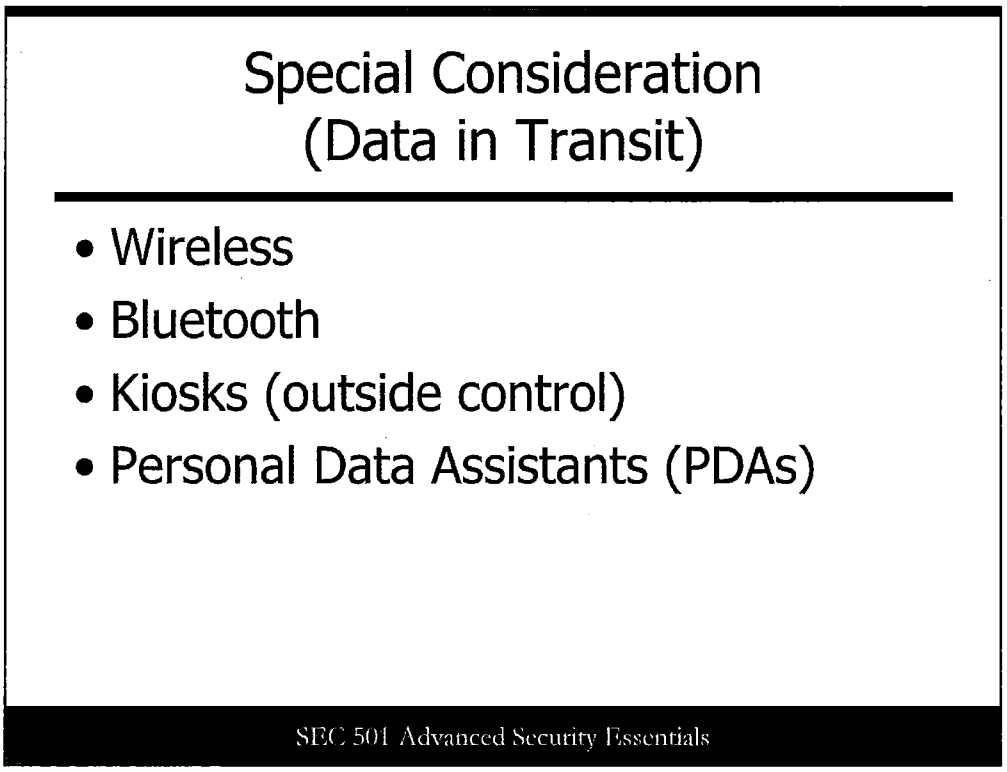

In the "Special Consideration" section, we discuss data threats targeting "data in transit," or, more specifically, the methods outside our control of security protection (or so it may seem).

- **Wireless:** This vulnerability or threat refers to using or relying on wireless devices that, either through peer-to-peer sharing or via access points, communicate to the Internet in hopes of reaching a known end point. During the data transfer and access, many options afford potential attackers the means to compromise the security of the traffic. Attacks can focus on the origination of communications and attack the initial handshaking or authentication of communications and present "man-in-the-middle" attacks in an effort to hijack the data. If control is not an issue, the attacker can simply monitor the traffic and analyze data within.
- **Bluetooth**: Threats to this service are similar to wireless, in that attacks against pairing (similar to the handshaking of wireless) are used in the monitoring of data. Bluetooth uses frequency-hopping technology (sends data in blocks across multiple frequencies randomly hopped across). Some tools developed and released to attack / monitor Bluetooth communications:
    - **Bluesnarf**: Tool able to monitor and record communication across Bluetooth connections.
    - **Bloover**: Tool designed to capture Bluetooth communication.
    - **Carwhisperer**: Designed to monitor and talk/text mobile hands-free devices using default passwords (0000, 1234, …).

While traveling, all of us have seen or used the public kiosks for the traveling Internet users. These areas range from allowing you to connect your device (most commonly wireless) to the ones where you are given guest accounts on their systems to use. Obviously they are not set up to record or store any of your precious personal communications, like say, your e-mail, your bank account PINs, or anything else.

Cloud Computing

- E-mail data and traffic
- P2P
- Voice over IP (VoIP)
- Online shared applications

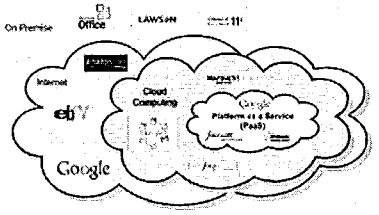

SEC 501 Advanced Security Essentials

Cloud computing refers to, in simple terms, data stored or transiting Internet devices. All business types and even reaching into our personal home environments, may offer a cloud computing option.

Our homes and Internet Service Providers (ISPs) offer web pages, on-line e-mail services, phone and voice mail, and some even offer the ability to run applications on their servers (for a price).

Our businesses have always had web pages, which over time have expanded to offer remote data base access & manipulation. Smaller businesses, in an effort to save licensing and update management issues, have opted to have their providers support them with application servers. Large businesses utilize web-based e-mail servers, data base servers, application servers, Share Point data storage, and now, Voice over Internet Protocol (VoIP).

VoIP offers new capabilities, challenges, and vulnerabilities. Newer protocols such as Session Initiation Protocol (SIP), while involved in communication between end-point devices, offers the potential for Man in the Middle (MITM) attacks similar to those targeting IP.

With more and more Internet-related protocols and services being offered, more businesses have to open more ports on their networks. These can get exploited, such as employees installing Peer-to-Peer (P2P) client applications and harmlessly sharing pictures of the company picnic. Did we say harmless? These same employees open networks and offer sharing of what they perceive as controlled directories and files to those who normally should not have access in any form.

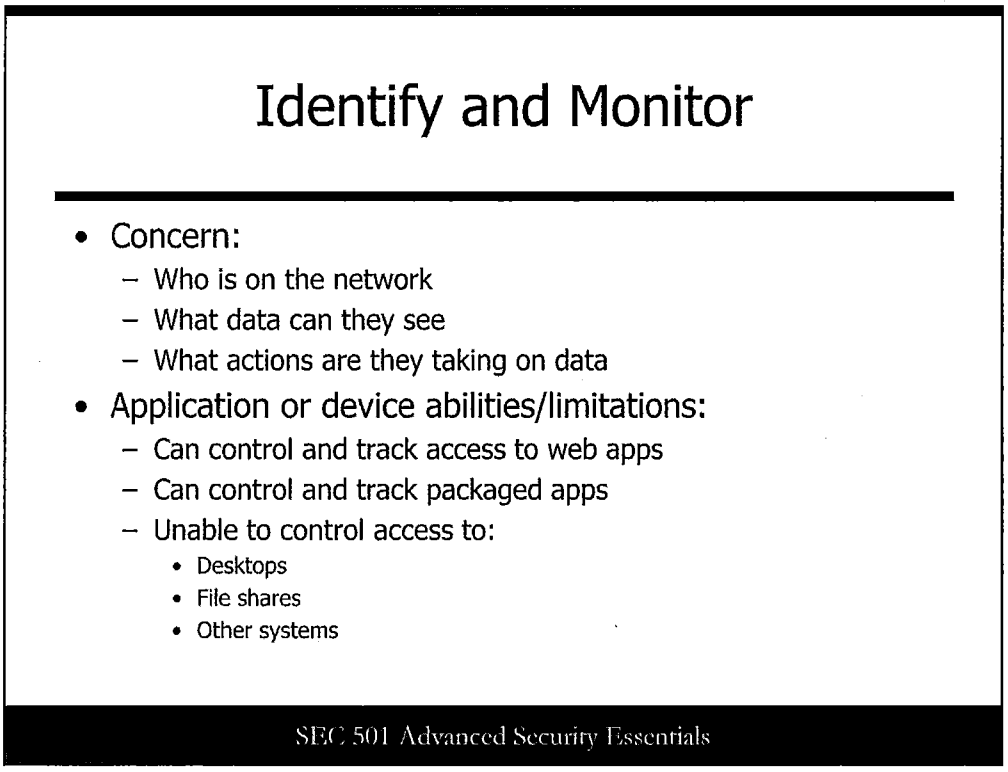

Identifying and monitoring for data loss have several characteristics to help with identifying the measures to use and prevention to incorporate. Monitoring applications and devices must have the ability to:

- Identify users on the network or system.
- Understand and monitor the objects and data each user has access to.
- Have the ability to log actions conducted by users based on their access.
- Identify and alert actions inappropriate with application designs.

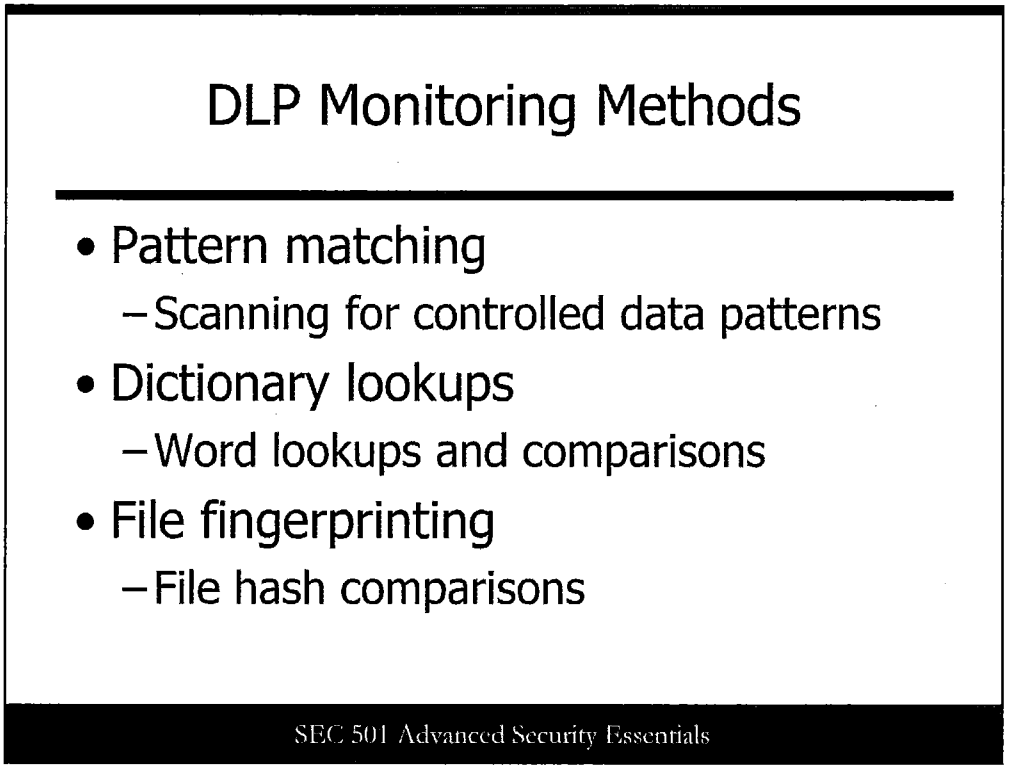

An in-depth look at monitoring methods, includes the ability to generate and/or incorporate normal patterns and compare them with ongoing data patterns and discern any anomalies from normal traffic.

Some monitoring methods use "dictionary lookups," which utilize critical Word files as comparison databases, monitor and view data content as it is accessed or transmitted, compare or search for trigger words, and alert to such action if noticed. This method has its limitations. If data is encrypted, this type of approach does not work, as the monitoring mechanism cannot view the data to search through. Also, if the data does not pass through or by the monitor device, it cannot afford appropriate monitoring of data.

File fingerprinting centers around the creation of file or object hashes and compares them to previous or known hashes of data to be protected. This method can be used for several data loss concerns. One, it can be used to identify data, if used as a hash comparison of data in transit, to identify known activity. Another is to ensure or manage data integrity during manipulation or destructive attacks.

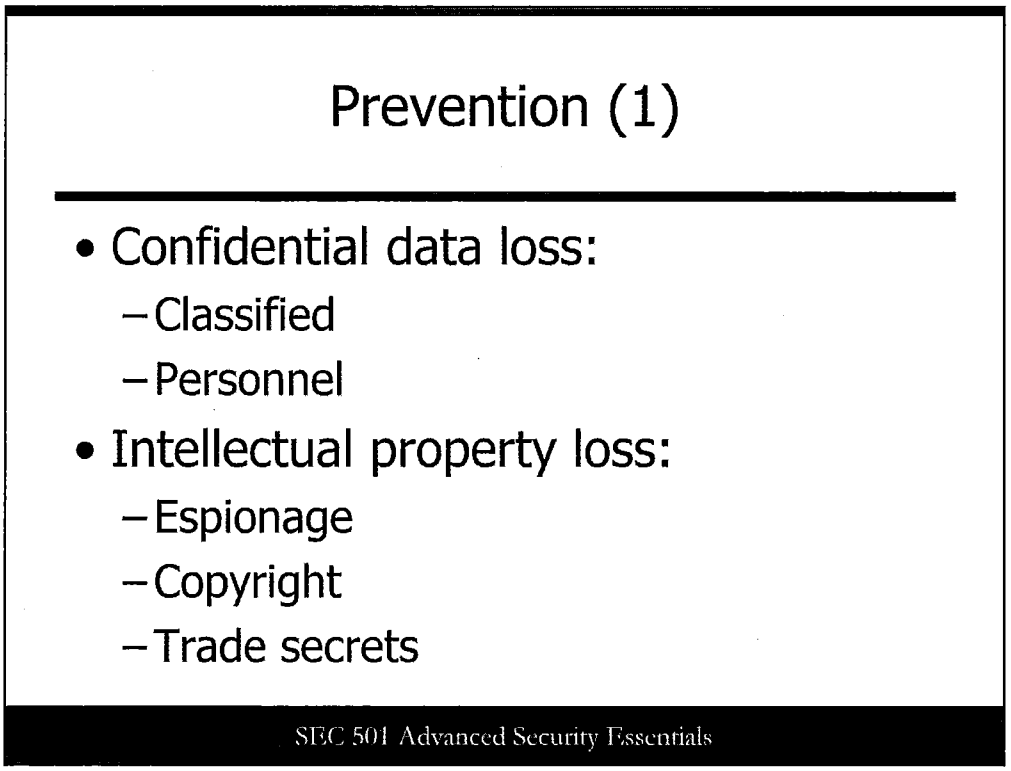

Prevention (1)

Prevention of data loss. This is the overall goal of this section and the methods and characteristics we have been discussing. No one system is 100% secure, nor can we guarantee it will be secure over extended time frames. We can identify data we wish to protect, remove it from data access and communication channels, and lock it away. This is not without its own limitations and vulnerabilities. What of the intent or access of the people we place in charge of its protection? What are their clearances or background check levels?

What we can do is implement ways to reduce or prevent potential data loss through a variety of mechanisms and methods. First, let's classify our data. Classifying information can refer to several different methods:

- **Governmental**: Classified as Top Secret, Secret, Confidential, FOUO, etc.
- **Corporate**: Copyright, Trade Secrets, Proprietary, Registered, etc.
- **Personal**: Sensitive, Close-Hold, Medical, Financial Sensitive, etc.

These classifications allow us to utilize applications that tag or manage and restrict data to specific classification of information in an effort to limit restrict access to authorized users only. Now the preventive measures for authorized users takes over. For instance, are background checks conducted and are roles of individuals, security clearances, medical or mental history, etc. considered?

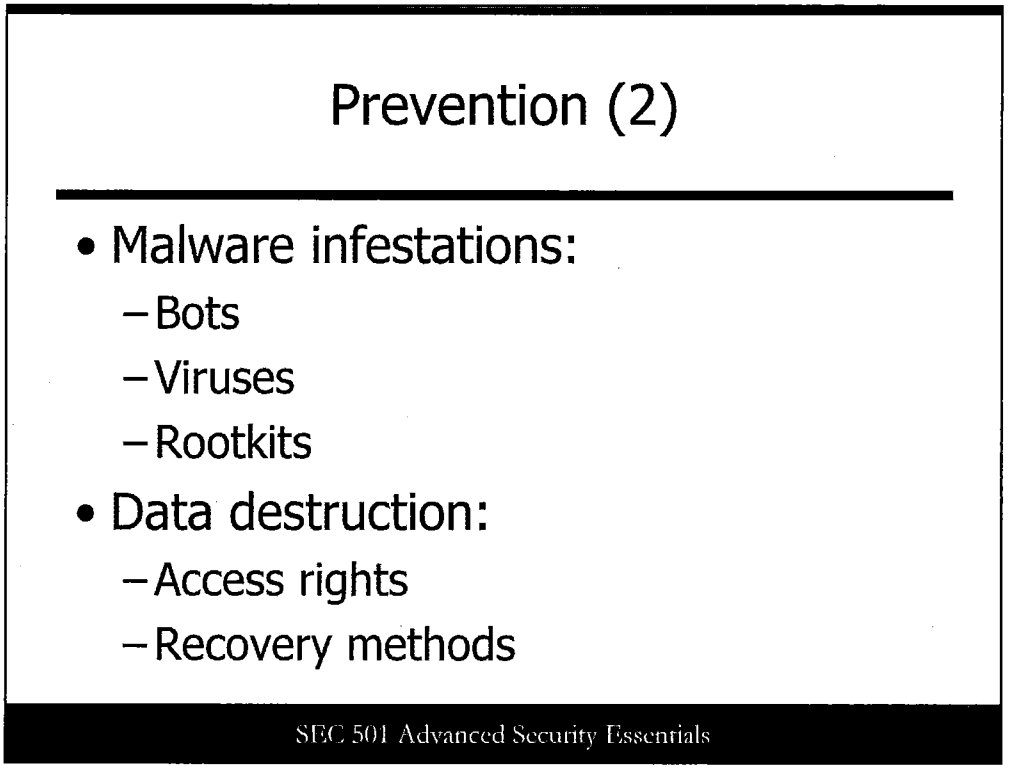

Prevention (2)

Prevention of malware can be handled by several methods. Perimeter defenses such as firewalls, packet filtering screening routers, enterprise wide virus applications, and host based firewalls. All of these offer the opportunity to stop, identify, and eradicate most malware infestations.

Data destruction is normally handled through:

- **Identification / authentication measures**: Prevent direct access to the network.
- **Data encryption**: Protects data, but with write or modify access to the file name, can lead to destruction.
- **Backups / data images**: Offer a means of recovery, not prevention of destruction.

With data loss comes questions that need to be resolved:

Legal:

- Who will bear the responsibility of data loss?
- If responsibilities are based on types of data, how is it addressed and who has responsibilities of each area?
- What steps are being or have been put in place to ensure no further data loss occurs?
- What steps to recovery have to be implemented?
- Who should be notified?

Personnel:

- Include data loss issues in security awareness training.
- Identify roles and responsibilities of data protection.
- Implement and test data classification access and monitor methods.
- Re-assess individual access and access roles and focus on least privileged.

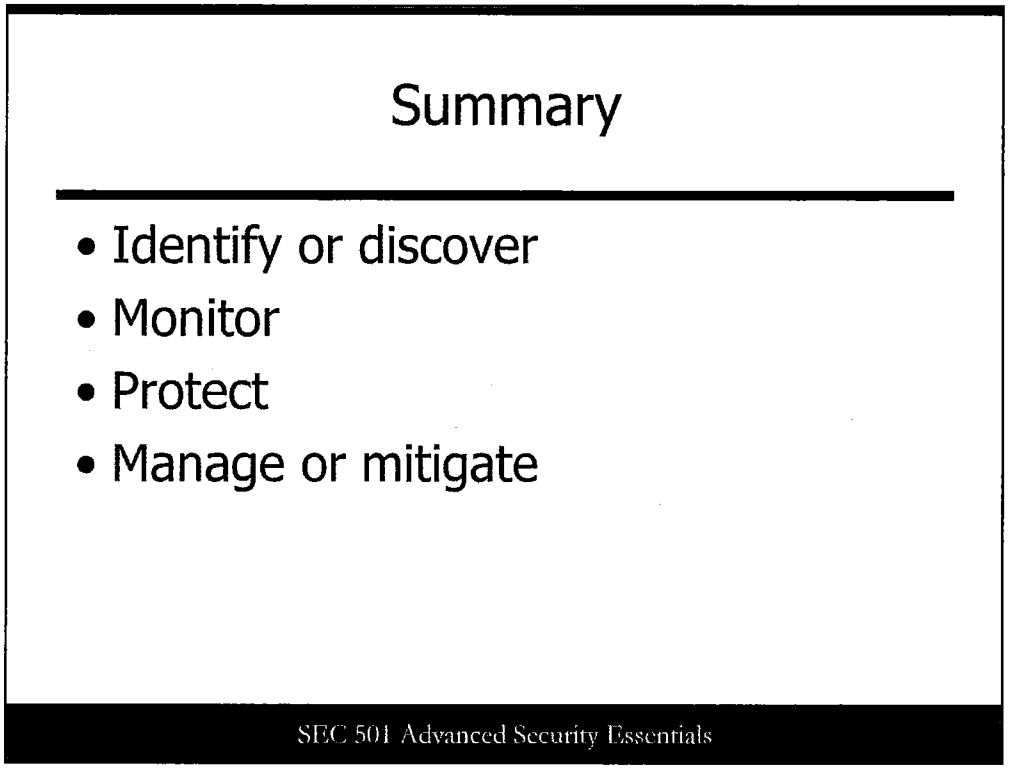

In this section, we discussed methods we can use in the protection and prevention of data loss.

- Identify or discover
- Monitor or audit
- Protect or prevent
- Manage or mitigate