

CISSP

Exam Study Guide



Brian Svidergol

RHEL3, VCP, NCIE-SAN, MCT, MCSE

Table of Contents

Introduction	6
Domain 1. Security and Risk Management	8
1.1 Understand and apply concepts of confidentiality, integrity and availability	8
1.2 Evaluate and apply security governance principles.....	8
1.3 Determine compliance requirements	10
1.4 Understand legal and regulatory issues that pertain to information security in a global context	10
1.5 Understand, adhere to, and promote professional ethics	11
1.6 Develop, document, and implement security policy, standards, procedures and guidelines	12
1.7 Identify, analyze, and prioritize Business Continuity (BC) requirements	12
1.8 Contribute to and enforce personnel security policies and procedures	13
1.9 Understand and apply risk management concepts	14
1.10 Understand and apply threat modeling concepts and methodologies	17
1.11 Apply risk-based management concepts to the supply chain	18
1.12 Establish and maintain a security awareness, education, and training program	19
Domain 1 Review Questions	20
Answers to Domain 1 Review Questions	21
Domain 2. Asset Security	22
2.1 Identify and classify information and assets	22
2.2 Determine and maintain information and asset ownership	23
2.3 Protect privacy	23
2.4 Ensure appropriate asset retention	24
2.5 Determine data security controls	24
2.6 Establish information and asset handling requirements	25
Domain 2 Review Questions	27
Answers to Domain 2 Review Questions	28

Domain 3. Security Architecture and Engineering	29
3.1 Implement and manage engineering processes using secure design principles	29
3.2 Understand the fundamental concepts of security models	30
3.3 Select controls based upon systems security requirements	30
3.4 Understand the security capabilities of information systems	31
3.5 Assess and mitigate the vulnerabilities of security architectures, designs and solution elements	32
3.6 Assess and mitigate vulnerabilities in web-based systems	34
3.7 Assess and mitigate vulnerabilities in mobile systems	34
3.8 Assess and mitigate vulnerabilities in embedded devices	35
3.9 Apply cryptography	35
3.10 Apply security principles to site and facility design	39
3.11 Implement site and facility security controls	39
Domain 3 Review Questions	42
Answers to Domain 3 Review Questions	43
Domain 4. Communication and Network Security	44
4.1 Implement secure design principles in network architecture	44
4.2 Secure network components	46
4.3 Implement secure communication channels according to design	48
Domain 4 Review Questions	49
Answers to Domain 4 Review Questions	50
Domain 5. Identity and Access Management (IAM)	51
5.1 Control physical and logical access to assets	51
5.2 Manage identification and authentication of people, devices and services	52
5.3 Integrate identity as a third-party service	54
5.4 Implement and manage authorization mechanisms	56
5.5 Manage the identity and access provisioning lifecycle	57
Domain 5 Review Questions	59

Answers to Domain 5 Review Questions	60
Domain 6. Security Assessment and Testing	61
6.1 Design and validate assessment, test and audit strategies	61
6.2 Conduct security control testing	61
6.3 Collect security process data	63
6.4 Analyze test output and generate reports	64
6.5 Conduct or facilitate security audits	64
Domain 6 Review Questions	65
Answers to Domain 6 Review Questions	66
Domain 7. Security Operations	67
7.1 Understand and support investigations	67
7.2 Understand the requirements for different types of investigations	68
7.3 Conduct logging and monitoring activities	69
7.4 Securely provision resources	70
7.5 Understand and apply foundational security operations concepts	70
7.6 Apply resource protection techniques	72
7.7 Conduct incident management	73
7.8 Operate and maintain detective and preventative measures	74
7.9 Implement and support patch and vulnerability management	75
7.10 Understand and participate in change management processes	76
7.11 Implement recovery strategies	77
7.12 Implement disaster recovery (DR) processes	78
7.13 Test disaster recovery plans (DRP)	79
7.14 Participate in business continuity (BC) planning and exercises	80
7.15 Implement and manage physical security	81
7.16 Address personnel safety and security concerns	81
Domain 7 Review Questions	83

Answers to Domain 7 Review Questions	84
Domain 8. Software Development Security	85
8.1 Understand and apply security in the software development lifecycle	85
8.2 Enforce security controls in development environments	87
8.3 Assess the effectiveness of software security	88
8.4 Assess security impact of acquired software	88
8.5 Define and apply secure coding guidelines and standards	88
Domain 8 Review Questions	90
Answers to Domain 8 Review Questions	91
Useful References	92
About the Author	93
About Netwrix	93

Introduction

Exam Overview

Preparing to take the Certified Information Systems Security Professional (CISSP) exam requires a great deal of time and effort. The exam covers eight domains:

1. Security and Risk Management
2. Asset Security
3. Security Architecture and Engineering
4. Communication and Network Security
5. Identity and Access Management (IAM)
6. Security Assessment and Testing
7. Security Operations
8. Software Development Security

To qualify to take the exam, you must generally have at least five years of cumulative, paid, full-time work experience in two or more of the eight domains. However, you can satisfy the eligibility requirement with four years of experience in at least two of the eight domains if you have either a four-year college degree or an approved credential or certification. See <https://www.isc2.org/Certifications/CISSP/Prerequisite-Pathway> for a complete list of approved credentials and certifications.

The exam is long, especially compared with other industry certifications. You can take it in English or another language:

- The English language exam is a computerized adaptive testing (CAT) exam, so it changes based on your answers. You get up to 3 hours to complete a minimum of 100 questions and a maximum of 150 questions.
- Exams in languages other than English remain in a linear format. You get up to 6 hours to complete a series of 250 questions.

You must score 700 points or more to pass the exam.

How to Use this Study Guide

Using multiple study sources and methods improves your chances of passing the CISSP exam. For example, instead of reading three or four books, you might read one book, watch a series of videos, take some practice test questions and read a study guide. Or you might take a class, take practice test questions and read a study guide. Or you might join a study group and read a book. The combination of reading, hearing and doing helps your brain process and retain information. If your plan is to read this study guide and then drive over to the exam center, you should immediately rethink your plan!

There are a couple of ways you can use this study guide:

- Use it before you do any other studying — Read it thoroughly. Assess your knowledge as you read. Do you already know everything being said? Or are you finding that you can't follow some of the topics easily? Based on how your reading of the study guide goes, you'll know which exam domains to focus on and how much additional study time you need.
- Use it as the last thing you read prior to taking the exam — Maybe you've taken a class, read a book and gone through a thousand practice test questions, and now you're wondering if you are ready. This study guide might help you answer that question. At a minimum, everything in this study guide should be known to you, make sense to you and not confuse you.

Note that a study guide like this doesn't dive deep enough to teach you a complete topic if you are new to that topic. But it is a very useful preparation tool because it enables you to review a lot of material in a short amount of time. In this guide, we've tried to provide the most important points for each of the topics, but it cannot include the background and details you might find in a 1,000-page book.

Recent Changes to the Exam

On April 15, 2018, the agency that provides the CISSP exam, the International Info System Security Certification Consortium, released an updated set of exam objectives (the exam blueprint). This blueprint is available at <https://www.isc2.org/-/media/ISC2/Certifications/Exam-Outlines/CISSP-Exam-Outline-121417--Final.ashx>.

While most of the exam topics remain the same, there are some minor changes to reflect the latest industry trends and information. Most books for the new version of the exam will be released in May 2018 or later. This study guide has been updated to reflect the new blueprint. The updates are minor: A few small topics have been removed, a few new ones have been added, and some items have been reworded.

What does this mean for you if you are preparing to take the exam? If you have already spent a good amount of time preparing, you might just need to supplement your study with some sources that explain the new and revised material. But if you are just starting to study, consider waiting until the updated guides are released.

Domain 1. Security and Risk Management

1.1 Understand and apply concepts of confidentiality, integrity and availability

Confidentiality, integrity and availability make up what's known as the CIA triad. The CIA triad is a security model that helps organizations stay focused on the important aspects of maintaining a secure environment.

- **Confidentiality.** Sensitive data, including personally identifiable information (PII) like identification numbers and bank account numbers, must be kept confidential. It's important to understand that confidentiality is different from secrecy. If you aren't aware something exists (such as data or a web service), then it is a secret. But keeping something secret, by itself, doesn't ensure confidentiality. You've probably seen stories of attackers (or even regular web surfers) stumbling across "secret" web sites or information, sometimes by accident. To ensure confidentiality, you must make certain that even if someone is aware that something valuable exists (such as a store that processes credit card transactions or a file share with sensitive data), they can't get to it. At a high level, you use access controls — locked doors, folder permissions and two-factor authentication — to maintain confidentiality. At a lower level, you use encryption to protect data at rest, hashing to protect data in motion, and physical security for data in use (privacy screens or physical separation between data in use and unauthorized persons). You can use a "default deny" configuration so that unless somebody has been expressly authorized to access data, they are denied access.
- **Integrity.** You also have to make certain that data isn't changed improperly. Encryption helps ensure the integrity of data at rest, but it isn't the best option for data in motion. Instead, hashing is typically used. Hashing data assigns the data a numeric value, which is calculated at the source before the transfer and then again by the recipient after the transfer; a match proves data integrity. Algorithms such as SHA256 and SHA512 are commonly used for hashing; older algorithms, such as SHA-1, have become susceptible to attack and therefore are rarely used.
- **Availability.** To ensure high availability of services and data, use techniques like failover clustering, site resiliency, automatic failover, load balancing, redundancy of hardware and software components, and fault tolerance. For example, they can help you thwart a denial of service (DoS) attack that aims to deny the availability of a service or data by overloading a system with invalid requests or requests that take a long time to process.

1.2 Evaluate and apply security governance principles

To establish security governance principles, adopt a framework such as the one from the National Institute of Standards and Technology (NIST). Be sure the framework you choose includes the following:

- **Alignment of security function to strategy, goals, mission, and objectives.** An organization has a mission and uses strategy, plans and objectives to try to meet that mission. These components flow down, with the ones below supporting the ones above. Business strategy is often focused 5 or more years out. In the shorter term, typically 1 to 2 years, you have tactical plans that are aligned with the strategic plan. Below that are operational plans — the detailed tactical plans that keep the business running day to day. Objectives are the closest to the ground and represent small efforts to help you achieve a mission. For example, a car manufacturer's mission might be to build and sell as many high-quality cars as possible. The objectives might include expanding automation to reduce the

total build time of a car and expanding from 2 factories to 3. A security framework must closely tie to the organization's mission and objectives, enabling the business to complete its objectives and advance the mission while securing the environment based on risk tolerance. Continuing with the car manufacturer example, the security framework must enable the expansion of automation. If the security framework is such that automation cannot be expanded, then the security framework isn't sufficiently aligned with the mission and objectives.

- **Organizational processes (acquisitions, divestitures, governance committees).** Be aware of the risks in acquisitions (since the state of the IT environment to be integrated is unknown, due diligence is critical) and divestitures (you need to determine how to split the IT infrastructure and what to do with identities and credentials). Understand the value of governance committees (vendor governance, project governance, architecture governance, etc.). Executives, managers and appointed individuals meet to review architecture, projects and incidents (security or otherwise), and provide approvals for new strategies or directions. The goal is a fresh set of eyes, often eyes that are not purely focused on information security.
- **Organizational roles and responsibilities.** There are multiple roles to consider. Management has a responsibility to keep the business running and to maximize profits and shareholder value. The security architect or security engineer has a responsibility to understand the organization's business needs, the existing IT environment, and the current state of security and vulnerability, as well as to think through strategies (improvements, configurations and countermeasures) that could maximize security and minimize risk. There is a need for people who can translate between technical and non-technical people. Costs must be justified and reasonable, based on the organization's requirements and risk.
- **Security control frameworks.** A control framework helps ensure that your organization is covering all the bases around securing the environment. There are many frameworks to choose from, such as Control Objectives for Information Technology (COBIT) and the ISO 27000 series (27000, 27001, 27002, etc.). These frameworks fall into four categories:
 - **Preventative** — Preventing security issues and violations through strategies such as policies and security awareness training
 - **Deterrent** — Discouraging malicious activities using access controls or technologies such as firewalls, intrusion detection systems and motion-activated cameras
 - **Detective** — Uncovering unauthorized activity in your environment
 - **Corrective** — Getting your environment back to where it was prior to a security incident
- **Due care / due diligence.** Ensure you understand the difference between these two concepts. Due care is about your legal responsibility within the law or within organizational policies to implement your organization's controls, follow security policies, do the right thing and make reasonable choices. Due diligence is about understanding your security governance principles (policies and procedures) and the risks to your organization. Due diligence often involves gathering information through discovery, risk assessments and review of existing documentation; creating documentation to establish written policies; and disseminating the information to the organization. Sometimes, people think of due diligence as the method by which due care can be exercised.

After you establish and document a framework for governance, you need security awareness training to bring everything together. All new hires should complete the security awareness training as they come on board, and existing employees should recertify on it regularly (typically yearly).

1.3 Determine compliance requirements

Many organizations need to comply with applicable laws and industry standards. Noncompliance can mean fines, jail time for executives or even the end of a business. To achieve compliance, you must focus on controls. Although most common standards are vague about implementation, a few provide detailed documentation to help organizations achieve compliance. For example, NIST provides a guide for complying with federal information standards.

- **Contractual, legal, industry standards, and regulatory requirements.** Understand the legal systems. Civil law is most common; rulings from judges typically do not set precedents that impact other cases. With common law, which is used in the USA, Canada, the UK and former British colonies, rulings from judges can set precedents that have significant impact on other cases. An example of religious law is Sharia (Islamic law), which use the Qur'an and Hadith for the foundation of laws. Customary law takes common, local and accepted practices and sometimes makes them laws. Within common law, you have criminal law (laws against society) and civil law (typically person vs. person and results in a monetary compensation from the losing party). Compliance factors into laws, regulations, and industry standards such as Sarbanes-Oxley (SOX), the Gramm-Leach-Bliley Act (GLBA), the Payment Card Industry Data Security Standard (PCI DSS), the Health Insurance Portability and Accountability Act (HIPAA), and the Federal Information Security Management Act (FISMA). As part of your exam preparation, familiarize yourself with these standards by reading their high-level summaries.
- **Privacy requirements.** Privacy is about protection of PII. Laws vary. The European Union has tough laws around privacy. Be familiar with the General Data Protection Regulation (GDPR). Be familiar with the requirements around healthcare data, credit card data and other PII data as it relates to various countries and their laws and regulations.

1.4 Understand legal and regulatory issues that pertain to information security in a global context

While you might be familiar with your local legal and regulatory issues, you must be familiar with legal and regulatory issues elsewhere too, at least at a high level.

- **Cyber crimes and data breaches.** Before your organization expands to other countries, perform due diligence to understand their legal systems and what changes might be required to the way that data is handled and secured. In particular, be familiar with the Council of Europe Convention on Cybercrime, a treaty signed by many countries that establishes standards for cybercrime policy. Be familiar with the various laws about data breaches, including notification requirements. In the United States, the Health Information Technology for Economic and Clinical Health (HITECH) Act requires notification of a data breach in some cases, such as when the exposed personal health information was not protected in accordance with the Health Insurance Portability and Accountability Act (HIPAA). The Gramm-Leach-Bliley Act (GLBA) applies to insurance and financial organizations; it requires notification to federal regulators, law enforcement agencies and customers when a data breach occurs. States in the United States also impose their own requirements concerning data breaches. The EU and other countries have their own requirements too. The GDPR has very strict data breach notification requirements: A data breach must be reported to the competent supervisory authority within 72 hours of its discovery. Some countries do not have any reporting requirements.
- **Licensing and intellectual property requirements.** Understand the rules around:

- **Trademarks** — A logo, symbol or mascot used for marketing a brand
- **Patents** — A temporary monopoly for producing a specific item such as a toy, which must be novel and unique to qualify for a patent
- **Copyright** — Exclusive use of artistic, musical or literary works that prevents unauthorized duplication, distribution or modification)
- **Licensing** — A contract between the software producer and the consumer that limits the use and/or distribution of the software
- **Import/export controls.** Every country has laws around the import and export of hardware and software. For example, the United States has restrictions around the export of cryptographic technology, and Russia requires a license to import encryption technologies manufactured outside the country.
- **Trans-border data flow.** If your organization adheres to specific security laws and regulations, then you should adhere to them no matter where the data resides — for example, even if you store a second copy of your data in another country. Be aware of the applicable laws in all countries where you store data and maintain computer systems. In some cases, data might need to remain in the country. In other cases, you need to be careful with your data because the technical teams might be unaware of the security and compliance requirements. The EU-US Privacy Shield (formerly the EU-US Safe Harbor agreement) controls data flow from the EU to the United States. The EU has more stringent privacy protections and without the Safe Harbor act, personal data flow from the EU to the United States would not be allowed.
- **Privacy.** Many laws include privacy protections for personal data. The new GDPR has strong privacy rules that apply to any organization anywhere that stores or processes the personal data of EU citizens; in particular, individuals must be told how their data is collected and used, and they must be able to opt out. The privacy guidelines of the Organization for Economic Co-operation and Development (OECD) require organizations to avoid unjustified obstacles to trans-border data flow, limit personal data collection, protect personal data with reasonable security, and more.

1.5 Understand, adhere to, and promote professional ethics

As a CISSP, you must understand and follow the (ISC)² code of ethics, as well as your organization's own code.

- **(ISC)² Code of Professional Ethics.** Take the time to read the code of ethics available at www.isc2.org/Ethics. At a minimum, know and understand the ethics canons:
 - **Protect society, the common good, necessary public trust and confidence, and the infrastructure.** This is “do the right thing.” Put the common good ahead of yourself. Ensure that the public can have faith in your infrastructure and security.
 - **Act honorably, honestly, justly, responsibly, and legally.** Always follow the laws. But what if you find yourself working on a project where conflicting laws from different countries or jurisdictions apply? In such a case, you should prioritize the local jurisdiction from which you are performing the services.
 - **Provide diligent and competent service to principles.** Avoid passing yourself as an expert or as qualified in areas that you aren't. Maintain and expand your skills to provide competent services.

- **Advance and protect the profession.** Don't bring negative publicity to the profession. Provide competent services, get training and act honorably. Think of it like this: If you follow the first three canons in the code of ethics, you automatically comply with this one.
- **Organizational code of ethics.** You must also support ethics at your organization. This can be interpreted to mean evangelizing ethics throughout the organization, providing documentation and training around ethics, or looking for ways to enhance the existing organizational ethics. Some organizations might have slightly different ethics than others, so be sure to familiarize yourself with your organization's ethics and guidelines.

1.6 Develop, document, and implement security policy, standards, procedures and guidelines

Develop clear security policy documentation, including the following:

- **Policy.** This is the high-level document, often written by the management team. Policy is mandatory. It is purposely vague. For example, a policy might require you to ensure the confidentiality of company data but not specify the method for doing so.
- **Standards.** These are more descriptive than policies and document the standards to be used by the company for things such as hardware and software. For example, an organization might standardize on virtual machines and not physical servers.
- **Procedures.** These are the step-by-step documents that detail how to perform specific tasks, such as how to restore a database. The person following the procedure uses the document to perform the task. Procedures are mandatory. If you have a procedure for restoring a database, then that procedure needs to be followed for every database restore.
- **Guidelines.** These are recommended but optional. For example, your organization might have a guideline that recommends storing passwords in an encrypted password vault. It is a good idea to do that. But somebody might choose to store passwords in their brain or using another secure storage mechanism.
- **Baselines.** Although baselines are not explicitly mentioned in this section of the exam, don't forget about them. Baselines automate implementation of your standards, thereby ensuring adherence to them. For example, if you have 152 configuration items for your server builds, you can configure all of them in a baseline that is applied to every server that is built. Group Policy objects (GPOs) are often used to comply with standards in a Windows network. Configuration management solutions can also help you establish baselines and spot configurations that drift away from them.

1.7 Identify, analyze, and prioritize Business Continuity (BC) requirements

Business continuity is the goal of remaining fully operational during an outage. ISO/IEC 27031 covers business continuity in detail (it provides a framework to build on, along with methods and processes covering the entire subject). Business continuity requires a lot of planning and preparation. The actual implementation of business continuity processes occurs quite infrequently. The primary facets of business continuity are resilience (within a data center and between sites or data

centers), recovery (if a service becomes unavailable, you need to recover it as soon as possible), and contingency (a last resort in case resilience and recovery prove ineffective).

- **Develop and document scope and plan.** Developing the project scope and plan starts with gaining support of the management team, making a business case (cost/benefit analysis, regulatory or compliance reasons, etc.), and ultimately gaining approval to move forward. Next, you need to form a team with representatives from the business as well as IT. Then you are ready to begin developing the plan. Start with a business continuity policy statement, then conduct a business impact analysis (as explained in the next bullet), and then develop the remaining components: preventive controls, relocation, the actual continuity plan, testing, training and maintenance). Be familiar with the difference between business continuity (resuming critical functions without regard for the site) and disaster recovery (recovering critical functions at the primary site, when possible).
- **Conduct a business impact analysis (BIA).** Identify the systems and services that the business relies on and figure out the impacts that a disruption or outage would cause, including the impacts on business processes like accounts receivable and sales. You also need to figure out which systems and services you need to get things running again (think foundational IT services such as the network and directory, which many other systems rely on). Be sure to prioritize the order in which critical systems and services are recovered or brought back online. As part of the BIA, you will establish the recovery time objectives (RTOs) (how long it takes to recover), the recovery point objectives (RPOs) (the maximum tolerable data loss), and maximum tolerable downtime (MTD), along with the costs of downtime and recovery.

1.8 Contribute to and enforce personnel security policies and procedures

In many organizations, the number one risk to the IT environment is people. And it's not just IT staff, but anyone who has access to the network. Malicious actors routinely target users with phishing and spear phishing campaigns, social engineering, and other types of attacks. Everybody is a target. And once attackers compromise an account, they can use that entry point to move around the network and elevate their privileges. The following strategies can reduce your risk:

- **Candidate screening and hiring.** Screening candidates thoroughly is a critical part of the hiring process. Be sure to conduct a full background check that includes a criminal records check, job history verification, education verification, certification validation and confirmation of other accolades when possible. Additionally, contact all references.
- **Employment agreements and policies.** An employment agreement specifies job duties, expectations, rate of pay, benefits and information about termination. Sometimes, such agreements are for a set period (for example, in a contract or short-term job). Employment agreements facilitate termination when needed for an underperforming employee. The more information and detail in an employment agreement, the less risk (risk of a wrongful termination lawsuit, for example) the company has during a termination proceeding. For instance, a terminated employee might take a copy of their email with them without thinking of it as stealing, but they are less likely to do so if an employment agreement or another policy document clearly prohibits it.
- **Onboarding and termination processes.** Onboarding comprises all the processes tied to a new employee starting at your organization. Having a documented process in place enables new employees to be integrated as quickly and consistently as possible, which reduces risk. For example, if you have five IT admins performing the various onboarding processes, you might get different results each time if you don't have the processes

standardized and documented; a new hire might end up with more access than required for their job. Termination is sometimes a cordial process, such as when a worker retires after 30 years. Other times, it can be a high-stress situation, such as when a person is being terminated unexpectedly. You need to have documented policies and procedures to handle all termination processes. The goal is to have a procedure to immediately revoke all access to all company resources. In a perfect world, you would push one button and all access would be revoked immediately.

- **Vendor, consultant, and contractor agreements and controls.** When workers who are not full-time employees have access to your network and data, you must take extra precautions. Consultants often work with multiple customers simultaneously, so you need to have safeguards in place to ensure that your company's data isn't mixed in with data from other organizations, or accidentally or deliberately transmitted to unauthorized people. In high-security organizations, it is common to have the organization issue a computing device to consultants and enable the consultant to access the network and data only through that device. Beyond the technical safeguards, you must also have a way to identify consultants, vendors and contractors. For example, maybe they have a different security badge than regular full-time employees. Perhaps they sit in the same area or their display names in the directory call out their status.
- **Compliance policy requirements.** Organizations have to adhere to different compliance mandates, depending on their industry, country and other factors. All of them need to maintain documentation about their policies and procedures for meeting those requirements. Employees should be trained on the company's compliance mandates at a high level upon hire and regularly thereafter (such as re-certifying once a year).
- **Privacy policy requirements.** Personally identifiable information about employees, partners, contractors, customers and other people should be stored in a secure way, accessible only to those who require the information to perform their jobs. For example, somebody in the Payroll department might need access to an employee's banking information to have their pay automatically deposited, but no one else should be able to access that data. Organizations should maintain a documented privacy policy that outlines the types of data covered by the policy and who the policy applies to. Employees, contractors and anyone else who might have access to the data should be required to read and agree to the privacy policy upon hire and on a regular basis thereafter (such as annually).

1.9 Understand and apply risk management concepts

Risk management involves three primary steps: identify threats and vulnerabilities, assess the risk (risk assessment), and choose whether and how to respond (often the choice is risk mitigation). As part of managing overall risk, the IT team strives to secure the IT environment, provide information to the management teams so that they can make informed decisions, and enable the management team to sign off on the IT environment based on the goals and requirements. Risk management also has a financial component: The management team must balance the risk with the budget. In a perfect world, the company would spend the minimum amount of money and time to minimize risk to an acceptable level for the organization.

- **Identify threats and vulnerabilities.** Threats and vulnerabilities are linked. A threat (such as a hacker taking over a client computer) is possible when a vulnerability (such as an unpatched client computer) is present. That is a known threat. But unknown threats also exist, such as when a hacker is aware of a bug that nobody else knows about in your anti-virus software and can remotely compromise your computer.

- **Assess risk.** You have a risk when you have a threat and a vulnerability. In those cases, you need to figure out the chances of the threat exploiting the vulnerability and the consequences if that does happen. Be familiar with the approaches:
 - **Qualitative.** This method uses a risk analysis matrix and assigns a risk value such as low, medium or high. For example, if the likelihood is rare and the consequences are low, then the risk is low. If the likelihood is almost certain and the consequences are major, then the risk is extreme.
 - **Quantitative.** This method is more objective than the qualitative method; it uses dollars or other metrics to quantify risk.
 - **Hybrid.** A mix of qualitative and quantitative. If you can easily assign a dollar amount, you do; if not, you don't. This can often provide a good balance between qualitative and quantitative.
- **Respond to risk.** You must formulate a plan of action for each risk you identify. For a given risk, you can choose risk mitigation (reduce the risk), risk assignment (assign the risk to a team or provider for action), risk acceptance (accept the risk) or risk rejection (ignore the risk).

Outside of the three primary steps for applying risk management, you should familiarize yourself with some of the details for those three steps:

- **Countermeasure selection and implementation.** You can use a software or hardware solution to reduce a particular risk by implementing a countermeasure, sometimes referred to as a "control" or a "safeguard." Suppose you have a password policy that a legacy application cannot technically meet (for example, the app is limited to 10 characters for the password). To reduce the likelihood of that password being compromised, you can implement any of several countermeasures: For instance, you can require that the password be changed more frequently than other (longer) passwords, or mandate that the password be stored in a secure password vault that requires two-factor authentication. For your exam preparation, don't just understand the words and definitions; understand how you implement the concepts in your environment. You don't have to provide a step-by-step technical configuration, but you must understand the implementation process — where you start, the order of the steps you take and how you finish.
- **Applicable types of controls.** Be familiar with the 6 types of controls:
 - **Preventive.** This type of control is intended to prevent a security incident from happening. For example, you add an anti-virus product to your computers.
 - **Detective.** This type of control is used to identify the details of a security incident, including (sometimes) the attacker.
 - **Corrective.** A corrective control implements a fix after a security incident occurs.
 - **Deterrent.** This type of control attempts to discourage attackers. For example, you lock your office whenever you leave for lunch or go home for the day.
 - **Recovery.** A recovery control tries to get the environment back to where it was prior to a security incident.
 - **Compensating.** A compensating control is an alternative control to reduce a risk. Suppose you need to enable outside users to get to your SharePoint site, which resides on your local area network. Instead of

opening the firewall to permit communication from the internet to your internal SharePoint servers, you can implement a compensating control, such as deploying a reverse proxy to the perimeter network and enabling SharePoint external access through the reverse proxy. In the end, the functionality is typically the same, but the method of getting there is different.

- **Security Control Assessment (SCA).** You need to periodically assess your security controls. What's working? What isn't working? As part of this assessment, the existing document must be thoroughly reviewed, and some of the controls must be tested at random. A report is typically produced to show the outcomes and enable the organization to remediate deficiencies.
- **Monitoring and measurement.** Monitoring and measurement are closely aligned with identifying risks. For example, if there are many invalid database query attempts coming from your web server, it might indicate an attack. At a minimum, it is worth investigating. Whether action is required will depend. Without the proper monitoring in place, you won't know about these types of events. You might not know when a person is probing your network. Even if you are capturing monitoring information, it isn't enough by itself. You also need a way to measure it. For example, if your monitoring shows 500 invalid logon attempts on your web server today, is that a cause for concern? Or is that typical because you have 75,000 users? While monitoring is used for more than security purposes, you need to tune it to ensure you are notified about potential security incidents as soon as possible. In some cases, it will be too late and a data breach might occur. That's when the monitoring data becomes valuable from a forensics perspective. You need to be able to look back at the data and figure out why you didn't see anything during the incident and what adjustments you need to make to minimize the chances of it happening again.
- **Asset valuation.** When you think of assets, don't just think of physical assets such as computers and office furniture (tangible assets). Assets also include the company's data and intellectual property (intangible assets). While tangible assets are easy to assess for value (for example, you bought the disk drive for \$250), data and intellectual property can be harder to place a value on. Be familiar with the following strategies of intangible asset valuation:
 - **Cost approach.** How much would it cost to replace the asset?
 - **Income approach.** How much income will the asset produce over its lifetime?
 - **Market approach.** How much does a similar asset cost?
 - **Quantitative approach.** Assigns a dollar value to assess risk.
 - **Qualitative approach.** Assigns a score to assess risk.
- **Reporting.** One of the foundations of an enterprise-grade security solution is the ability to report on your environment (what you have, what the top risks are, what's happening right now, what happened 3 days ago, etc.). Reporting provides information. And that information is sometimes used to start a continuous improvement process.
- **Continuous improvement.** Continuous improvement is an ongoing, never-ending effort to take what you have and improve it. Often, improvements are small and incremental. However, over time, small improvements can add up. Continuous improvement can be applied to products (for example, upgrading to the latest version), services

(for example, expanding your internal phishing testing) or processes (for example, automating processes to save time and improve consistency).

- **Risk frameworks.** A risk framework documents how your organization handles risk assessment, risk resolution and ongoing monitoring. See <http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-37r1.pdf> for an example of a risk framework. There are other risk frameworks, such as the British Standard BS 31100. Be familiar with risk frameworks and their goals. The NIST framework identifies the following steps: categorize, select, implement, assess, authorize and monitor.

1.10 Understand and apply threat modeling concepts and methodologies

When you perform threat modeling for your organization, you document potential threats and prioritize those threats (often by putting yourself in an attacker's mindset). There are four well-known methods. STRIDE, introduced at Microsoft in 1999, focuses on spoofing of user identity, tampering, repudiation, information disclosure, denial of service and elevation of privilege. PASTA (process for attack simulation and threat analysis) provides dynamic threat identification, enumeration and scoring. Trike uses threat models based on a requirements model. VAST (visual, agile and simple threat modeling) applies across IT infrastructure and software development without requiring security experts.

- **Threat modeling methodologies.** Part of the job of the security team is to identify threats. You can identify threats using different methods:
 - **Focus on attackers.** This is a useful method in specific situations. For example, suppose that a developer's employment is terminated. After extracting data from the developer's computer, you determine that the person was disgruntled and angry at the management team. You now know this person is a threat and can focus on what he or she might want to achieve. However, outside of specific situations like this, organizations are usually not familiar with their attackers.
 - **Focus on assets.** Your organization's most valuable assets are likely to be targeted by attackers. For example, if you have a large number of databases, the database with the HR and employee information might be the most sought after.
 - **Focus on software.** Many organizations develop applications in house, either for their own use or for customer use. You can look at your software as part of your threat identification efforts. The goal isn't to identify every possible attack, but instead to focus on the big picture, such as whether the applications are susceptible to DoS or information disclosure attacks.
- **Threat modeling concepts.** If you understand the threats to your organization, then you are ready to document the potential attack vectors. You can use diagramming to list the various technologies under threat. For example, suppose you have a SharePoint server that stores confidential information and is therefore a potential target. You can diagram the environment integrating with SharePoint. You might list the edge firewalls, the reverse proxy in the perimeter network, the SharePoint servers in the farm and the database servers. Separately, you might have a diagram showing SharePoint's integration with Active Directory and other applications. You can use these diagrams to identify attack vectors against the various technologies.

1.11 Apply risk-based management concepts to the supply chain

Organizations must use risk-based management concepts when they contract out tasks (such as hiring an air conditioning company to maintain the air conditioning in their data centers), bring on new suppliers or utilize service companies to transport their goods. Many of these concepts apply to mergers and acquisitions too.

- **Risks associated with hardware, software, and services.** The company should perform due diligence, which includes looking at the IT infrastructure of the supplier. When thinking about the risk considerations, you must consider:
 - **Hardware.** Is the company using antiquated hardware that introduces potential availability issues? Is the company using legacy hardware that isn't being patched by the vendor? Will there be integration issues with the hardware?
 - **Software.** Is the company using software that is out of support, or from a vendor that is no longer in business? Is the software up to date on security patches? Are there other security risks associated with the software?
 - **Services.** Does the company provide services for other companies or to end users? Is the company reliant on third-party providers for services (such as SaaS apps)? Did the company evaluate service providers in a way that enables your company to meet its requirements? Does the company provide services to your competitors? If so, does that introduce any conflicts of interest?
- **Third-party assessment and monitoring.** Before agreeing to do business with another company, your organization needs to learn as much as it can about that company. Often, third-party assessments are used to help gather information and perform the assessment. An on-site assessment is useful to gain information about physical security and operations. During the document review, your goal is to thoroughly review all the architecture, designs, implementations, policies, procedures, etc. You need to have a good understanding of the current state of the environment, especially so you can understand any shortcomings or compliance issues prior to integrating the IT infrastructures. You need to ensure that the other company's infrastructure meets all your company's security and compliance requirements. The level of access and depth of information you are able to gain is often directly related to how closely your companies will work together. For example, if a company is your primary supplier of a critical hardware component, then a thorough assessment is critical. If the company is one of 3 delivery companies used to transport goods from your warehouse, then the assessment is important but does not have to be as deep.
- **Minimum security requirements.** As part of the assessment, the minimum security requirements must be established. In some cases, the minimum security requirements are your company's security requirements. In other cases, new minimum security requirements are established. In such scenarios, the minimum security requirements should have a defined period, such as 12 months.
- **Service-level requirements.** A final area to review involves service level agreements (SLAs). Companies have SLAs for internal operations (such as how long it takes for the helpdesk to respond to a new ticket), for customers (such as the availability of a public-facing service), and for partner organizations (such as how much support a vendor provides a partner). All the SLAs of the company should be reviewed. Your company sometimes has an SLA standard that should be applied, when possible, to the SLAs as part of working with another company. This can

sometimes take time, as the acquiring company might have to support established SLAs until they expire or renewal comes up.

1.12 Establish and maintain a security awareness, education, and training program

This section of the exam covers all the aspects of ensuring that everybody in your organization is security conscious and familiar with the organization's policies and procedures. In general, it is most effective to start with an awareness campaign and then provide detailed training. For example, teaching everybody about malware or phishing campaigns before they understand the bigger picture of risk isn't very effective.

- **Methods and techniques to present awareness and training.** While the information security team is typically well-versed on security, the rest of the organization often isn't. As part of having a well-rounded security program, the organization must provide security education, training and awareness to the entire staff. Employees need to understand what to be aware of (types of threats, such as phishing or free USB sticks), understand how to perform their jobs securely (encrypt sensitive data, physically protect valuable assets), and how security plays a role in the big picture (company reputation, profits and losses). Training should be mandatory and provided both to new employees and yearly (at a minimum) for ongoing training. Routine tests of operational security should be performed (such as tailgating at company doors and social engineering tests like phishing campaigns).
- **Periodic content reviews.** Threats are complex and the training needs to be relevant and interesting to be effective. This means updating training materials and awareness training, and changing out the ways which security is tested and measured. If you always use the same phishing test campaign or send it from the same account on the same day of the year, it isn't effective. The same applies to other material. Instead of relying on long and detailed security documentation for training and awareness, consider using internal social media tools, videos and interactive campaigns.
- **Program effectiveness evaluation.** Time and money must be allocated for evaluating the company's security awareness and training. The company should track key metrics, such as the percentage of employees clicking on a link in a test phishing email. Is the awareness and training bringing the total number of clicks down? If so, the program is effective. If not, you need to re-evaluate it.

Domain 1 Review Questions

Read and answer the following questions. If you do not get at least one of them correct, spend more time with the subject. Then move on to Domain 2.

1. You are a security consultant. A large enterprise customer hires you to ensure that their security operations are following industry standard control frameworks. For this project, the customer wants you to focus on technology solutions that will discourage malicious activities. Which type of control framework should you focus on?
 - a. Preventative
 - b. Deterrent
 - c. Detective
 - d. Corrective
 - e. Assessment

2. You are performing a risk analysis for an internet service provider (ISP) that has thousands of customers on its broadband network. Over the past 5 years, some customers have been compromised or experienced data breaches. The ISP has a large amount of monitoring and log data for all customers. You need to figure out the chances of additional customers experiencing a security incident based on that data. Which type of approach should you use for the risk analysis?
 - a. Qualitative
 - b. Quantitative
 - c. STRIDE
 - d. Reduction
 - e. Market

3. You are working on a business continuity project for a company that generates a large amount of content each day for use in social networks. Your team establishes 4 hours as the maximum tolerable data loss in a disaster recovery or business continuity event. In which part of the business continuity plan should you document this?
 - a. Recovery time objective (RTO)
 - b. Recovery point objective (RPO)
 - c. Maximum tolerable downtime (MTD)
 - d. Maximum data tolerance (MDT)

Answers to Domain 1 Review Questions

1. Answer: B

Explanation: Deterrent frameworks are technology-related and used to discourage malicious activities. For example, an intrusion prevention system or a firewall would be appropriate in this framework.

There are three other primary control frameworks. A preventative framework helps establish security policies and security awareness training. A detective framework is focused on finding unauthorized activity in your environment after a security incident. A corrective framework focuses on activities to get your environment back after a security incident. There isn't an assessment framework.

2. Answer: B

Explanation: You have three risk analysis methods to choose from: qualitative (which uses a risk analysis matrix), quantitative (which uses money or metrics to compute), or hybrid (a combination of qualitative and quantitative but not an answer choice in this scenario). Because the ISP has monitoring and log data, you should use a quantitative approach; it will help quantify the chances of additional customers experiencing a security risk.

STRIDE is used for threat modeling. A market approach is used for asset valuation. A reduction analysis attempts to eliminate duplicate analysis and is tied to threat modeling.

3. Answer: B

Explanation: The RTO establishes the maximum amount of time the organization will be down (or how long it takes to recover), the RPO establishes the maximum data loss that is tolerable, the MTD covers the maximum tolerable downtime, and MDT is just a made-up phrase used as a distraction. In this scenario, with the focus on the data loss, the correct answer is RPO.

Domain 2. Asset Security

When we think about assets, some people consider only physical assets, such as buildings, land and computers. But asset security for the CISSP exam focuses on virtual assets such as intellectual property and data. (In Domain 3, there will be some physical security topics.)

2.1 Identify and classify information and assets

To improve security, you need to identify both your data and your physical assets and classify them according to their importance or sensitivity, so you can specify procedures for handling them appropriately based on their classification.

- **Data classification.** Organizations classify their data using labels. You might be familiar with two government classification labels, Secret and Top Secret. Non-government organizations generally use classification labels such as Public, Internal Use Only, Partner Use Only, or Company Confidential. However, data classification can be more granular; for example, you might label certain information as HR Only.
- **Asset classification.** You also need to identify and classify physical assets, such as computers, smartphones, desks and company cars. Unlike data, assets are typically identified and classified by asset type. Often, asset classification is used for accounting purposes, but it can also be tied to information security. For example, an organization might designate a set of special laptops with particular software installed, and assign them to employees when they travel to high-risk destinations, so their day-to-day assets can remain safely at home.

Classification labels help users disseminate data and assets properly. For example, if Sue has a document classified as Partner Use Only, she knows that it can be distributed only to partners; any further distribution is a violation of security policy. In addition, some data loss prevention solutions can use classification data to help protect company data automatically. For example, an email server can prevent documents classified as Internal Use Only from being sent outside of the organization.

People with the right clearance can view certain classifications of data or check out certain types of company equipment (such as a company truck). While clearance is often associated with governments or the military, it is also useful for organizations. Some organizations use it routinely throughout their environments, while other organizations use it for special scenarios, such as a merger or acquisition. When studying for this section, concentrate on understanding the following concepts:

- **Clearance.** Clearance dictates who has access to what. Generally, a certain clearance provides access to a certain classification of data or certain types of equipment. For example, Secret clearance gives access to Secret documents, and a law enforcement organization might require a particular clearance level for use of heavy weaponry.
- **Formal access approval.** Whenever a user needs to gain access to data or assets that they don't currently have access to, there should be a formal approval process. The process should involve approval from the data owner, who should be provided with details about the access being requested. Before a user is granted access to the data, they should be told the rules and limits of working with it. For example, they should be aware that they must not send documents outside the organization if they are classified as Internal Only.

- **Need to know.** Suppose your company is acquiring another company but it hasn't been announced yet. The CIO, who is aware of the acquisition, needs to have IT staff review some redacted network diagrams as part of the due diligence process. In such a scenario, the IT staff is given only the information they need to know (for example, that it is a network layout and the company is interested in its compatibility with its own network). The IT staff do not need to know about the acquisition at that time. This is "need to know."

2.2 Determine and maintain information and asset ownership

If you don't know who owns a piece of data, how can you go through a formal access approval process? You can't, at least not as effectively. Similarly, you can't properly account for assets if you don't know which department owns them, or assign the right type of laptop for high-risk travel if you don't have the assets classified.

Data owners are responsible for classifying the data they own. In larger companies, an asset management department handles asset classification. A custodian is a hands-on role that implements and operates solutions for data (e.g., backups and restores). A system owner is responsible for the computer environment (hardware, software) that houses data; this is typically a management role with operational tasks handed off to the custodian.

2.3 Protect privacy

All workers need to be aware of the company's privacy policies and procedures and know how to contact data owners in the event of an issue. Key terms to understand include the following:

- **Data owners.** Data owners are usually members of the management or senior management team. They approve access to data (usually by approving the data access policies that are used day to day).
- **Data processors.** Data processors are the users who read and edit the data regularly. Users must clearly understand their responsibilities with data based on its classification. Can they share it? What happens if they accidentally lose it or destroy it?
- **Data remanence.** Data remanence occurs when data is deleted but remains recoverable. Whenever you delete a file, the operating system marks the space the file took up as available. But the data is still there, and with freely downloadable tools, you can easily extract that data. Organizations need to account for data remanence to ensure they are protecting their data. There are a few options:
 - **Secure deletion or overwriting of data.** You can use a tool to overwrite the space that a file was using with random 1s and 0s, either in one pass or in multiple passes. The more passes you use, the less likely it is that the data can be recovered.
 - **Destroying the media.** You can shred disk drives, smash them into tiny pieces, or use other means to physically destroy them. This is effective but renders the media unusable thereafter.
 - **Degaussing.** Degaussing relies on the removal or reduction of magnetic fields on the disk drives. It is very effective and complies with many government requirements for data remanence.

- **Collection limitation.** Security often focuses on protecting the data you already have. But part of data protection is limiting how much data your organization collects. For example, if you collect users' birthdates or identification card numbers, you then must protect that data. If your organization doesn't need the data, it shouldn't collect it. Many countries are enacting laws and regulations to limit the collection of data. But many organizations are unaware and continue to collect vast amounts of sensitive data. You should have a privacy policy that specifies what information is collected, how it is used and other pertinent details.

2.4 Ensure appropriate asset retention

There are two aspects to data retention: You should ensure that your organization holds data for as long as required — and also that it securely deletes data that is no longer required, in order to reduce the risk of its exposure.

To determine how long to keep certain data, you need to consider both whether the data is still useful to your organization and whether there are any regulations, legal reasons or company policies requiring its retention. In many cases, a company must keep data for longer than the data provides value; for example, your organization might have a policy to retain email data for 7 years regardless of its value. As part of your comprehensive security policies, you should ensure the destruction of unneeded data.

Besides data, this section also covers the hardware and personnel required to use the data. These are quite important.

- **Hardware.** Even if you maintain data for the appropriate retention period, it won't do you any good if you don't have hardware that can read the data. For example, if you have data on backup tapes and hold them for 10 years, you run the risk of not being able to read the tapes toward the end of the retention period because tape hardware changes every few years. Thus, you must ensure you have the hardware and related software (tape drives, media readers and so on) needed to get to the data that you are saving.
- **Personnel.** Suppose your company is retaining data for the required time periods and maintaining hardware to read the data. But what happens if the only person who knew how to operate your tape drives and restore data from them no longer works at the company, and the new team is only familiar with disk-to-disk backup? You might not be able to get to your data! By documenting all the procedures and architecture, you can minimize this risk.

2.5 Determine data security controls

You need data security controls that protect your data as it is stored, used and transmitted.

- **Understanding data states.** The industry identifies three data states:
 - **Data at rest** is data stored on a storage medium (disk, tape, etc.).
 - **Data in motion** is data moving from a source (such as a computer) to a destination (such as another computer).
 - **Data in use** is data that is actively being worked on (for example, a person editing a spreadsheet).
- **Scoping and tailoring.** Scoping is the process of finalizing which controls are in scope and which are out of scope (not applicable). Tailoring is the process of customizing the implementation of controls for an organization.

- **Standards selection.** Standards selection is the process by which organizations plan, choose and document technologies and/or architectures for implementation. For example, you might evaluate three vendors for an edge firewall solution. You could use a standards selection process to help determine which solution best fits the organization. Vendor selection is closely related to standards selection but focuses on the vendors, not the technologies or solutions. The overall goal is to have an objective and measurable selection process. If you repeat the process with a totally different team, then they should come up with the same selection as the first team. In such a scenario, you would know that your selection process is working as expected.
- **Data protection methods.** The options for protecting data depend on its state:
 - **Data at rest.** You can encrypt data at rest. You should consider encryption for operating system volumes and data volumes, and you should encrypt backups, too. Be sure to consider all locations for data at rest, such as tapes, USB drives, external drives, RAID arrays, SAN, NAS and optical media.
 - **Data in motion.** Data is in motion when it is being transferred from one place to another. Sometimes, it is moving from your local area network to the internet, but it can also be internal to your network, such as from a server to a client computer. You can encrypt data in motion to protect it. For example, a web server uses a certificate to encrypt data being viewed by a user, and you can use IPsec to encrypt communications. There are many options. The most important point is to use encryption whenever possible, including for internal-only web sites available only to workers connected to your local area network.
 - **Data in use.** Data in use is often in memory because it is being used by, say, a developer working on some code updates or a user running reports on company sales. The data must be available to the relevant applications and operating system functions. There are some third-party solutions for encrypting data in memory, but the selection is limited. In addition to keeping the latest patches deployed to all computing devices, maintaining a standard computer build process, and running anti-virus and anti-malware software, organizations often use strong authentication, monitoring and logging to protect data in use.

2.6 Establish information and asset handling requirements

This section covers how people and systems work with data. This includes any action you can take with the data, such as read, copy, edit or delete. The key subtopics are important to know:

- **Markings and labels.** You should mark data to ensure that users are following the proper handling requirements. The data could be printouts or media like disks or backup tapes. For example, if your employee review process is on paper, the documents should be labeled as sensitive, so that anyone who stumbles across them accidentally will know not to read them but turn them over to the data owner or a member of the management or security team. You also might restrict the movement of confidential data, such as backup tapes, to certain personnel or to certain areas of your facility. Without labels, the backup tapes might not be handled in accordance with company requirements.
- **Storage.** You can store data in many ways, including on paper, disk or tape. For each scenario, you must define the acceptable storage locations and inform users about those locations. It is common to provide a vault or safe for backup tapes stored on premises, for example. Personnel who deal with sensitive papers should have a locked

cabinet or similar secure storage for those documents. Users should have a place to securely store files, such as an encrypted volume or an encrypted shared folder.

- **Destruction.** Your organization should have a policy for destruction of sensitive data. The policy should cover all the mediums that your organization uses for storing data — paper, disk, tape, etc. Some data classifications, such as those that deal with sensitive or confidential information, should require the most secure form of data destruction, such as physical destruction or secure data deletion with multiple overwrite passes. Other classifications might require only a single overwrite pass. The most important thing is to document the requirement for the various forms of media and the classification levels. When in doubt, destroy data as though it were classified as the most sensitive data at your organization.

Domain 2 Review Questions

Read and answer the following questions. If you do not get at least one correct, then spend more time with the subject. Then move on to Domain 3.

1. You are performing a security audit for a customer. During the audit, you find several instances of users gaining access to data without going through a formal access approval process. As part of the remediation, you recommend establishing a formal access approval process. Which role should you list to approve policies that dictate which users can gain access to data?
 - a. Data creator
 - b. Data processor
 - c. Data custodian
 - d. Data owner
 - e. System owner
2. Your organization has a goal to maximize the protection of organizational data. You need to recommend 3 methods to minimize data remanence in the organization. Which 3 of the following methods should you recommend?
 - a. Formatting volumes
 - b. Overwriting of data
 - c. Data encryption
 - d. Degaussing
 - e. Physical destruction
3. You are preparing to build a hybrid cloud environment for your organization. Three vendors present their proposed solution. Which methodology should your team use to select the best solution?
 - a. Standards selection
 - b. Standards deviation
 - c. Vendor screening
 - d. Vendor reviewing

Answers to Domain 2 Review Questions

1. Answer: D

Explanation: Each data owner is responsible for approving access to data that they own. This is typically handled via approving data access policies that are then implemented by the operations team. As part of a formal access approval process, a data owner should be the ultimate person responsible for the data access.

2. Answer: B, D, E

Explanation: When you perform a typical operating system deletion, the data remains on the media but the space on the media is marked as available. Thus, the data is often recoverable. There are 3 established methods for preventing data recovery: overwriting the data (sometimes referred to as a “secure deletion” or “wiping”), degaussing with magnets and physical destruction.

Formatting a volume does not render data unrecoverable, and neither does data encryption (if somebody had the decryption key, the data is at risk).

3. Answer: A

Explanation: In this scenario, your goal is to evaluate the solutions presented, not the vendors, so you should use a standards selection process. This will enable the team to select the solution that best fits the organization's needs. While a vendor selection process is part of engaging with a vendor, this scenario specifically calls for the evaluation of the solutions.

Domain 3. Security Architecture and Engineering

This domain is more technical than some of the others. If you already work in a security engineering role, then you have an advantage in this domain. If you don't, allocate extra time to be sure you have a firm understanding of the topics. Note that some of the concepts in this domain are foundational in nature, so you'll find aspects of them throughout the other domains.

3.1 Implement and manage engineering processes using secure design principles

When managing projects or processes, you need to use proven principles to ensure you end up with a functional solution that meets or exceeds the requirements, stays within the budget, and does not introduce unnecessary risk to the organization. The following are the high-level phases of a project:

- **Idea or concept.** You might want to create an app or a new web site, or deploy a new on-premises virtualized infrastructure. At this stage, the priority is to stay at a high level, without details. You need to document what the idea or concept will amount to. For example, you want to develop an app that will enable customers to schedule appointments, manage their accounts and pay their bills.
- **Requirements.** It is important to document all the requirements from the various business units and stakeholders. Establish both functional requirements (for example, the app will enable users to pay bills by taking a picture of their credit card) and non-functional requirements (for example, the app must be PCI DSS compliant).
- **Design.** Next, establish a design to meet the requirements. A design cannot be completed without all requirements. For example, to know how robust an infrastructure to design, you need to know how many users need to use the system simultaneously. Part of the design phase must be focused around security. For example, you must account for the principle of least privilege, fail-safe defaults and segregation of duties.
- **Develop and implement in a non-production environment.** In this phase, you create and deploy hardware, software and code as applicable for your project into a non-production environment (typically a development environment).
- **Initial testing.** Teams test the non-production implementation. The goal is to find and eliminate major bugs, missing functionality and other issues. It is common to go back to the previous phase to make necessary changes. Occasionally, you might have to even go back to the design phase.
- **Implementation.** Once all requirements have been met and the team is satisfied, you can move to a quality assurance (QA) environment. There, you'll repeat the "develop and implement" phase and the testing phase. Then you will move the app or service to the production environment.
- **Support.** After you implement your solution, you must operationalize it. Support teams and escalation paths should have been identified as part of the design.

There are many other phases, such as user training, communication and compliance testing. Remember that skipping any of these steps reduces the chances of having a successful and secure solution.

3.2 Understand the fundamental concepts of security models

Security models enable people to access only the data classified for their clearance level. There are many models. We will cover Bell-LaPadula and Biba, both of which use mathematical formulas. You don't need to know the formulas or other details for the exam, but you should be familiar with the models and their pros and cons.

- **Bell-LaPadula.** This model was established in 1973 for the United States Air Force. It focuses on confidentiality. The goal is to ensure that information is exposed only to those with the right level of classification. For example, if you have a Secret clearance, you can read data classified as Secret, but not Top Secret data. This model has a "no read up" (users with a lower clearance cannot read data classified at a higher level) and a "no write down" (users with a clearance higher than the data cannot modify that data) methodology. Notice that Bell-LaPadula doesn't address "write up," which could enable a user with a lower clearance to write up to data classified at a higher level. To address this complexity, this model is often enhanced with other models that focus on integrity. Another downside to this model is that it doesn't account for covert channels. A covert channel is a way of secretly sending data across an existing connection. For example, you can send a single letter inside the IP identification header. Sending a large message is slow. But often such communication isn't monitored or caught.
- **Biba.** Released in 1977, this model was created to supplement Bell-LaPadula. Its focus is on integrity. The methodology is "no read down" (for example, users with a Top Secret clearance can't read data classified as Secret) and "no write up" (for example, a user with a Secret clearance can't write data to files classified as Top Secret). By combining it with Bell-LaPadula, you get both confidentiality and integrity.

There are other models; for example, the [Clark-Wilson model](#) also focuses on integrity.

3.3 Select controls based upon systems security requirements

For this section of the exam, you should be familiar with the Common Criteria for Information Technology Security Evaluation. The Common Criteria (CC) unifies older standards (CTCPEC, ITSEC and TCSEC) to provide a standard to evaluate systems against. CC evaluations are focused on security-related systems and products. The important concepts for this section are:

- To perform an evaluation, you need to select the target of evaluation (TOE). This might be a firewall or an anti-malware app.
- The evaluation process will look at the protection profile (PP), which is a document that outlines the security needs. A vendor might opt to use a specific protection profile for a particular solution.
- The evaluation process will look at the security target (ST), which identifies the security properties for the TOE. The ST is usually published to customers and partners and available to internal staff.
- The evaluation will attempt to gauge the confidence level of a security feature. Security assurance requirements (SARs) are documented and based on the development of the solution. Key actions during development and testing should be captured along the way. An evaluation assurance level (EAL) is a numerical rating used to assess the rigor of an evaluation. The scale is EAL 1 (cheap and easy) to EAL7 (expensive and complex).

3.4 Understand the security capabilities of information systems

This section focuses on the capabilities of specific computing components. Thus, it isn't a section where hands-on experience can give you an advantage. Some of these components are discussed in other sections, sometimes in more detail. Ensure that you are familiar with all the information in this section. For any topic in this section that is new to you, plan to dive deeper into the topic outside of this study guide.

- **Memory protection.** At any given time, a computing device might be running multiple applications and services. Each one occupies a segment of memory. The goal of memory protection is to prevent one application or service from impacting another application or service. There are two popular memory protection methods:
 - **Process isolation.** Virtually all modern operating systems provide process isolation, which prevents one process from impacting another process.
 - **Hardware segmentation.** Hardware isolation is stricter than process isolation; the operating system maps processes to dedicated memory locations.
- **Virtualization.** In virtualized environments, there are special considerations to maximize security. The goal is to prevent attacks on the hypervisors and ensure that a compromise of one VM does not result in a compromise of all VMs on the host. Many organizations choose to deploy their high-security VMs to dedicated high-security hosts. In some cases, organizations have teams (such as the team responsible for identity and access management) manage their own virtualization environment to minimize the chances of an internal attack.
- **Trusted Platform Module.** A Trusted Platform Module (TPM) is a cryptographic chip that is sometimes included with a client computer or server. A TPM expands the capabilities of the computer by offering hardware-based cryptographic operations. Many security products and encryption solutions require a TPM. For example, BitLocker Drive Encryption (a built-in volume encryption solution) requires a TPM to maximize the security of the encryption.
- **Interfaces.** In this context, an interface is the method by which two or more systems communicate. For example, when an LDAP client communicates with an LDAP directory server, it uses an interface. When a VPN client connects to a VPN server, it uses an interface. For this section, you need to be aware of the security capabilities of interfaces. There are a couple of common capabilities across most interfaces:
 - **Encryption.** When you encrypt communications, a client and server can communicate privately without exposing information over the network. For example, if you use encryption between two email servers, then the SMTP transactions are encrypted and unavailable to attackers (compared to a default SMTP transaction which takes place in plain text). In some cases, an interface (such as LDAP) provides a method (such as LDAPS) for encrypting communication. When an interface doesn't provide such a capability, then IPsec or another encrypted transport mechanism can be used.
 - **Signing.** You can also sign communication, whether or not you encrypt the data. Signing communications tells the receiver, without a doubt, who the sender (client) is. This provides non-repudiation. In a high-security environment, you should strive to encrypt and sign all communications, though this isn't always feasible.

- **Fault tolerance.** Fault tolerance is a capability used to keep a system available. In the event of an attack (such as a DoS attack), fault tolerance helps keep a system up and running. Complex attacks can target a system, knowing that the fallback method is an older system or communication method that is susceptible to attack.

3.5 Assess and mitigate the vulnerabilities of security architectures, designs and solution elements

This section represents the vulnerabilities present in a plethora of technologies in an environment. You should feel comfortable reviewing an IT environment, spotting the vulnerabilities and proposing solutions to mitigate them. To do this, you need to understand the types of vulnerabilities often present in an environment and be familiar with mitigation options.

- **Client-based systems.** Client computers are the most attacked entry point. An attacker tries to gain access to a client computer, often through a phishing attack. Once a client computer is compromised, the attacker can launch attacks from the client computer, where detection is more difficult compared to attacks originating from the internet. Productivity software (word processors, spreadsheet applications) and browsers are constant sources of vulnerabilities. Even fully patched client computers are at risk due to phishing and social engineering attacks. To mitigate client-based issues, you should run a full suite of security software on each client computer, including anti-virus, anti-malware, anti-spyware and a host-based firewall.
- **Server-based systems.** While attackers often target client computer initially, their goal is often gaining access to a server, from which they can gain access to large amounts of data and potentially every other device on the network. To mitigate the risk of server-based attacks (whether attacking a server or attacking from a server), you should patch servers regularly — within days of new patches being released, and even sooner for patches for remote code execution vulnerabilities. In addition, you should use a hardened operating system image for all server builds. Last, you should use a host-based firewall to watch for suspicious traffic going to or from servers.
- **Database systems.** Databases often store a company's most important and sensitive data, such as credit card transactions, employees' personally identifiable information, customer lists, and confidential supplier and pricing information. Attackers, even those with low-level access to a database, might try to use inference and aggregation to obtain confidential information. Attackers might also use valid database transactions to work through data using data mining and data analytics.
- **Cryptographic systems.** The goal of a well-implemented cryptographic system is to make a compromise too time-consuming (such as 5,000 years) or too expensive (such as millions of dollars). Each component has vulnerabilities:
 - **Software.** Software is used to encrypt and decrypt data. It can be a standalone application with a graphical interface, or software built into the operating system or other software. As with any software, there are sometimes bugs or other issues, so regular patching is important.
 - **Keys.** A key dictates how encryption is applied through an algorithm. A key should remain secret; otherwise, the security of the encrypted data is at risk. Key length is an important consideration. To defend against quick brute-force attacks, you need a long key. Today, a 256-bit key is typically the minimum recommended for symmetric encryption, and a 2048-bit key is typically the minimum recommended for asymmetric encryption. However, the length should be based on your requirements and the sensitivity of the data being handled.

- **Algorithms.** There are many algorithms (or ciphers) to choose from. It is a good practice to use an algorithm with a large key space (a key space represents all possible permutations of a key) and a large random key value (a key value is a random value used by an algorithm for the encryption process). Algorithms are not secret, but instead well known.
- **Protocols.** There are different protocols for performing cryptographic functions. Transport Layer Security (TLS) is a very popular protocol used across the internet, such as for banking sites or sites that require encryption. Today, most sites (even Google) use encryption. Other protocols include Kerberos and IPsec.
- **Industrial Control Systems (ICS).** Supervisory control and data acquisition (SCADA) systems are used to control physical devices such as those found in an electrical power plant or factory. SCADA systems are well suited for distributed environments, such as those spread out across continents. Some SCADA systems still rely on legacy or proprietary communications. These communications are at risk, especially as attackers are gaining knowledge of such systems and their vulnerabilities.
- **Cloud-based systems.** Unlike systems on-premises, cloud-based systems are mainly controlled by cloud vendors. You often will not have access to or control of the hardware, software or supporting systems. When working with cloud-based systems, you need to focus your efforts on areas that you can control, such as the network entry and exit points (use firewalls and similar security solutions), encryption (use for all network communication and data at rest), and access control (use a centralized identity access and management system with multi-factor authentication). You should also gather diagnostic and security data from the cloud-based systems and store that information in your security information and event management system. With some cloud vendors, you might be able to configure aspects of the service, such as networking or access. In such scenarios, ensure that your cloud configuration matches or exceeds your on-premises security requirements. In high-security environments, your organization should have a dedicated cloud approach. Last, don't forget to look at the cloud vendors and understand their security strategy and tactics. You should be comfortable with the vendor's approach before you use their cloud services.
- **Distributed systems.** Distributed systems are systems that work together to perform a common task, such as storing and sharing data, computing, or providing a web service. Often, there isn't centralized management (especially with peer-to-peer implementations). In distributed systems, integrity is sometimes a concern because data and software are spread across various systems, often in different locations. To add to the trouble, there is often replication that is duplicating data across many systems.
- **Internet of Things (IoT).** Like cloud-based systems, you will have limited control over IoT devices. Mostly, you will have control of the configuration and updating. And you should spend extra time understanding both. Keeping IoT devices up to date on software patches is critically important. Without the latest updates, devices are often vulnerable to remote attacks from the internet. This is riskier than internal-only devices. On the configuration side, you should disable remote management and enable secure communication only (such as over HTTPS), at a minimum. As with cloud-based systems, review the IoT vendor to understand their history with reported vulnerabilities, response time to vulnerabilities and overall approach to security. Not all IoT devices are suitable for enterprise networks!

3.6 Assess and mitigate vulnerabilities in web-based systems

Web-based systems are systems you reach through the internet, often (but not always) through a web browser. Web-based systems are often meant to be public-facing, so they are exposed to the entire internet. This makes them vulnerable to attackers looking for easy targets, such as older and unpatched versions of web server software. There are several areas to review when you assess and mitigate vulnerabilities in web-based systems:

- **Web server software.** The web server software must be running the latest security patches. Running the latest version of the software can provide enhanced (and optional) security features. You need to have logging, auditing and monitoring for your web servers. The goal of these isn't to prevent attacks but instead to recognize warning signs early, before an attack or as early in the attack as possible. After an attack, the logs can provide critical information about the vulnerability, the date of compromise and sometimes even the identity of the attacker.
- **Endpoint security.** You also need to manage the client side. Clients that visit a compromised web server could become compromised. To minimize the risk of compromise, you need a multi-layered approach that includes a standardized browser configured for high security, web proxy servers to blacklist known bad web servers and track web traffic, host-based firewalls to block suspicious traffic, and anti-malware/anti-spyware/anti-virus software to watch for suspicious activity.
- **OWASP Top 10.** The Open Web Application Security Project (OWASP) publishes a list of the top 10 critical web application security risks. You should read through it and be familiar with these risks. See https://www.owasp.org/images/7/72/OWASP_Top_10-2017_%28en%29.pdf.pdf for more information. Here are two of the most important:
 - **Injection flaws (OWASP Top 10, #1).** Injection flaws have been around a long time. Two of the most common are SQL injection attacks and cross-site scripting (XSS) attacks. In an injection attack, an attacker provides invalid input to a web application, which is then processed by an interpreter. For example, an attacker might use special characters in a web-based form to alter how the form is processed (for example, comment out the password check). Input validation can help minimize the chances of an injection attack. But you need more than that. You need to properly test these types of scenarios prior to going live. One common mitigation strategy for SQL injection attacks is using prepared statements and parameterized queries; this enables the database to differentiate between code and data.
 - **XML External Entities / XXE (OWASP Top 10, #4).** In this type of attack, the goal is to pass invalid input (containing a reference to an external entity) to an XML parsing application. To minimize the potential for this attack, you can disable document type definitions (DTDs).

3.7 Assess and mitigate vulnerabilities in mobile systems

Today, mobile systems such as smartphones and tablets are full-blown computers. You can use them to connect to corporate networks and to produce, consume and share content. Therefore, these devices need to be treated like computers. You need to deploy and maintain security software, such as anti-malware and anti-virus software. You need to use encryption for storing data on the devices and for sending and receiving data, especially with the corporate network. You need to apply your organization's standards and security policies, when applicable. For example, you need to ensure that the devices are running the latest version of the software and have the latest patches. To deploy and maintain the

devices with a secure configuration, you need centralized management software so you can report on vulnerabilities and risk, and manage devices in bulk or with automation. At the device level, you need to require screen locks, strong authentication and encryption. You need to be able to remotely lock and wipe devices in the event a device is lost or stolen. Even with these things in place, you should restrict mobile systems to non-sensitive data, so they can't read or store PII or other confidential information.

3.8 Assess and mitigate vulnerabilities in embedded devices

In addition to managing security for your computing infrastructure and computers, you also should think about other systems that interact with your computing infrastructure. Today, that includes everything from coffee makers to smart white boards to copiers. These devices are becoming more and more connected, and some of them are even IoT devices. While these devices have had computers embedded in them for some time, they used to be standalone devices, not connected to your network, so a compromise was extremely limited and quite rare. Today, you need to consider the following information when managing your embedded devices:

- Some devices are configured by default to contact the manufacturer to report health information or diagnostic data. You need to be aware of such communication. Disable it when possible. At a minimum, ensure that the configuration is such that additional information cannot be sent out alongside the expected information.
- Some devices, by default, accept remote connections from anywhere. Sometimes the connections are for remote management. You should eliminate remote connectivity options for devices that do not need to be managed remotely.
- Many embedded systems and IoT systems are built for convenience, functionality and compatibility — security is often last on the list, so authentication and authorization are sometimes non-existent. Additionally, many systems are small and have limited battery life, so encryption is often not used because it drains the batteries too fast and requires ample CPU power. And your existing systems for managing device security and managing patches are not likely to be compatible with IoT devices, which makes managing software versions and patches difficult. Attackers have already exploited flaws in IoT devices; for example, one company was infected with malware that originated from a coffeemaker. As the number and sophistication of the devices increases, hackers will likely explore this attack vector even more.

3.9 Apply cryptography

Cryptography is present in several technologies. Applying cryptography is a big topic that covers several independent technologies. For the exam, be familiar with the high-level concepts around applying cryptography and its related technologies more so than understanding the details of implementing or supporting them. The subtopics below are the key topics outlined for this section. If you are new to cryptography or have limited exposure to it, consider additional sources to dive deeper.

- **Cryptographic lifecycle (e.g., cryptographic limitations, algorithm/protocol governance).** When we think about the lifecycle of technologies, we often think about the hardware and software support, performance and reliability. When it comes to cryptography, things are a bit different: The lifecycle is focused squarely around security. As computing power goes up, the strength of cryptographic algorithms goes down. It is only a matter of

time before there is enough computing power to brute-force through existing algorithms with common key sizes. You must think through the effective life of a certificate or certificate template, and of cryptographic systems. Beyond brute force, you have other issues to think through, such as the discovery of a bug or an issue with an algorithm or system. NIST defines the following terms that are commonly used to describe algorithms and key lengths: approved (a specific algorithm is specified as a NIST recommendation or FIPS recommendation), acceptable (algorithm + key length is safe today), deprecated (algorithm and key length is OK to use, but brings some risk), restricted (use of the algorithm and/or key length is deprecated and should be avoided), legacy (the algorithm and/or key length is outdated and should be avoided when possible), and disallowed (algorithm and/or key length is no longer allowed for the indicated use).

- **Cryptographic methods.** This subtopic covers the following three types of encryption. Be sure you know the differences.
 - **Symmetric.** Symmetric encryption uses the same key for encryption and decryption. Symmetric encryption is faster than asymmetric encryption because you can use smaller keys for the same level of protection. The downside is that users or systems must find a way to securely share the key and then hope that the key is used only for the specified communication.
 - **Asymmetric.** Asymmetric encryption uses different keys for encryption and decryption. Since one is a public key that is available to anybody, this method is sometimes referred to as “public key encryption.” Besides the public key, there is a private key that should remain private and protected. Asymmetric encryption doesn’t have any issues with distributing public keys. While asymmetric encryption is slower, it is best suited for sharing between two or more parties. RSA is one common asymmetric encryption standard.
 - **Elliptic curves.** Elliptic Curve Cryptography (ECC) is a newer implementation of asymmetric encryption. The primary benefit is that you can use smaller keys, which enhances performance.
- **Public key infrastructure (PKI).** A PKI is a foundational technology for applying cryptography. A PKI issues certificates to computing devices and users, enabling them to apply cryptography (for example, send encrypted email messages, encrypt web sites, or use IPsec to encrypt data communications). There are multiple vendors providing PKI services. You can run a PKI privately and solely for your own organization, you can acquire certificates from a trusted third-party provider, or you can do both, which is very common. A PKI is made up of certification authorities (CAs) (servers that provide one or more PKI functions, such as providing policies or issuing certificates), certificates (issued to other certification authorities or to devices and users), policies and procedures (such as how the PKI is secured), and templates (a predefined configuration for specific uses, such as a web server template). There are other components and concepts you should know for the exam:
 - A PKI can have multiple tiers. Having a single tier means you have one or more servers that perform all the functions of a PKI. When you have two tiers, you often have an offline root CA (a server that issues certificates to the issuing CAs but remains offline most of the time) in one tier, and issuing CAs (the servers that issue certificates to computing devices and users) in the other tier. The servers in the second tier are often referred to as intermediate CAs or subordinate CAs. Adding a third tier means you can have CAs that are only responsible for issuing policies (and they represent the second tier in a three-tier hierarchy). In such a scenario, the policy CAs should also remain offline and brought online only as needed. In general,

the more tiers, the more security (but proper configuration is critical). The more tiers you have, the more complex and costly the PKI is to build and maintain.

- A PKI should have a certificate policy and a certificate practice statement (CSP). A certificate policy documents how your company handles items like requestor identities, the uses of certificates and storage of private keys. A CSP documents the security configuration of your PKI and is usually available to the public.
- Besides issuing certificates, a PKI has other duties. For example, your PKI needs to be able to provide certificate revocation information to clients. If an administrator revokes a certificate that has been issued, clients must be able to get that information from your PKI. Another example is the storage of private keys and information about issued certificates. You can store these in a database or a directory.
- **Key management practices.** Remember, key management can be difficult with symmetric encryption but is much simpler with asymmetric encryption. There are several tasks related to key management:
 - **Key creation and distribution.** Key creation is self-explanatory. Key distribution is the process of sending a key to a user or system. It must be secure and it must be stored in a secure way on the computing device; often, it is stored in a secured store, such as the Windows certificate store.
 - **Key protection and custody.** Keys must be protected. You can use a method called split custody which enables two or more people to share access to a key — for example, with two people, each person can hold half the password to the key.
 - **Key rotation.** If you use the same keys forever, you are at risk of having the keys lost or stolen or having your information decrypted. To mitigate these risks, you should retire old keys and implement new ones.
 - **Key destruction.** A key can be put in a state of suspension (temporary hold), revocation (revoked with no reinstatement possible), expiration (expired until renewed), or destruction (such as at the end of a lifecycle or after a compromise).
 - **Key escrow and key backup recovery.** What happens if you encrypt data on your laptop but then lose your private key (for example, through profile corruption)? Normally, you lose the data. But key escrow enables storage of a key for later recovery. This is useful if a private key is lost or a court case requires escrow pending the outcome of a trial. You also need to have a method to back up and recover keys. Many PKIs offer a backup or recovery method, and you should take advantage of that if requirements call for it.
- **Digital signatures.** Digital signatures are the primary method for providing non-repudiation. By digitally signing a document or email, you are providing proof that you are the sender. Digital signatures are often combined with data encryption to provide confidentiality.
- **Non-repudiation.** For this section, non-repudiation refers to methods to ensure that the origin of data is can be deduced with certainty. The most common method for asserting the source of data is to use digital signatures, which rely on certificates. If User1 sends a signed email to User2, User2 can be sure that the email came from User1. It isn't foolproof though. For example, if User1 shares his credentials to his computer with User3, then User3 can send an email to User2 purporting to be User1, and User2 wouldn't have a way to deduce that. It is common to combine non-repudiation with confidentiality (data encryption).

- **Integrity.** A hash function implements encryption with a specified algorithm but without a key. It is a one-way function. Unlike encryption, where you can decrypt what's been encrypted, hashing isn't meant to be decrypted in the same way. For example, if you hash the word "hello", you might end up with "4cd21dba5fb0a60e26e83f2ac1b9e29f1b161e4c1fa7425e73048362938b4814". When apps are available for download, the install files are often hashed. The hash is provided as part of the download. If the file changes, the hash changes. That way, you can figure out if you have the original install file or a bad or modified file. Hashes are also used for storing passwords, with email and for other purposes. Hashes are susceptible to brute force. If you try to hash every possible word and phrase, eventually you will get the hash value that matches whatever hash you are trying to break. Salting provides extra protection for hashing by adding an extra, usually random, value to the source. Then, the hashing process hashes the original value of the source plus the salt value. For example, if your original source value is "Hello" and your salt value is "12-25-17-07:02:32", then "hello12-25-17-07:02:32" gets hashed. Salting greatly increased the strength of hashing.
- **Methods of cryptanalytic attacks.** There are several methods to attack cryptography. Each has strengths and weaknesses. The primary methods are:
 - **Brute force.** In a brute-force attack, every possible combination is attempted. Eventually, with enough time, the attack will be successful. For example, imagine a game where you have to guess the number between 1 and 1,000 that I chose. A brute-force attack would try all numbers between 1 and 1,000 until it found my number. This is a very simplified version of a brute-force attack, but the key point is that a brute-force attack will eventually be successful, provided it is using the correct key space. For example, if an attempt is made to brute force a password, the key space must include all the characters in the password; if the key space includes only letters but the password includes a number, the attack will fail.
 - **Ciphertext only.** In a ciphertext-only attack, you obtain samples of ciphertext (but not any plaintext). If you have enough ciphertext samples, the idea is that you can decrypt the target ciphertext based on the ciphertext samples. Today, such attacks are very difficult.
 - **Known plaintext.** In a known plaintext attack, you have an existing plaintext file and the matching ciphertext. The goal is to derive the key. If you derive the key, you can use it to decrypt other ciphertext created by the same key.
- **Digital rights management.** When people think of digital rights management (DRM), they think of protections placed on movies and games. But for the CISSP exam, it is really about protection of data, such as spreadsheets and email messages. Organizations often refer to data protection as enterprise digital rights management (E-DRM) or information rights management (IRM). Several vendors offer solutions to protect data in individual files. The solutions all provide a common set of foundational features:
 - Restrict viewing of a document to a defined set of people
 - Restrict editing of a document to a defined set of people
 - Expire a document (rendering it unreadable after a specified date)
 - Restrict printing of a document to a defined set of people
 - Provide portable document protection such that the protection remains with the document no matter where it is stored, how it is stored, or which computing device or user opens it

You can use DRM, E-DRM or IRM to protect data for your organization. Many of the solutions also enable you to securely share data with external organizations. Sometimes, this sharing is enabled through federation. Other times, the use of a public cloud provider enables cross-organization sharing. DRM, E-DRM and IRM provide companies with a way to provide confidentiality to sensitive documents. Additionally, some of the solutions enable you to track when and where documents were viewed. Last, some solutions enable you to update the protection of a document (such as removing a previously authorized viewer) even after a document has been sent and shared with external parties.

3.10 Apply security principles to site and facility design

This section applies to applying secure principles to data centers, server rooms, network operations centers and offices across an organization's locations. While some areas must be more secure than others, you must apply secure principles throughout your site to maximize security and reduce risk. Crime Prevention through Environmental Design (CPTED) is a well known set of guidelines for the secure design of buildings and office spaces. CPTED stresses three principles:

- **Natural surveillance.** Natural surveillance enables people to observe what's going on around the building or campus while going about their day-to-day work. It also eliminates hidden areas, areas of darkness and obstacles such as solid fences. Instead, it stresses low or see-through fencing, extra lighting, and the proper place of doors, windows and walkways to maximize visibility and deter crime.
- **Territoriality.** Territoriality is the sectioning of areas based on the area's use. For example, you might have a private area in the basement of your building for long-term company storage. It should be clearly designated as private, with signs, different flooring and other visible artifacts. The company's parking garage should have signs indicating that it is private parking only. People should recognize changes in the design of the space and be aware that they might be moving into a private area.
- **Access control.** Access control is the implementation of impediments to ensure that only authorized people can gain access to a restricted area. For example, you can put a gate at the driveway to the parking lot. For an unmanned server room, you should have a secure door with electronic locks, a security camera and signs indicating that the room is off limits to unauthorized people.

The overall goal is to deter unauthorized people from gaining access to a location (or a secure portion of a location), prevent unauthorized people from hiding inside or outside of a location, and prevent unauthorized people from committing attacks against the facility or personnel. There are several smaller activities tied to site and facility design, such as upkeep and maintenance. If your property is run down, unkempt or appears to be in disrepair, it gives attackers the impression that they can do whatever they want on your property.

3.11 Implement site and facility security controls

Physical security is a topic that covers all the interior and exterior of company facilities. While the subtopics are focused on the interior, many of the same common techniques are applicable to the exterior too.

- **Wiring closets.** A wiring closet is typically a small room that holds IT hardware. It is common to find telephony and network devices in a wiring closet. Occasionally, you also have a small number of servers in a wiring closet. Access

to the wiring closet should be restricted to the people responsible for managing the IT hardware. You should use some type of access control for the door, such as an electronic badge system or electronic combination lock. From a layout perspective, wiring closets should be accessible only in private areas of the building interior; people must pass through a visitor center and a controlled doorway prior to be able to enter a wiring closet.

- **Server rooms and data centers.** A server room is a bigger version of a wiring closet but not nearly as big as a data center. A server room typically houses telephony equipment, network equipment, backup infrastructure and servers. A server room should have the same minimum requirements as a wiring closet. While the room is bigger, it should have only one entry door; if there is a second door, it should be an emergency exit door only. It is common to use door alarms for server rooms: If the door is propped open for more than 30 seconds, the alarm goes off. All attempts to enter the server room without authorization should be logged. After multiple failed attempts, an alert should be generated.

Data centers are protected like server rooms, but often with a bit more protection. For example, in some data centers, you might need to use your badge both to enter and to leave, whereas with a server room, it is common to be able to walk out by just opening the door. In a data center, it is common to have one security guard checking visitors in and another guard walking the interior or exterior. Some organizations set time limits for authorized people to remain inside the data center. Inside a data center, you should lock everything possible, such as storage cabinets and IT equipment racks.

- **Media storage facilities.** Media storage facilities often store backup tapes and other media, so they should be protected just like a server room. It is common to have video surveillance too.
- **Evidence storage.** An evidence storage room should be protected like a server room or media storage facility.
- **Restricted work area.** Restricted work areas are used for sensitive operations, such as network operations or security operations. The work area can also be non-IT related, such as a bank vault. Protection should be like a server room, although video surveillance is typically limited to entry and exit points.
- **Utilities and HVAC.** When it comes to utilities such as HVAC, you need to think through the physical controls. For example, a person should not be able to crawl through the vents or ducts to reach a restricted area. For the health of your IT equipment, you should use separate HVAC systems. All utilities should be redundant. While a building full of cubicles might not require a backup HVAC system, a data center does, to prevent IT equipment from overheating and failing. In a high-security environment, the data center should be on a different electrical system than other parts of the building. It is common to use a backup generator just for the data center, whereas the main cubicle and office areas have only emergency lighting.
- **Environmental issues.** Some buildings use water-based sprinklers for fire suppression. In a fire, shut down the electricity before turning on the water sprinklers (this can be automated). Water damage is possible; by having individual sprinklers turn on, you can minimize the water damage to only what is required to put out a fire. Other water issues include flood, a burst pipe or backed up drains. Besides water issues, there are other environmental issues that can create trouble, such as earthquakes, power outages, tornados and wind. These issues should be considered before deciding on a data center site or a backup site. It is a good practice to have your secondary data center far enough away from your primary data center so it is not at risk from any environmental issues affecting the primary data center. For example, you should avoid building your backup data center on the same earthquake fault line as your primary data center, even if they are hundreds of miles away from each other.

- **Fire prevention, detection and suppression.** The following key points highlight things to know for this section:
 - **Fire prevention.** To prevent fires, you need to deploy the proper equipment, test it and manage it. This includes fire detectors and fire extinguishers. You also need to ensure that workers are trained about what to do if they see a fire and how to properly store combustible material. From a physical perspective, you can use firewalls and fire suppressing doors to slow the advancement of a fire and compartmentalize it.
 - **Fire detection.** The goal is to detect a fire as soon as possible. For example, use smoke detectors, fire detectors and other sensors (such as heat sensors).
 - **Fire suppression.** You need a way to suppress a fire once a fire breaks out. Having emergency pull levers for employees to pull down if they see a fire can help expedite the suppression response (for example, by automatically calling the fire department when the lever is pulled). You can use water-based fire-suppression system, or minimize the chances of destroying IT equipment by choosing non-water fire suppressants, such as foams, powders CO2-based solutions, or an FM-200 system. FM-200 systems replace Halon, which was banned for depleting the ozone layer. FM-200 is more expensive than water sprinklers.

Domain 3 Review Questions

Read and answer the following questions. If you do not get at least one correct, then spend more time with the subject. Then move on to Domain 4.

1. You are a security consultant tasked with reviewing a company's security model. The current model has the following characteristics:
 - It establishes confidentiality such that people cannot read access classified at a higher level than their clearance.
 - It forbids users with a specific clearance from writing data to a document with a lower clearance level.

You note that the current model does not account for somebody with a low clearance level from writing data to a document classified at a higher level than their clearance. You need to implement a model to mitigate this. Which of the following security tenets should the new model focus on?

- a. Availability
 - b. Governance
 - c. Integrity
 - d. Due diligence
 - e. Due care
2. You are documenting the attempted attacks on your organization's IT systems. The top type of attack was injection attacks. Which definition should you use to describe an injection attack?
 - e. Overloading a system or network
 - f. Plugging in infected portable hard drives
 - g. Capturing packets on a network
 - h. Providing invalid input
 - i. Intercepting and altering network communications
 3. You are designing a public key infrastructure for your organization. The organization has issued the following requirements for the PKI:
 - Maximize security of the PKI architecture
 - Maximize the flexibility of the PKI architecture

You need to choose a PKI design to meet the requirements. Which design should you choose?

- a. A two-tier hierarchy with an offline root CA being in the first tier and issuing CAs in the second tier
- b. A two-tier hierarchy with an online root CA being in the first tier and issuing CAs in the second tier
- c. A three-tier hierarchy with an offline root CA being in the first tier, offline policy CAs being in the second tier, and issuing CAs being in the third tier
- d. A three-tier hierarchy with an offline root CA being in the first tier, online policy CAs being in the second tier, and issuing CAs being in the third tier

Answers to Domain 3 Review Questions

1. Answer: C

Explanation: In this scenario, the existing model focused on confidentiality. To round out the model and meet the goal of preventing “write up,” you need to supplement the existing model with a model that focuses on integrity (such as Biba). Focusing on integrity will ensure that you don’t have “write up” (or “read down” either, although that wasn’t a requirement in this scenario).

2. Answer: D

Explanation: An injection attack provides invalid input to an application or web page. The goal is to craft that input so that a backend interpreter either performs an action not intended by the organization (such as running administrative commands) or crashes. Injection attacks are mature and routinely used, so it is important to be aware of them and how to protect against them.

3. Answer: C

Explanation: When designing a PKI, keep in mind the basic security tenets — the more tiers, the more security, and the more flexibility. Of course, having more tiers also means more cost and complexity. In this scenario, to maximize security and flexibility, you need to use a three-tier hierarchy with the root CAs and the policy CAs being offline. Offline CAs enhance security. Multiple tiers, especially with the use of policy CAs, enhance flexibility because you can revoke one section of the hierarchy without impacting the other (for example, if one of the issuing CAs had a key compromised).

Domain 4. Communication and Network Security

Networking can be one of the most complex topics on the CISSP exam. If you are lucky enough to have a network background, then you won't find this domain difficult. However, if your background doesn't have much networking, spend extra time in this section and consider diving deep into topics that still don't make sense after you go through this section.

4.1 Implement secure design principles in network architecture

This section addresses the design aspects of networking, focusing on security. While networking's primary function is to enable communication, security will ensure that the communication is between authorized devices only and the communication is private when needed.

- **Open System Interconnection (OSI) and Transmission Control Protocol/Internet Protocol (TCP/IP) models.** The Open Systems Interconnection (OSI) model is the more common of the two prevailing network models. However, in the context of CISSP, you must also be aware of the TCP/IP model and how it compares to the OSI model. The TCP/IP model uses only four layers, while the OSI model uses seven. The following table summarizes the layers of each model.

Layer Number	OSI Model	TCP/IP Model
7	Application	Applications
6	Presentation	
5	Session	
4	Transport	TCP (host to host)
3	Network	IP
2	Data link	Network access
1	Physical	

Many people use mnemonics to memorize the OSI layers. One popular mnemonic for the OSI layers is "All People Seem To Need Data Processing."

- **Internet Protocol (IP) networking.** IP networking is what enables devices to communicate. IP provides the foundation for other protocols to be able to communicate. IP itself is a connectionless protocol. IPv4 is for 32-bit addresses, and IPv6 is for 128-bit addresses. Regardless of which version you use to connect devices, you then typically use TCP or UDP to communicate over IP. TCP is a connection-oriented protocol that provides reliable communication, while UDP is a connectionless protocol that provides best-effort communication. Both protocols use standardized port numbers to enable applications to communicate over the IP network.
- **Implications of multilayer protocols.** Some protocols simultaneously use multiple layers of the OSI or TCP/IP model to communicate, and traverse the layers at different times. The process of traversing these layers is called encapsulation. For example, when a Layer 2 frame is sent through an IP layer, the Layer 2 data is encapsulated into

a Layer 3 packet, which adds the IP-specific information. Additionally, that layer can have other TCP or UDP data added to it for Layer 4 communication.

- **Converged protocols.** Like encapsulation, converged protocols enable communication over different mediums. For example, FCoE sends typical fibre channel control commands over Ethernet. Voice over IP (VoIP) sends SIP or other voice protocols over typical IP networks. In most cases, this provides simplicity, since the same infrastructure can be used for multiple scenarios. However, it can also add complexity by introducing more protocols and devices to manage and maintain on that same infrastructure.
- **Software-defined networks.** As networks, cloud services and multi-tenancy grow, the need to manage these networks has changed. Many networks follow either a two-tier (spine/leaf or core/access) or a three-tier (core, distribution, edge/access) topology. While the core network might not change that frequently, the edge or access devices can communicate with a variety of devices types and tenants. Increasingly, the edge or access switch is a virtual switch running on a hypervisor or virtual machine manager. You must be able to add a new subnet or VLAN or make other network changes on demand. You must be able to make configuration changes programmatically across multiple physical devices, as well as across the virtual switching devices in the topology. A software-defined network enables you to make these changes for all devices types with ease.
- **Wireless networks.** Wireless networks can be broken into the different 802.11 standards. The most common protocols within 802.11 are shown in the table below. Additional protocols have been proposed to IEEE, including ad, ah, aj, ax, ay and az. You should be aware of the frequency that each protocol uses.

802.11 protocol	Frequency	Data stream rate
a	5 GHz	Up to 54 Mbps
b	2.4 GHz	Up to 11 Mbps
g	2.4 GHz	Up to 54 Mbps
n	2.4–5 GHz	Up to 600 Mbps
ac	5 GHz	Up to 3466 Mbps

You should also be familiar with the wireless security standards:

- **Wired Equivalent Privacy (WEP).** WEP is a legacy security algorithm for wireless networks. Originally, it was the only encryption protocol for 802.11a and 802.11b networks. WEP used 64-bit to 256-bit keys, but with a weak stream cipher. WEP was deprecated in 2004 in favor of WPA and WPA2. Today, WEP should be avoided.
- **Wi-Fi Protected Access (WPA).** WPA uses Temporal Key Integrity Protocol (TKIP) with a 128-bit per-packet key. However, WPA is still vulnerable to password cracking from packet spoofing on a network. WPA typically uses a pre-shared key (PSK) and Temporal Key Integrity Protocol (TKIP) for encryption. This is known as WPA Personal (which is typically used in a home environment). There is also a WPA Enterprise which can use certificate authentication or an authentication server (such as a RADIUS server).
- **Wi-Fi Protected Access II (WPA 2).** WPA2 is the current standard for wireless encryption. WPA2 is based on the Advanced Encryption Standard (AES) cipher with message authenticity and integrity checking. AES is stronger than TKIP. Like WPA, WPA2 offers a PSK mode (for home or small business) and an enterprise mode (known as WPA2-ENT). WPA2-ENT uses a new encryption key each time a user connects. The

password is not stored on the client devices (unlike PSK mode, which stores the passwords locally on clients).

Regardless of the security method you use, you should also use TLS or IPsec for network communication. Finally, remember that wireless networks use collision avoidance, instead of the collision detection used on wired networks.

4.2 Secure network components

The components of a network make up the backbone of the logical infrastructure for an organization. These components are often critical to day-to-day operations, and an outage or security issue can cause millions of dollars in business losses. Here are issues to pay attention to:

- **Operation of hardware.** Modems are a type of Channel Service Unit/Data Service Unit (CSU/DSU) typically used for converting analog signals into digital. In this scenario, the CSU handles communication to the provider network, while the DSU handles communication with the internal digital equipment (in most cases, a router). Modems typically operate on Layer 2 of the OSI model. Routers operate on Layer 3 of the OSI model, and make the connection from a modem available to multiple devices in a network topology, including switches, access points and endpoint devices. Switches are typically connected to a router to enable multiple devices to use the connection. Switches help provide internal connectivity, as well as create separate broadcast domains when configured with VLANs. Switches typically operate at Layer 2 of the OSI model, but many switches can operate at both Layer 2 and Layer 3. Access points can be configured in the network topology to provide wireless access using one of the protocols and encryption algorithms discussed in section 4.1.
- **Transmission media.** Wired transmission media can typically be described in three categories: coaxial, Ethernet and fiber. Coaxial is typically used with cable modem installations to provide connectivity to an ISP, and requires a modem to convert the analog signals to digital. While Ethernet can be used to describe many mediums, it is typically associated with Category 5 and Category 6 unshielded twisted-pair (UTP) or shielded twisted pair (STP), and can be plenum-rated for certain installations. Fiber typically comes in two options, single-mode or multi-mode. Single-mode is typically used for long-distance communication, over several kilometers or miles. Multi-mode fiber is typically used for faster transmission, but with a distance limit depending on the desired speed. Fiber is most often used in the datacenter for backend components.
- **Network access control (NAC) devices.** Much as you need to control physical access to equipment and wiring, you need to use logical controls to protect a network. There are a variety of devices that provide this type of protection, including the following:
 - **Stateful and stateless firewalls** can perform inspection of the network packets that traverse it and use rules, signatures and patterns to determine whether the packet should be delivered. Reasons for dropping a packet could include addresses that don't exist on the network, ports or addresses that are blocked, or the content of the packet (such as malicious packets that have been blocked by administrative policy).
 - **Intrusion detection and prevention devices.** These devices monitor the network for unusual network traffic and MAC or IP address spoofing, and then either alert on or actively stop this type of traffic.

- **Proxy or reverse proxy servers.** Proxy servers can be used to proxy internet-bound traffic to the internet, instead of having clients going directly to the internet. Reverse proxies are often deployed to a perimeter network. They proxy communication from the internet to an internal server, such as a web server. Like a firewall, a reverse proxy can have rules and policies to block certain types of communication.
- **Endpoint security.** The saying “a chain is only as strong as its weakest link” can also apply to your network. Endpoint security can be the most difficult to manage and maintain, but also the most important part of securing a network. It can include authentication on endpoint devices, multifactor authentication, volume encryption, VPN tunnels and network encryption, remote access, anti-virus and anti-malware software, and more. Unauthorized access to an endpoint device is one of the easiest backdoor methods into a network because the attack surface is so large. Attackers often target endpoint devices hoping to use the compromised device as a launching spot for lateral movement and privilege escalation. Beyond the traditional endpoint protection methods, there are others that provide additional security:
 - **Application whitelisting.** Only applications on the whitelist can run on the endpoint. This can minimize the chances of malicious applications being installed or run.
 - **Restricting the use of removable media.** In a high-security organization, you should minimize or eliminate the use of removable media, including any removable storage devices that rely on USB or other connection methods. This can minimize malicious files coming into the network from the outside, as well as data leaving the company on tiny storage mechanisms.
 - **Automated patch management.** Patch management is the most critical task for maintaining endpoints. You must patch the operating system as well as all third-party applications. Beyond patching, staying up to date on the latest versions can bring enhanced security.
- **Content-distribution networks (CDNs).** CDNs are used to distribute content globally. They are typically used for downloading large files from a repository. The repositories are synchronized globally, and then each incoming request for a file or service is directed to the nearest service location. For example, if a request comes from Asia, a local repository in Asia, rather than one in the United States, would provide the file access. This reduces the latency of the request and typically uses less bandwidth. CDNs are often more resistant to denial of service (DoS) attacks than typical corporate networks, and they are often more resilient.
- **Physical devices.** Physical security is one of the most important aspects of securing a network. Most network devices require physical access to perform a reset, which can cause configurations to be deleted and grant the person full access to the device and an easy path to any devices attached to it. The most common methods for physical access control are code-based or card-based access. Unique codes or cards are assigned to individuals to identify who accessed which physical doors or locks in the secure environment. Secure building access can also involve video cameras, security personnel, reception desks and more. In some high-security organizations, it isn't uncommon to physically lock computing devices to a desk. In the case of mobile devices, it is often best to have encryption and strong security policies to reduce the impact of stolen devices because physically protecting them is difficult.

4.3 Implement secure communication channels according to design

This section focuses on securing data in motion. You need to understand both design and implementation aspects.

- **Voice.** As more organizations switch to VoIP, voice protocols such as SIP have become common on Ethernet networks. This has introduced additional management, either by using dedicated voice VLANs on networks, or establishing quality of service (QoS) levels to ensure that voice traffic has priority over non-voice traffic. Other web-based voice applications make it more difficult to manage voice as a separate entity. The consumer Skype app, for example, allows for video and voice calls over the internet. This can cause additional bandwidth consumption that isn't typically planned for in the network topology design or purchased from an ISP.
- **Multimedia collaboration.** There are a variety of new technologies that allow instant collaboration with colleagues. Smartboards and interactive screens make meeting in the same room more productive. Add in video technology, and someone thousands of miles away can collaborate in the same meeting virtually. Instant messaging through Microsoft Teams, Slack and other applications enables real-time communication. Mobile communication has become a huge market, with mobile apps such as WhatsApp, WeChat and LINE making real-time communication possible anywhere in the world.
- **Remote access.** Because of the abundance of connectivity, being productive in most job roles can happen from anywhere. Even in a more traditional environment, someone working outside of the office can use a VPN to connect and access all the internal resources for an organization. Taking that a step further, Remote Desktop Services (RDS) and virtual desktop infrastructure (VDI) can give you the same experience whether you're in the office or at an airport: If you have an internet connection, you can access the files and applications that you need to be productive. A screen scraper is a security application that captures a screen (such as a server console or session) and either records the entire session or takes a screen capture every couple of seconds. Screen scraping can help establish exactly what a person did when they logged into a computer. Screen scrapers are most often used on servers or remote connectivity solutions (such as VDI or Remote Desktop farms).
- **Data communications.** Whether you are physically in an office or working remotely, the communication between the devices being used should be encrypted. This prevents any unauthorized device or person from openly reading the contents of packets as they are sent across a network. Corporate networks can be segmented into multiple VLANs to separate different resources. For example, the out-of-band management for certain devices can be on a separate VLAN so that no other devices can communicate unless necessary. Production and development traffic can be segmented on different VLANs. An office building with multiple departments or building floors can have separate VLANs for each department or each floor in the building. Logical network designs can tie into physical aspects of the building as necessary. Even with VLAN segments, the communication should be encrypted using TLS, SSL or IPsec.
- **Virtualized networks.** Many organizations use hypervisors to virtualize servers and desktops for increased density and reliability. However, to host multiple servers on a single hypervisor, the Ethernet and storage networks must also be virtualized. VMware vSphere and Microsoft Hyper-V both use virtual network and storage switches to allow communication between virtual machines and the physical network. The guest operating systems running in the VMs use a synthetic network or storage adapter, which is relayed to the physical adapter on the host. The software-defined networking on the hypervisor can control the VLANs, port isolation, bandwidth and other aspects just as if it was a physical port.

Domain 4 Review Questions

Read and answer the following questions. If you do not get at least one correct, then spend more time with the subject. Then move on to Domain 5.

1. You are troubleshooting some anomalies with network communication on your network. You notice that some communication isn't taking the expected or most efficient route to the destination. Which layer of the OSI model you should troubleshoot?
 - a. Layer 2
 - b. Layer 3
 - c. Layer 4
 - d. Layer 5
 - e. Layer 7

2. A wireless network has a single access point and two clients. One client is on the south side of the building toward the edge of the network. The other client is on the north side of the building, also toward the edge of the network. The clients are too far from each other to see each other. In this scenario, which technology can be used to avoid collisions?
 - a. Collision detection
 - b. Collision avoidance
 - c. Channel service unit
 - d. Data service unit

3. Your company uses VoIP for internal telephone calls. You are deploying a new intrusion detection system and need to capture traffic related to internal telephone calls only. Which protocol should you capture?
 - a. H.264
 - b. DNS
 - c. H.263
 - d. HTTPS
 - e. SIP

Answers to Domain 4 Review Questions

1. Answer: B

Explanation: In this scenario, the information indicates that the issue is with the routing of the network communication. Routing occurs at Layer 3 of the OSI model. Layer 3 is typically handled by a router or the routing component of a network device.

2. Answer: B

Explanation: In this scenario, collision avoidance is used. Wireless networks use collision avoidance specifically to address the issue described in the scenario (which is known as the "hidden node problem").

3. Answer: E

Explanation: SIP is a communications protocol used for multimedia communication such as internal voice calls. In this scenario, you need to capture SIP traffic to ensure that you are only capturing traffic related to the phone calls.

Domain 5. Identity and Access Management (IAM)

This section covers technologies and concepts related to authentication and authorization, for example, usernames, passwords and directories. While it isn't a huge domain, it is technical and there are many important details related to the design and implementation of the technologies.

5.1 Control physical and logical access to assets

There are some common methods for controlling access without regard for the asset type. For example, we need a way to authenticate users — validate that they are who they say they are. Then we need a way to authorize the users — figure out whether they are authorized to perform the requested action for the specific asset (such as read or write a given file or enter a particular server room). Let's take a closer look at how authentication and authorization typically work.

- **Authentication.** Traditional authentication systems rely on a username and password, especially for authenticating to computing devices. LDAP directories are commonly used to store user information, authenticate users and authorize users. But there are newer systems that enhance the authentication experience. Some replace the traditional username and password systems, while others (such as single sign-on, or SSO), extend them. Biometrics is an emerging authentication method that includes (but is not limited to) fingerprints, retina scans, facial recognition and iris scans.
- **Authorization.** Traditional authorization systems rely on security groups in a directory, such as an LDAP directory. Based on your group memberships, you have a specific type of access (or no access). For example, administrators might grant one security group read access to an asset, while a different security group might get read/write/execute access to the asset. This type of system has been around a long time and is still the primary authorization mechanism for on-premises technologies. Newer authorization systems incorporate dynamic authorization or automated authorization. For example, the authorization process might check to see if you are in the Sales department and in a management position before you can gain access to certain sales data. Other information can be incorporated into authorization. For example, you can authenticate and get read access to a web-based portal, but you can't get into the admin area of the portal unless you are connected to the corporate network.

Next, let's look at some key details around controlling access to specific assets.

- **Information.** "Information" and "data" are interchangeable here. Information is often stored in shared folders or in storage available via a web portal. In all cases, somebody must configure who can gain access and which actions they can perform. The type of authentication isn't relevant here. Authorization is what you use to control the access.
- **Systems.** In this context, "systems" can refer to servers or applications, either on premises or in the cloud. You need to be familiar with the various options for controlling access. In a hybrid scenario, you can use federated authentication and authorization in which the cloud vendor trusts your on-premises authentication and

authorization solutions. This centralized access control is quite common because it gives organizations complete control no matter where the systems are.

- **Devices.** Devices include computers, smartphones and tablets. Today, usernames and passwords (typically from an LDAP directory) are used to control access to most devices. Fingerprints and other biometric systems are common, too. In high-security environments, users might have to enter a username and password and then use a second authentication factor (such as a code from a smartcard) to gain access to a device. Beyond gaining access to devices, you also need to account for the level of access. In high-security environments, users should not have administrative access to devices, and only specified users should be able to gain access to particular devices.
- **Facilities.** Controlling access to facilities (buildings, parking garages, server rooms, etc.) is typically handled via badge access systems. Employees carry a badge identifying them and containing a chip. Based on their department and job role, they will be granted access to certain facilities (such as the main doors going into a building) but denied access to other facilities (such as the power plant or the server room). For high-security facilities, such as a data center, it is common to have multi-factor authentication. For example, you must present a valid identification card to a security guard and also go through a hand or facial scan to gain access to the data center. Once inside, you still need to use a key or smartcard to open racks or cages.

5.2 Manage identification and authentication of people, devices and services

This section builds on the previous section. The subtopics are more operational in nature and go into more detail.

- **Identity management implementation.** We looked briefly at SSO and LDAP. Now, we will look at them in more detail.
 - **SSO.** Single sign-on provides an enhanced user authentication experience as the user accesses multiple systems and data across a variety of systems. It is closely related to federated identity management (which is discussed later in this section). Instead of authenticating to each system individually, the recent sign-on is used to create a security token that can be reused across apps and systems. Thus, a user authenticates once and then can gain access to a variety of systems and data without having to authenticate again. Typically, the SSO experience will last for a specified period, such as 4 hours or 8 hours. SSO often takes advantage of the user's authentication to their computing device. For example, a user signs into their device in the morning, and later when they launch a web browser to go to a time-tracking portal, the portal accepts their existing authentication. SSO can be more sophisticated. For example, a user might be able to use SSO to seamlessly gain access to a web-based portal, but if the user attempts to make a configuration change, the portal might prompt for authentication before allowing the change. Note that using the same username and password to access independent systems is not SSO. Instead, it is often referred to as "same sign-on" because you use the same credentials. The main benefit of SSO is also its main downside: It simplifies the process of gaining access to multiple systems for everyone. For example, if attackers compromise a user's credentials, they can sign into the computer and then seamlessly gain access to all apps using SSO. Multi-factor authentication can help mitigate this risk.
 - **LDAP.** Lightweight Directory Access Protocol (LDAP) is a standards-based protocol (RFC 4511) that traces its roots back to the X.500 standard that came out in the early 1990s. Many vendors have implemented LDAP-compliant systems and LDAP-compliant directories, often with vendor-specific enhancements. LDAP

is especially popular for on-premises corporate networks. An LDAP directory stores information about users, groups, computers, and sometimes other objects such as printers and shared folders. It is common to use an LDAP directory to store user metadata, such as their name, address, phone numbers, departments, employee number, etc. Metadata in an LDAP directory can be used for dynamic authentication systems or other automation. The most common LDAP system today is Microsoft Active Directory (Active Directory Domain Services or AD DS). It uses Kerberos (an authentication protocol that offers enhanced security) for authentication, by default.

- **Single- or multi-factor authentication.** There are three different authentication factors — something you know, something you have and something you are. Each factor has many different methods. Something you know could be a username and password or the answer to a personal question; something you have could be a smartcard or a phone, and something you are could be a fingerprint or retinal scan. Single-factor authentication requires only one method from any of the three factors — usually a username and password. Multi-factor authentication (MFA) requires a method from each of two or three different factors, which generally increases security. For example, requiring you to provide a code sent to a hard token in addition to a username and password increases security because an attacker who steals your credentials is unlikely to also have access to the hard token. Different methods provide different levels of security, though. For example, the answer to a personal question isn't as secure as a token from a security app on your phone, because a malicious user is much more likely to be able to discover the information to answer the question on the internet than to get access to your phone. One downside to multi-factor authentication is the complexity it introduces; for instance, if a user doesn't have their mobile phone or token device with them, they can't sign in. To minimize issues, you should provide options for the second method (for example, the user can opt for a phone call to their landline).
- **Accountability.** In this context, accountability is the ability to track users' actions as they access systems and data. You need to be able to identify the users on a system, know when they access it, and record what they do while on the system. This audit data must be captured and logged for later analysis and troubleshooting. Important information can be found in this data. For example, if a user successfully authenticates to a computer in New York and then successfully authenticates to a computer in London a few minutes later, that is suspicious and should be investigated. If an account has repeated bad password attempts, you need data to track down the source of the attempts. Today, many companies are centralizing accountability. For example, all servers and apps send their audit data to the centralized system, so admins can gain insight across multiple systems with a single query. Because of the enormous amount of data in these centralized systems, they are usually "big data" systems, and you can use analytics and machine learning to unearth insights into your environment.
- **Session management.** After users authenticate, you need to manage their sessions. If a user walks away from the computer, anybody can walk up and assume their identity. To reduce the chances of that happening, you can require users to lock their computers when stepping away. You can also use session timeouts to automatically lock computers. You can also use password-protected screen savers that require the user to re-authenticate. You also need to implement session management for remote sessions. For example, if users connect from their computers to a remote server over Secure Shell (SSH) or Remote Desktop Protocol (RDP), you can limit the idle time of those sessions.
- **Registration and proofing of identity.** With some identity management systems, users must register and provide proof of their identity. For example, with self-service password reset apps, it is common for users to register and prove their identity. If they later forget their password and need to reset it, they must authenticate using an

alternative method, such as providing the same answers to questions as they provided during registration. Note that questions are often insecure and should be used only when questions can be customized or when an environment doesn't require a high level of security. One technique users can use to enhance question and answer systems is to use false answers. For example, if the question wants to know your mother's maiden name, you enter another name which is incorrect but serves as your answer for authentication. Alternatively, you can treat the answers as complex passwords. Instead of directly answering the questions, you can use a long string of alphanumeric characters such as "Vdsfh2873423#@\$wer78wreuy23143ya".

- **Federated Identity Management (FIM).** Note that this topic does not refer to Microsoft Forefront Identity Manager, which has the same acronym. Traditionally, you authenticate to your company's network and gain access to certain resources. When you use identity federation, two independent organizations share authentication and/or authorization information with each other. In such a relationship, one company provides the resources (such as a web portal) and the other company provides the identity and user information. The company providing the resources trusts the authentication coming from the identity provider. Federated identity systems provide an enhanced user experience because users don't need to maintain multiple user accounts across multiple apps. Federated identity systems use Security Assertion Markup Language (SAML), OAuth, or other methods for exchanging authentication and authorization information. SAML is the most common method for authentication in use today. It is mostly limited to use with web browsers, while OAuth isn't limited to web browsers. Federated identity management and SSO are closely related. You can't reasonably provide SSO without a federated identity management system. Conversely, you use federated identities without SSO, but the user experience will be degraded because everyone must re-authenticate manually as they access various systems.
- **Credentials management systems.** A credentials management system centralizes the management of credentials. Such systems typically extend the functionality of the default features available in a typical directory service. For example, a credentials management system might automatically manage the passwords for account passwords, even if those accounts are in a third-party public cloud or in a directory service on premises. Credentials management systems often enable users to temporarily check out accounts to use for administrative purposes. For example, a database administrator might use a credentials management system to check out a database admin account in order to perform some administrative work using that account. When they are finished, they check the account back in and the system immediately resets the password. All activity is logged and access to the credentials is limited. Without a credentials management system, you run the risk of having multiple credentials management approaches in your organization. For example, one team might use an Excel spreadsheet to list accounts and passwords, while another team might use a third-party password safe application. Having multiple methods and unmanaged applications increases risks for your organization. Implementing a single credentials management system typically increases efficiency and security.

5.3 Integrate identity as a third-party service

There are many third-party vendors that offer identity services that complement your existing identity store. For example, Ping Identity provides an identity platform that you can integrate with your on-premises directory (such as Active Directory) and your public cloud services (such as Microsoft Azure or Amazon AWS). Third-party identity services can help manage identities both on premises and in the cloud:

- **On premises.** To work with your existing solutions and help manage identities on premises, identity services often put servers, appliances or services on your internal network. This ensures a seamless integration and provides additional features, such as single sign-on. For example, you might integrate your Active Directory domain with a third-party identity provider and thereby enable certain users to authenticate through the third-party identity provider for SSO.
- **Cloud.** Organizations that want to take advantage of software-as-a-service (SaaS) and other cloud-based applications need to also manage identities in the cloud. Some of them choose identity federation — they federate their on-premises authentication system directly with the cloud providers. But there is another option: using a cloud-based identity service, such as Microsoft Azure Active Directory or Amazon AWS Identity and Access Management. There are some pros with using a cloud-based identity service:
 - You can have identity management without managing the associated infrastructure.
 - You can quickly start using a cloud-based identity service, typically within just a few minutes.
 - Cloud-based identity services are relatively inexpensive.
 - Cloud-based identity services offer services worldwide, often in more places and at a bigger scale than most organizations can.
 - The cloud provider often offers features not commonly found in on-premises environments. For example, a cloud provider can automatically detect suspicious sign-ins attempts, such as those from a different type of operating system than normal or from a different location than usual, because they have a large amount of data and can use artificial intelligence to spot suspicious logins.
 - For services in the cloud, authentication is local, which often results in better performance than sending all authentication requests back to an on-premises identity service.

You also need to be aware of the potential downsides:

- You lose control of the identity infrastructure. Because identity is a critical foundational service, some high-security organizations have policies that require complete control over the entire identity service. There is a risk in using an identity service in a public cloud, although the public cloud can sometimes be as secure or more secure than many corporate environments.
- You might not be able to use only the cloud-based identity service. Many companies have legacy apps and services that require an on-premises identity. Having to manage an on-premises identity infrastructure and a cloud-based identity system requires more time and effort than just managing an on-premises environment.
- If you want to use all the features of a cloud identity service, the costs rise. On-premises identity infrastructures are not expensive compared to many other foundational services such as storage or networking.
- There might be a large effort required to use a cloud-based identity service. For example, you need to figure out new operational processes. You need to capture the auditing and log data and often bring it back to your on-premises environment for analysis. You might have to update, upgrade or deploy new

software and services. For example, if you have an existing multi-factor authentication solution, it might not work seamlessly with your cloud-based identity service.

- **Federated.** Federation enables your organization to use their existing identities (such as those used to access your internal corporate systems) to access systems and resources outside of the company network. For example, if you use a cloud-based HR application on the internet, you can configure federation to enable employees to sign into the application with their corporate credentials. You can federate with vendors or partners. Federating between two organizations involves an agreement and software to enable your identities to become portable (and thus usable based on who you federate with). Federation typically provides the best user experience because users don't have to remember additional passwords or manage additional identities.

Other key facts about third-party identity services include:

- Often, you still need an on-premises directory service.
- Many third-party identity services started off as solutions for web-based applications. They have since to cover other use cases but still can't be used for many day-to-day authentication scenarios. For example, most of them can't authenticate users to their corporate laptops.
- Third-party identity services often offer single sign-on, multi-factor authentication and meta-directory services (pulling data from multiple directories into a single third-party directory).
- Many of the offerings are cloud-based, with a minimal on-premises footprint.
- Third-party identity services typically support SAML, OpenID Connect, WS-Federation, OAuth and WS-Trust.

5.4 Implement and manage authorization mechanisms

This section focuses on access control methods. To prepare for the exam, you should understand the core methods and the differences between them.

- **Role-based access control (RBAC).** RBAC is a common access control method. For example, one role might be a desktop technician. The role has rights to workstations, the anti-virus software and a software installation shared folder. For instance, if a new desktop technician starts at your company, you simply add them to the role group and they immediately have the same access as other desktop technicians. RBAC is a non-discretionary access control method because there is no discretion — each role has what it has. RBAC is considered an industry-standard good practice and is in widespread use throughout organizations.
- **Rule-based access control.** Rule-based access control implements access control based on predefined rules. For example, you might have a rule that permits read access to marketing data for anyone who is in the marketing department, or a rule that permits only managers to print to a high-security printer. Rule-based access control systems are often deployed to automate access management. Many rule-based systems can be used to implement access dynamically. For example, you might have a rule that allows anybody in the New York office to access a file server in New York. If a user tries to access the file server from another city, they will be denied access, but if they travel to the New York office, access will be allowed. Rule-based access control methods simplify access control in some scenarios. For example, imagine a set of rules based on department, title and location. If somebody transfers

to a new role or a new office location, their access is updated automatically. In particular, their old access goes away automatically, addressing a major issue that plagues many organizations.

- **Mandatory access control (MAC).** MAC is a method to restrict access based on a person's clearance and the data's classification or label. For example, a person with a Top Secret clearance can read a document classified as Top Secret. The MAC method ensures confidentiality. MAC is not in widespread use but is considered to provide higher security than DAC because individual users cannot change access.
- **Discretionary access control (DAC).** When you configure a shared folder on a Windows or Linux server, you use DAC. You assign somebody specific rights to a volume, a folder or a file. Rights could include read-only, write, execute, list and more. You have granular control over the rights, including whether the rights are inherited by child objects (such as a folder inside another folder). DAC is flexible and easy. It is in widespread use. However, anybody with rights to change permissions can alter the permissions. It is difficult to reconcile all the various permissions throughout an organization. It can also be hard to determine all the assets that somebody has access to, because DAC is very decentralized.
- **Attribute-based access control (ABAC).** Many organizations use attributes to store data about users, such as their department, cost center, manager, location, employee number and date of hire. These attributes can be used to automate authorization and to make it more secure. For example, you might configure authorization to allow only users who have "Paris" as their office location to use the wireless network at your Paris office. Or you might strengthen security for your HR folder by checking not only that users are members of a specific group, but also that their department attribute is set to "HR".

5.5 Manage the identity and access provisioning lifecycle

The identity lifecycle extends from the creation of users, to the provisioning of access, to the management of users, to the deprovisioning of access or users. While there are several methods to manage this lifecycle, the following ordered steps provide an overview of the typical implementation process:

1. A new user is hired at a company.
2. The HR department creates a new employee record in the human capital management (HCM) system, which is the authoritative source for identity information such as legal name, address, title and manager.
3. The HCM syncs with the directory service. As part of the sync, any new users in HCM are provisioned in the directory service.
4. The IT department populates additional attributes for the user in the directory service. For example, the users' email address and role might be added.
5. The IT department performs maintenance tasks such as resetting the user's password and changing the user's roles when they move to a new department.

6. The employee leaves the company. The HR department flags the user as terminated in the HCM, and the HCM performs an immediate sync with the directory service. The directory service disables the user account to temporarily remove access.
7. The IT department, after a specific period (such as 7 days), permanently deletes the user account and all associated access.

Beyond these steps, there are additional processes involved in managing identity and access:

- **User access review.** You should perform periodic access reviews in which appropriate personnel attest that each user has the appropriate rights and permissions. Does the user have only the access they need to perform their job? Were all permissions granted through the company's access request process? Is the granting of access documented and available for review? You should also review the configuration of your identity service to ensure it adheres to known good practices. You should review the directory service for stale objects (for example, user accounts for employees who have left the company). The primary goal is to ensure that users have the access permissions they need and nothing more. If a terminated user still has a valid user account, then you are in violation of your primary goal.
- **System account access review.** System accounts are accounts that are not tied one-to-one to humans. They are often used to run automated processes, jobs, and tasks. System accounts sometimes have elevated access. In fact, it isn't uncommon to find system accounts with the highest level of access (root or administrative access). System accounts require review similar to user accounts. You need to find out if system accounts have the minimum level of permissions required for what they are used for. And you need to be able to show the details — who provided the access, the date it was granted, and what the permissions provide access to.
- **Provisioning and deprovisioning.** Account creation and account deletion — provisioning and deprovisioning — are key tasks in the account lifecycle. Create accounts too early and you have dormant accounts that can be targeted. Wait too long to disable and delete accounts and you also have dormant accounts that can be targeted. When feasible, it is a good practice to automate provisioning and deprovisioning. Automation helps reduce the time to create and delete accounts. It also reduces human error (although the automation code could have human error). Your company should establish guidelines for account provisioning and deprovisioning. For example, your company might have a policy that an account must be disabled while the employee is in the meeting being notified of their termination.

Domain 5 Review Questions

Read and answer the following questions. If you do not get at least one of them correct, then spend more time with the subject. Then move on to Domain 6.

1. You are implementing a multi-factor authentication solution. As part of the design, you are capturing the three authentication factors. What are they?
 - a. Something you make
 - b. Something you know
 - c. Something you have
 - d. Something you need
 - e. Something you are
 - f. Something you do
2. Your company is rapidly expanding its public cloud footprint, especially with Infrastructure as a Service (IaaS), and wants to update its authentication solution to enable users to be authenticated to services in the cloud that are yet to be specified. The company issues the following requirements:
 - Minimize the infrastructure required for the authentication.
 - Rapidly deploy the solution.
 - Minimize the overhead of managing the solution.

You need to choose the authentication solution for the company. Which solution should you choose?

- a. A federated identity solution
 - b. A cloud-based identity service
 - c. A multi-factor authentication solution
 - d. A third-party identity service
3. A user reports that they cannot gain access to a shared folder. You investigate and find the following information:
 - Neither the user nor any groups the user is a member of have been granted permissions to the folder.
 - Other users and groups have been granted permissions to the folder.
 - Another IT person on your team reports that they updated the permissions on the folder recently.

Based on the information in this scenario, which type of access control is in use?

- a. RBAC
- b. Rule-based access control
- c. MAC
- d. DAC

Answers to Domain 5 Review Questions

1. Answer: B, C, E

Explanation: The three factors are something you know (such as a password), something you have (such as a smartcard or authentication app), and something you are (such as a fingerprint or retina). Using methods from multiple factors for authentication enhances security and mitigates the risk of a stolen or cracked password.

2. Answer: B

Explanation: With the rapid expansion to the cloud and the type of services in the cloud unknown, a cloud-based identity service, especially one from your public cloud vendor, is the best choice. Such services are compatible with IaaS, SaaS and PaaS solutions. While a third-party identity service can handle SaaS, it will not be as capable in non-SaaS scenarios. A federated identity solution is also limited to certain authentication scenarios and requires more time to deploy and more work to manage.

3. Answer: D

Explanation: Because you found individual users being granted permissions, and an IT administrator had manually changes permissions on the folder, DAC is in use. RBAC uses roles, and rule-based access control relies on rules and user attributes, so you would not find individual users configured with permissions on the folder with either of these. MAC is based on clearance levels, so, again, users aren't individually granted permissions on a folder; instead, a group for each clearance is used.

Domain 6. Security Assessment and Testing

This section covers assessments and audits, along with all the technologies and techniques you will be expected to know to perform them.

6.1 Design and validate assessment, test and audit strategies

An organization's assessment, testing and audit strategies will depend on its size, industry, financial status and other factors. For example, a small non-profit, a small private company and a small public company will all have different requirements and goals. Like any procedure or policy, the audit strategy should be assessed and tested regularly to ensure that the organization is not doing a disservice to itself with the current strategy. There are three types of audit strategies:

- **Internal.** An internal audit strategy should be aligned to the organization's business and day-to-day operations. For example, a publicly traded company will have a more rigorous auditing strategy than a privately held company. However, the stakeholders in both companies have an interest in protecting intellectual property, customer data and employee information. Designing the audit strategy should include laying out applicable regulatory requirements and compliance goals.
- **External.** An external audit strategy should complement the internal strategy, providing regular checks to ensure that procedures are being followed and the organization is meeting its compliance goals.
- **Third-party.** Third-party auditing provides a neutral and objective approach to reviewing the existing design, methods for testing and overall strategy for auditing the environment. A third-party audit can also ensure that both internal and external auditors are following the processes and procedures that are defined as part of the overall strategy.

6.2 Conduct security control testing

Security control testing can include testing of the physical facility, logical systems and applications. Here are the common testing methods:

- **Vulnerability assessment.** The goal of a vulnerability assessment is to identify elements in an environment that are not adequately protected. This does not always have to be from a technical perspective; you can also assess the vulnerability of physical security or the external reliance on power, for instance. These assessments can include personnel testing, physical testing, system and network testing, and other facilities tests.
- **Penetration testing.** A penetration test is a purposeful attack on systems to attempt to bypass automated controls. The goal of a penetration test is to uncover weaknesses in security so they can be addressed to mitigate risk. Attack techniques can include spoofing, bypassing authentication, privilege escalation and more. Like vulnerability assessments, penetration testing does not have to be purely logical. For example, you can use social engineering to try to gain physical access to a building or system.

- **Log reviews.** IT systems can log anything that occurs on the system, including access attempts and authorizations. The most obvious log entries to review are any series of “deny” events, since someone is attempting to access something that they don’t have permissions for. It’s more difficult to review successful events, since there are generally thousands of them, and almost all of them follow existing policies. However, it can be important to show that someone or something did indeed access a resource that they weren’t supposed to, either by mistake or through privilege escalation. A procedure and software to facilitate frequent review of logs is essential.
- **Synthetic transactions.** While user monitoring captures actual user actions in real time, synthetic — scripted or otherwise artificial — transactions can be used to test system performance or security.
- **Code review and testing.** Security controls are not limited to IT systems. The application development lifecycle must also include code review and testing for security controls. These reviews and controls should be built into the process just as unit tests and function tests are; otherwise, the application is at risk of being insecure.
- **Misuse case testing.** Software and systems can both be tested for use for something other than its intended purpose. From a software perspective, this could be to reverse engineer the binaries or to access other processes through the software. From an IT perspective, this could be privilege escalation, sharing passwords and accessing resources that should be denied.
- **Test coverage analysis.** You should be aware of the following coverage testing types:
 - **Black box testing.** The tester has no prior knowledge of the environment being tested.
 - **White box testing.** The tester has full knowledge prior to testing.
 - **Dynamic testing.** The system that is being tested is monitored during the test.
 - **Static testing.** The system that is being tested is not monitored during the test.
 - **Manual testing.** Testing is performed manually by humans.
 - **Automated testing.** A script performs a set of actions.
 - **Structural testing.** This can include statement, decision, condition, loop and data flow coverage.
 - **Functional testing.** This includes normal and anti-normal tests of the reaction of a system or software. Anti-normal testing goes through unexpected inputs and methods to validate functionality, stability and robustness.
 - **Negative testing.** This test purposely uses the system or software with invalid or harmful data, and verifies that the system responds appropriately.
- **Interface testing.** This can include the server interfaces, as well as internal and external interfaces. The server interfaces include the hardware, software and networking infrastructure to support the server. For applications, external interfaces can be a web browser or operating system, and internal components can include plug-ins, error handling and more. You should be aware of the different testing types for each system.

6.3 Collect security process data

Organizations should collect data about policies and procedures and review it on a regular basis to ensure that the established goals are being met. Additionally, they should consider whether new risks have appeared since the creation of the process that must now be addressed.

- **Account management.** Every organization should have a defined procedure for maintaining accounts that have access to systems and facilities. This doesn't just mean documenting the creation of a user account, but can include when that account expires and the logon hours of the account. This should also be tied to facilities access. For example, was an employee given a code or key card to access the building? Are there hours that the access method is also prevented? There should also be separate processes for managing accounts of vendors and other people who might need temporary access.
- **Management review and approval.** Management plays a key role in ensuring that these processes are distributed to employees, and that they are followed. The likelihood of a process or procedure succeeding without management buy-in is minimal. The teams that are collecting the process data should have the full support of the management team, including periodic reviews and approval of all data collection techniques.
- **Key performance and risk indicators.** You can associate key performance and risk indicators with the data that is being collected. The risk indicators can be used to measure how risky the process, account, facility access or other action is to the organization. The performance indicators can be used to ensure that a process or procedure is successful and measure how much impact it has on the organization's day-to-day operations.
- **Backup verification data.** A strict and rigorous backup procedure is almost useless without verification of the data. Backups should be restored regularly to ensure that the data can be recovered successfully. When using replication, you should also implement integrity checks to ensure that the data was not corrupted during the transfer process.
- **Training and awareness.** Training and awareness of security policies and procedures are half the battle when implementing or maintaining these policies. This extends beyond the security team that is collecting the data, and can impact every employee or user in an organization. The table below outlines different levels of training that can be used for an organization.

	Awareness	Training	Education
Knowledge level	The "what" of a policy or procedure	The "how" of a policy or procedure	The "why" of a policy or procedure
Objective	Knowledge retention	Ability to complete a task	Understanding the big picture
Typical training methods	Self-paced e-learning, web-based training (WBT), videos	Instructor-led training (ILT), demos, hands-on activities	Seminars and research
Testing method	Short quiz after training	Application-level problem solving	Design-level problem solving and architecture exercises

- **Disaster recovery (DR) and business continuity (BC).** Two areas that must be heavily documented are disaster recovery and business continuity. Because these processes are infrequently used, the documentation plays a key role helping teams understand what to do and when to do it. As part of your security assessment and testing, you should review DR and BC documentation to ensure it is complete and represents a disaster from beginning to end. The procedures should adhere to the company's established security policies and answer questions such as, how do administrators obtain system account passwords during a DR scenario? If some sensitive information is required during a DR or BC tasks, you need to ensure this information is both secure and accessible to those who need it.

6.4 Analyze test output and generate reports

The teams that analyze the security procedures should be aware of the output and reporting capabilities for the data. Any information that is of concern must be reported to the management teams immediately so that they are aware of possible risks or alerts. The level of detail given to the management teams might vary depending on their roles and involvement.

The type of auditing being performed can also determine the type of reports that must be used. For example, for an SSAE 16 audit, a Service Organization Control (SOC) report is required. There are four types of SOC reports:

- **SOC 1 Type 1.** This report outlines the findings of an audit, as well as the completeness and accuracy of the documented controls, systems and facilities.
- **SOC 1 Type 2.** This report includes the Type 1 report, along with information about the effectiveness of the procedures and controls in place for the immediate future.
- **SOC 2.** This report includes the testing results of an audit.
- **SOC 3.** This report provides general audit results with a datacenter certification level.

6.5 Conduct or facilitate security audits

Security audits should occur on a routine basis according to the policy set in place by the organization. Internal auditing typically occurs more frequently than external or third-party auditing.

- **Internal.** Security auditing should be an ongoing task of the security team. There are dozens of software vendors that simplify the process of aggregating log data. The challenge is knowing what to look for once you have collected the data.
- **External.** External security auditing should be performed on a set schedule. This could be aligned with financial reporting each quarter or some other business-driven reason.
- **Third-party.** Third-party auditing can be performed on a regular schedule in addition to external auditing. The goal of third-party auditing can either be to provide checks and balances for the internal and external audits, or to perform a more in-depth auditing procedure.

Domain 6 Review Questions

Read and answer the following questions. If you do not get at least one them correct, spend more time with the subject. Then move on to Domain 7.

1. Your company recently implemented a pre-release version of a new email application. The company wants to perform testing against the application and has issued the following requirements:
 - Testers must test all aspects of the email application.
 - Testers must not have any knowledge of the new e-mail environment.

Which type of testing should you use to meet the company requirements?

- a. White box testing
 - b. Black box testing
 - c. Negative testing
 - d. Static testing
 - e. Dynamic testing
2. You are working with your company to validate assessment and audit strategies. The immediate goal is to ensure that all auditors are following the processes and procedures defined by the company's audit policies. Which type of audit should you use for this scenario?
 - a. Internal
 - b. External
 - c. Third-party
 - d. Hybrid
 3. Your company is planning to perform some security control testing. The following requirements have been established:
 - The team must try to bypass controls in the systems.
 - The team can use technical methods or non-technical methods in attempting to bypass controls.

Which type of testing should you perform to meet the requirements?

- a. Vulnerability assessment testing
- b. Penetration testing
- c. Synthetic transaction testing
- d. Misuse case testing

Answers to Domain 6 Review Questions

1. Answer: B

Explanation: In black box testing, testers have no knowledge of the system they are testing.

2. Answer: C

Explanation: Third-party testing is specifically geared to ensuring that the other auditors (internal and external) are properly following your policies and procedures.

3. Answer: B

Explanation: In a penetration test, teams attempt to bypass controls, whether technically or non-technically.

Domain 7. Security Operations

This domain is focused on the day-to-day tasks of securing your environment. If you are in a role outside of operations (such as in engineering or architecture), you should spend extra time in this section to ensure familiarity with the information. You'll notice more hands-on sections in this domain, specifically focused on how to do things instead of the design or planning considerations found in previous domains.

7.1 Understand and support investigations

This section discusses concepts related to supporting security investigations. You should be familiar with the processes in an investigation. You should know all the fundamentals of collecting and handling evidence, documenting your investigation, reporting the information, performing root cause analysis, and performing digital forensic tasks.

- **Evidence collection and handling.** Like a crime scene investigation, a digital investigation involving potential computer crimes has rules and processes to ensure that evidence is usable in court. At a high level, you need to ensure that your handling of the evidence doesn't alter the integrity of the data or environment. To ensure consistency and integrity of data, your company should have an incident response policy that outlines the steps to take in the event of a security incident, with key details such as how employees report an incident. Additionally, the company should have an incident response team that is familiar with the incident response policy and that represents the key areas of the organization (management, HR, legal, IT, etc.). The team doesn't have to be dedicated but instead could have members who have regular work and are called upon only when necessary. With evidence collection, documentation is key. The moment a report comes in, the documentation process begins. As part of the documentation process, you must document each time somebody handles evidence and how that evidence was gathered and moved around; this is known as the chain of custody. Interviewing is often part of evidence collection. If you need to interview an internal employee as a suspect, an HR representative should be present. Consider recording all interviews, if that's legal.
- **Reporting and documenting.** There are two types of reporting: one for IT with technical details and one for management without technical details. Both are critical. The company must be fully aware of the incident and kept up to date as the investigation proceeds. Capture everything possible, including dates, times and pertinent details.
- **Investigative techniques.** When an incident occurs, you need to find out how it happened. A part of this process is the root cause analysis, in which you pinpoint the cause (for example, a user clicked on a malicious link in an email, or a web server was missing a security update and an attacker used an unpatched vulnerability to compromise the server). Often, teams are formed to help determine the root cause. Incident handling is the overall management of the investigation — think of it as project management but on a smaller level. NIST and others have published guidelines for incident handling. At a high level, it includes the following steps: detect, analyze, contain, eradicate and recover. Of course, there are other smaller parts to incident handling, such as preparation and post-incident analysis, like a "lessons learned" review meeting.
- **Digital forensics tools, tactics and procedures.** Forensics should preserve the crime scene, though in digital forensics, this means the computers, storage and other devices, instead of a room and a weapon, for example. Other investigators should be able to perform their own analyses and come to the same conclusions because they

have the same data. This requirement impacts many of the operational procedures. In particular, instead of performing scans, searches and other actions against the memory and storage of computers, you should take images of the memory and storage, so you can thoroughly examine the contents without modifying the originals. For network forensics, you should work from copies of network captures acquired during the incident. For embedded devices, you need to take images of memory and storage and note the configuration. In all cases, leave everything as is, although your organization might have a policy to have everything removed from the network or completely shut down. New technologies can introduce new challenges in this area because sometimes existing tools don't work (or don't work as efficiently) with new technologies. For example, when SSDs were introduced, they presented challenges for some of the old ways of working with disk drives.

7.2 Understand the requirements for different types of investigations

Your investigation will vary based on the type of incident you are investigating. For example, if you work for a financial company and there was a compromise of a financial system, you might have a regulatory investigation. If a hacker defaces your company website, you might have a criminal investigation. Each type of investigation has special considerations:

- **Administrative.** The primary purpose of an administrative investigation is to provide the appropriate authorities with all relevant information so they can determine what, if any, action to take. Administrative investigations are often tied to HR scenarios, such as when a manager has been accused of improprieties.
- **Criminal.** A criminal investigation occurs when a crime has been committed and you are working with a law enforcement agency to convict the alleged perpetrator. In such a case, it is common to gather evidence for a court of law, and to have to share the evidence with the defense. Therefore, you need to gather and handle the information using methods that ensure that the evidence can be used in court. We covered some key points earlier, such as chain of custody. Be sure to remember that in a criminal case, a suspect must be proven guilty beyond a reasonable doubt. This is more difficult than showing a preponderance of evidence, which is often the standard in a civil case.
- **Civil.** In a civil case, one person or entity sues another person or entity; for example, one company might sue another for a trademark violation. A civil case typically seeks monetary damages, not incarceration or a criminal record. As we just saw, the burden of proof is less in a civil case.
- **Regulatory.** A regulatory investigation is conducted by a regulating body, such as the Securities and Exchange Commission (SEC) or Financial Industry Regulatory Authority (FINRA), against an organization suspected of an infraction. In such cases, the organization is required to comply with the investigation, for example, by not hiding or destroying evidence.
- **Industry standards.** An industry standards investigation is intended to determine whether an organization is adhering to a specific industry standard or set of standards, such as logging and auditing failed logon attempts. Because industry standards represent well-understood and widely implemented best practices, many organizations try to adhere to them even when they are not required to do so in order to reduce security, operational and other risks.

7.3 Conduct logging and monitoring activities

This section covers logging and monitoring.

- **Intrusion detection and prevention.** There are two technologies that you can use to detect and prevent intrusions. You should use both. Some solutions combine them into a single software package or appliance.
 - An **intrusion detection system (IDS)** is a technology (typically software or an appliance) that attempts to identify malicious activity in your environment. Solutions often rely on patterns, signatures, or anomalies. There are multiple types of IDS solutions. For example, there are solutions specific to the network (network IDS or NIDS) and others specific to computers (host-based IDS or HIDS).
 - An **intrusion prevention system (IPS)** can help block an attack before it gets inside your network. In the worst case, it can identify an attack in progress. Like an IDS, an IPS is often a software or appliance. However, an IPS is typically placed in line on the network so it can analyze traffic coming into or leaving the network, whereas an IDS typically sees intrusions after they've occurred.
- **Security information and event management (SIEM).** Companies have security information stored in logs across multiple computers and appliances. Often, the information captured in the logs is so extensive that it can quickly become hard to manage and use. Many companies deploy a security information and event management (SIEM) solution to centralize the log data and make it simpler to work with. For example, if you need to find all failed logon attempts on your web servers, you could look through the logs on each web server individually. But if you have a SIEM solution, you can go to a portal and search across all web servers with a single query. A SIEM is a critical technology in large and security-conscious organizations.
- **Continuous monitoring.** Continuous monitoring is the process of streaming information related to the security of the computing environment in real time (or close to real time). Some SIEM solutions offer continuous monitoring or at least some features of continuous monitoring.
- **Egress monitoring.** Egress monitoring is the monitoring of data as it leaves your network. One reason is to ensure that malicious traffic doesn't leave the network (for example, in a situation in which a computer is infected and trying to spread malware to hosts on the internet). Another reason is to ensure that sensitive data (such as customer information or HR information) does not leave the network unless authorized. The following strategies can help with egress monitoring:
 - **Data loss prevention (DLP)** solutions focus on reducing or eliminating sensitive data leaving the network.
 - **Steganography** is the art of hiding data inside another file or message. For example, steganography enables a text message to be hidden inside a picture file (such as a .jpg). Because the file appears innocuous, it can be difficult to detect.
 - **Watermarking** is the act of embedding an identifying marker in a file. For example, you can embed a company name in a customer database file or add a watermark to a picture file with copyright information.

7.4 Securely provision resources

This section covers provisioning of resources. Therefore, we cover the topics from a provisioning standpoint, rather than an overall management or design standpoint.

- **Asset inventory.** You need to have a method for maintaining an accurate inventory of your company's assets. For example, you need to know how many computers you have and how many installations of each licensed software application you have. Asset inventory helps organizations protect physical assets from theft, maintain software licensing compliance, and account for the inventory (for example, depreciating the assets). There are other benefits too. For example, if a vulnerability is identified in a specific version of an application, you can use your asset inventory to figure out whether you have any installations of the vulnerable version.
- **Asset management.** Assets, such as computers, desks and software applications, have a lifecycle — simply put, you buy it, you use it and then you retire it. Asset management is the process of managing that lifecycle. You keep track of all your assets, including when you got it, how much you paid for it, its support model and when you need to replace it. For example, asset management can help your IT team figure out which laptops to replace during the next upgrade cycle. It can also help you control costs by finding overlap in hardware, software or other assets.
- **Configuration management.** Configuration management helps you standardize a configuration across your devices. For example, you can use configuration management software to ensure that all desktop computers have anti-virus software and the latest patches, and that the screen will automatically be locked after 5 minutes of inactivity. The configuration management system should automatically remediate most changes users make to a system. The benefits of configuration management include having a single configuration (for example, all servers have the same baseline services running and the same patch level), being able to manage many systems as a single unit (for example, you can deploy an updated anti-malware application to all servers the same amount of time it takes to deploy it to a single server), and being able to report on the configuration throughout your network (which can help to identify anomalies). Many configuration management solutions are OS-agnostic, meaning that they can be used across Windows, Linux and Mac computers. Without a configuration management solution, the chances of having a consistent and standardized deployment plummets, and you lose the efficiencies of configuring many computers as a single unit.

7.5 Understand and apply foundational security operations concepts

This section covers some of the foundational items for security operations. Many of these concepts apply to several other sections on the exam. You should have a very firm grasp of these topics so that you can navigate them effectively throughout the other sections.

- **Need-to-know and least privilege.** Access should be given based on a need to know. For example, a system administrator who is asked to disable a user account doesn't need to know that the user was terminated, and a systems architect who is asked to evaluate an IT inventory list doesn't need to know that his company is considering acquiring another company. The principle of least privilege means giving users the fewest privileges they need to perform their job tasks; entitlements are granted only after a specific privilege is deemed necessary. It is a good practice and almost always a recommend practice. Two other concepts are important here:
 - **Aggregation.** The combining of multiple things into a single unit is often used in role-based access control.

- **Transitive trust.** From a Microsoft Active Directory perspective, a root or parent domain automatically trusts all child domains. Because of the transitivity, all child domains also trust each other. Transitivity makes it simpler to have trusts. But it is important to be careful. Consider outside of Active Directory: If Chris trusts Terry and Pat trusts Terry, should Chris trust Pat? Probably not. In high-security environments, it isn't uncommon to see non-transitive trusts used, depending on the configuration and requirements.
- **Separation of duties and responsibilities.** Separation of duties refers to the process of separating certain tasks and operations so that a single person doesn't control all them. For example, you might dictate that one person is the security administrator and another is the email administrator. Each has administrative access to only their area. You might have one administrator responsible for authentication and another responsible for authorization. The goal with separation of duties is to make it more difficult to cause harm to the organization (via destructive actions or data loss, for example). With separation of duties, it is often necessary to have two or more people working together (colluding) to cause harm to the organization. Separation of duties is not always practical, though. For example, in a small company, you might only have one person doing all the IT work, or one person doing all the accounting work. In such cases, you can rely on compensating controls or external auditing to minimize risk.
- **Privileged account management.** A special privilege is a right not commonly given to people. For example, certain IT staff might be able to change other users' passwords or restore a system backup, and only certain accounting staff can sign company checks. Actions taken using special privileges should be closely monitored. For example, each user password reset should be recorded in a security log along with pertinent information about the task: date and time, source computer, the account that had its password changed, the user account that performed the change, and the status of the change (success or failure). For high-security environments, you should consider a monitoring solution that offers screen captures or screen recording in addition to the text log.
- **Job rotation.** Job rotation is the act of moving people between jobs or duties. For example, an accountant might move from payroll to accounts payable and then to accounts receivable. The goal of job rotation is to reduce the length of one person being in a certain job (or handling a certain set of responsibilities) for too long, which minimizes the chances of errors or malicious actions going undetected. Job rotation can also be used to cross-train members of teams to minimize the impact of an unexpected leave of absence.
- **Information lifecycle.** Information lifecycle is made up of the following phases:
 - **Collect data.** Data is gathered from sources such as log files and inbound email, and when users produce data such as a new spreadsheet.
 - **Use data.** Users read, edit and share data.
 - **Retain data (optional).** Data is archived for the time required by the company's data retention policies. For example, some companies retain all email data for 7 years by archiving the data to long-term storage until the retention period has elapsed.
 - **Legal hold (occasional).** A legal hold requires you to maintain one or more copies of specified data in an unalterable form during a legal scenario (such as a lawsuit) or an audit or government investigation. A legal hold is often narrow; for example, you might have to put a legal hold on all email to or from the accounts payable department. In most cases, a legal hold is invisible to users and administrators who are not involved in placing the hold.

- **Delete data.** The default delete action in most operating systems is not secure: The data is marked as deleted, but it still resides on the disks and can be easily recovered with off-the-shelf software. To have an effective information lifecycle, you must use secure deletion techniques such as disk wiping (for example, by overwriting the data multiple times), degaussing and physical destruction (shredding a disk).
- **Service-level agreements (SLAs).** An SLA is an agreement between a provider (which could be an internal department) and the business that defines when a service provided by the department is acceptable. For example, the email team might have an SLA that dictates that they will provide 99.9% uptime each month or that spam email will represent 5% or less of the email in user mailboxes. SLAs can help teams design appropriate solutions. For example, if an SLA requires 99.9% uptime, a team might focus on high availability and site resiliency. Sometimes, especially with service providers, not adhering to SLAs can result in financial penalties. For example, an internet service provider (ISP) might have to reduce its monthly connection charges if it does not meet its SLA.

7.6 Apply resource protection techniques

This section covers media, hardware and software management. We will look at some key tips for managing media and using asset management for software and hardware.

- **Media management.** Media management is the act of maintaining media for your software and data. This includes operating system images, installation files and backup media. Any media that you use in your organization potentially falls under this umbrella. There are some important media management concepts to know:
 - **Source files.** If you rely on software for critical functions, you need to be able to reinstall that software at any time. Despite the advent of downloadable software, many organizations rely on legacy software that they purchased on disk years ago and that is no longer available for purchase. To protect your organization, you need to maintain copies of the media along with copies of any license keys.
 - **Operating system images.** You need a method to manage your operating system images so that you can maintain clean images, update the images regularly (for example, with security updates), and use the images for deployments. Not only should you maintain multiple copies at multiple sites, but you should also test the images from time to time. While you can always rebuild an image from your step-by-step documentation, that lost time could cost your company money during an outage or other major issue.
 - **Backup media.** Backup media is considered sensitive media. While many organizations encrypt backups on media, you still need to treat the backup media in a special way to reduce the risk of it being stolen and compromised. Many companies lock backup media in secure containers and store the containers in a secure location. It is also common to use third-party companies to store backup media securely in off-site facilities.
- **Hardware and software asset management.** At first glance, asset management might not seem related to security operations, but it actually is. For example, if a vendor announces a critical vulnerability in a specific version of a product that allows remote code execution, you need to quickly act to patch your devices — which means you need to be able to quickly figure out if you have any devices that are vulnerable. You can't do that without effective asset management (and, in some cases, configuration management). Here are some key tasks for an asset management solution:

- **Capture as much data as you reasonably can.** You need to know where a given product is installed. But you also need to know when it was installed (for example, whether a vulnerable version was installed after the company announced the vulnerability), the precise version number (because without that, you might not be able to effectively determine whether you are susceptible), and other details.
- **Have a robust reporting system.** You need to be able to use all the asset management data you collect, so you need a robust reporting system that you can query on demand. For example, you should be able to quickly get a report listing all computers running a specific version of a specific software product. And you should then be able to filter that data to only corporate-owned devices or laptop computers.
- **Integrate asset management with other automation software.** If your asset management solution discovers 750 computers running a vulnerable version of a piece of software, you need an automated way to update the software to the latest version. You can do that by integrating your asset management system with your configuration management system. Some vendors offer an all-in-one solution that performs both asset management and configuration management.

7.7 Conduct incident management

Incident management is the management of incidents that are potentially damaging to an organization, such as a distributed denial of service attack. Not all incidents are computer-related; for example, a break-in at your CEO's office is also an incident.

- **Detection.** It is critical to be able to detect incidents quickly because they often become more damaging at time passes. It is important to have a robust monitoring and intrusion detection solution in place. Other parts of a detection system include security cameras, motion detectors, smoke alarms and other sensors. If there is a security incident, you want to be alerted (for example, if an alarm is triggered at your corporate headquarters over a holiday weekend).
- **Response.** When you receive a notification about an incident, you should start by verifying the incident. For example, if an alarm was triggered at a company facility, a security guard can physically check the surroundings for an intrusion and check the security cameras for anomalies. For computer-related incidents, it is advisable to keep compromised systems powered on to gather forensic data. Along with the verification process, during the response phase you should also kick off the initial communication with teams or people that can help with mitigation. For example, you should contact the information security team initially during a denial-of-service attack.
- **Mitigation.** The next step is to contain the incident. For example, if a computer has been compromised and is actively attempting to compromise other computers, the compromised computer should be removed from the network to mitigate the damage.
- **Reporting.** Next, you should disseminate data about the incident. You should routinely inform the technical teams and the management teams about the latest findings regarding the incident.
- **Recovery.** In the recovery phase, you get the company back to regular operations. For example, for a compromised computer, you re-image it or restore it from a backup. For a broken window, you replace it.

- **Remediation.** In this phase, you take additional steps to minimize the chances of the same or a similar attack being successful. For example, if you suspect that an attacker launched attacks from the company's wireless network, you should update the wireless password or authentication mechanism. If an attacker gained access to sensitive plain text data during an incident, you should encrypt the data in the future.
- **Lessons learned.** During this phase, all team members who worked on the security incident gather to review the incident. You want to find out which parts of the incident management were effective and which were not. For example, you might find that your security software detected an attack immediately (effective) but you were unable to contain the incident without powering off all the company's computers (less effective). The goal is to review the details to ensure that the team is better prepared for the next incident.

7.8 Operate and maintain detective and preventative measures

This section deals with the hands-on work of operating and maintaining security systems to block attacks on your company's environment or minimize their impact.

- **Firewalls.** While operating firewalls often involves adding and editing rules and reviewing logs, there are other tasks that are important, too. For example, review the firewall configuration change log to see which configuration settings have been changed recently.
- **Intrusion detection and prevention systems.** You need to routinely evaluate the effectiveness of your IDS and IPS systems. You also need to review and fine-tune the alerting functionality. If too many alerts are sent (especially false positive or false negatives), administrators will often ignore or be slow to respond to alerts, causing response to a real incident alert to be delayed.
- **Whitelisting and blacklisting.** Whitelisting is the process of marking applications as allowed, while blacklisting is the process of marking applications as disallowed. Whitelisting and blacklisting can be automated. It is common to whitelist all the applications included on a corporate computer image and disallow all others.
- **Security services provided by third parties.** Some vendors offer security services that ingest the security-related logs from your entire environment and handle detection and response using artificial intelligence or a large network operations center. Other services perform assessments, audits or forensics. Finally, there are third-party security services that offer code review, remediation or reporting.
- **Sandboxing.** Sandboxing is the act of totally segmenting an environment or a computer from your production networks and computers; for example, a company might have a non-production environment on a physically separate network and internet connection. Sandboxes help minimize damage to a production network. Because computers and devices in a sandbox aren't managed in the same way as production computers, they are often more vulnerable to attacks and malware. By segmenting them, you reduce the risk of those computers infecting your production computers. Sandboxes are also often used for honeypots and honeynets, as explained in the next bullet.
- **Honeypots and honeynets.** A honeypot or a honeynet is a computer or network purposely deployed to lure would-be attackers and record their actions. The goal is to understand their methods and use that knowledge to design more secure computers and networks. There are important and accepted uses; for example, an anti-virus software company might use honeypots to validate and strengthen their anti-virus and anti-malware software.

However, honeypots and honeynets have been called unethical because of their similarities to entrapment. While many security-conscious organizations stay away from running their own honeypots and honeynets, they can still take advantage of the information gained from other companies that use them.

- **Anti-malware.** Anti-malware is a broad term that often includes anti-virus, anti-spam and anti-malware (with malware being any other code, app or service created to cause harm). You should deploy anti-malware to every possible device, including servers, client computers, tablets and smartphones, and be vigilant about product and definition updates.

7.9 Implement and support patch and vulnerability management

While patch management and vulnerability management seem synonymous, there are some key differences:

- **Patch management.** The updates that software vendors provide to fix security issues or other bugs are called patches. Patch management is the process of managing all the patches in your environment, from all vendors. A good patch management system tests and implements new patches immediately upon release to minimize exposure. Many security organizations have released studies claiming that the single most important part of securing an environment is having a robust patch management process that moves swiftly. A patch management system should include the following processes:
 - **Automatic detection and download of new patches.** Detection and downloading should occur at least once per day. You should monitor the detection of patches so that you are notified if detection or downloading is not functional.
 - **Automatic distribution of patches.** Initially, deploy patches to a few computers in a lab environment and run them through system testing. Then expand the distribution to a larger number of non-production computers. If everything is functional and no issues are found, distribute the patches to the rest of the non-production environment and then move to production. It is a good practice to patch your production systems within 7 days of a patch release. In critical scenarios where there is known exploit code for a remote code execution vulnerability, you should deploy patches to your production systems the day of the patch release to maximize security.
 - **Reporting on patch compliance.** Even if you might have an automatic patch distribution method, you need a way to assess your overall compliance. Do 100% of your computers have the patch? Or 90%? Which specific computers are missing a specific patch? Reporting can be used by the management team to evaluate the effectiveness of a patch management system.
 - **Automatic rollback capabilities.** Sometimes, vendors release patches that create problems or have incompatibilities. Those issues might not be evident immediately but instead show up days later. Ensure you have an automated way of rolling back or removing the patch across all computers. You don't want to figure that out a few minutes before you need to do it.
- **Vulnerability management.** A vulnerability is a way in which your environment is at risk of being compromised or degraded. The vulnerability can be due to a missing patch. But it can also be due to a misconfiguration or other factors. For example, when SHA-1 certificates were recently found to be vulnerable to attack, many companies

suddenly found themselves vulnerable and needed to take action (by replacing the certificates). Many vulnerability management solutions can scan the environment looking for vulnerabilities. Such solutions complement, but do not replace, patch management systems and other security systems (such as anti-virus or anti-malware systems). Be aware of the following definitions:

- **Zero-day vulnerability.** A vulnerability is sometimes known about before a patch is available. Such zero-day vulnerabilities can sometimes be mitigated with an updated configuration or other temporary workaround until a patch is available. Other times, no mitigations are available and you have to be especially vigilant with logging and monitoring until the patch is available.
- **Zero-day exploit.** Attackers can release code to exploit a vulnerability for which no patch is available. These zero-day exploits represent one of the toughest challenges for organizations trying to protect their environments.

7.10 Understand and participate in change management processes

Change management represents a structured way of handling changes to an environment. The goals include providing a process to minimize risk, improving the user experience, and providing consistency with changes. While many companies have their own change management processes, there are steps that are common across most organizations:

- **Identify the need for a change.** For example, you might find out that your routers are vulnerable to a denial of service attack and you need to update the configuration to remedy that.
- **Test the change in a lab.** Test the change in a non-production environment to ensure that the proposed change does what you think it will. Also use the test to document the implementation process and other key details.
- **Put in a change request.** A change request is a formal request to implement a change. You specify the proposed date of the change (often within a pre-defined change window), the details of the work, the impacted systems, notification details, testing information, rollback plans and other pertinent information. The goal is to have enough information in the request that others can determine whether there will be any impact to other changes or conflicts with other changes and be comfortable moving forward. Many companies require a change justification for all changes.
- **Obtain approval.** Often, a change control board (a committee that runs change management), will meet weekly or monthly to review change requests. The board and the people that have submitted the changes meet to discuss the change requests, ask questions and vote on approval. If approval is granted, you move on to the next step. If not, you restart the process.
- **Send out notifications.** A change control board might send out communications about upcoming changes. In some cases, the implementation team handles the communications. The goal is to communicate to impacted parties, management and IT about the upcoming changes. If they see anything unusual after a change is made, the notifications will help them begin investigating by looking at the most recent changes.
- **Perform the change.** While most companies have defined change windows, often on the weekend, sometimes a change can't wait for that window (such as an emergency change). During the change process, capture the existing

configuration, capture the changes and steps, and document all pertinent information. If a change is unsuccessful, perform the rollback plan steps.

- **Send out “all clear” notifications.** These notifications indicate success or failure.

7.11 Implement recovery strategies

A recovery operation takes place following an outage, security incident or other disaster that takes an environment down or compromises it in a way that requires restoration. Recovery strategies are important because they have a big impact on how long your organization will be down or have a degraded environment, which has an impact on the company's bottom line. Note that this section focuses on strategies rather than tactics, so be thinking from a design perspective, not from a day-day-day operational perspective.

- **Backup storage strategies.** While most organizations back up their data in some way, many do not have an official strategy or policy regarding where the backup data is stored or how long the data is retained. In most cases, backup data should be stored offsite. Offsite backup storage provides the following benefits:
 - If your data center is destroyed (earthquake, flood, fire), your backup data isn't destroyed with it. In some cases, third-party providers of off-site storage services also provide recovery facilities to enable organizations to recover their systems to the provider's environment.
 - Offsite storage providers provide environmentally sensitive storage facilities with high-quality environmental characteristics around humidity, temperature and light. Such facilities are optimal for long-term backup storage.
 - Offsite storage providers provide additional services that your company would have to manage otherwise, such as tape rotation (delivery of new tapes and pickup of old tapes), electronic vaulting (storing backup data electronically), and organization (cataloging of all media, dates and times).
- **Recovery site strategies.** When companies have multiple data centers, they can often use one as a primary data center and one another as a recovery site (either a cold standby site or a warm standby site). An organization with 3 or more data centers can have a primary data center, a secondary data center (recovery site) and regional data centers. With the rapid expansion of public cloud capabilities, having a public cloud provider be your recovery site is feasible and reasonable. One key thing to think about is cost. While cloud storage is inexpensive and therefore your company can probably afford to store backup data there, trying to recover your entire data center from the public cloud might not be affordable or fast enough.
- **Multiple processing sites.** Historically, applications and services were highly available within a site such as a data center, but site resiliency was incredibly expensive and complex. Today, it is common for companies to have multiple data centers, and connectivity between the data centers is much faster and less expensive. Because of these advances, many applications provide site resiliency with the ability to have multiple instances of an application spread across 3 or more data centers. In some cases, application vendors are recommending backup-free designs in which an app and its data are stored in 3 or more locations, with the application handling the multi-site syncing. The public cloud can be the third site, which is beneficial for companies that lack a third site or that have apps and services already in the public cloud.

- **System resilience, high availability, quality of service (QoS) and fault tolerance.** To prepare for the exam, it is important to know the differences between these related terms:
 - **System resilience.** Resilience is the ability to recover quickly. For example, site resilience means that if Site 1 goes down, Site 2 quickly and seamlessly comes online. Similarly, with system resilience, if a disk drive fails, another (spare) disk drive is quickly and seamlessly added to the storage pool. Resilience often comes from having multiple functional components (for example, hardware components).
 - **High availability.** While resilience is about recovering with a short amount of downtime or degradation, high availability is about having multiple redundant systems that enable zero downtime or degradation for a single failure. For example, if you have a highly available database cluster, one of the nodes can fail and the database cluster remains available without an outage or impact. While clusters are often the answer for high availability, there are many other methods available too. For instance, you can provide a highly available web application by using multiple web servers without a cluster. Many organizations want both high availability and resiliency.
 - **Quality of service (QoS).** QoS is a technique that helps enable specified services to receive a higher quality of service than other specified services. For example, on a network, QoS might provide the highest quality of service to the phones and the lowest quality of service to social media. QoS has been in the news because of the net neutrality discussion taking place in the United States. The new net neutrality law gives ISPs a right to provide higher quality of services to a specified set of customers or for a specified service on the internet. For example, an ISP might opt to use QoS to make its own web properties perform wonderfully while ensuring the performance of its competitors' sites is subpar.
 - **Fault tolerance.** As part of providing a highly available solution, you need to ensure that your computing devices have multiple components — network cards, processors, disk drives, etc. —of the same type and kind to provide fault tolerance. Fault tolerance, by itself, isn't valuable. For example, imagine a server with fault-tolerant CPUs. The server's power supply fails. Now the server is done even though you have fault tolerance. As you can see, you must account for fault tolerance across your entire system and across your entire network.

7.12 Implement disaster recovery (DR) recovery processes

Trying to recover from a disaster without a documented disaster recovery processes is difficult, if not impossible. Thus, you should establish clear disaster recovery processes to minimize the effort and time required to recover from a disaster. Testing the plans is also important and is discussed separately in the next section (7.13).

- **Response.** When you learn about an incident, the first step is to determine whether it requires a disaster recovery procedure. Timeliness is important because if a recovery is required, you need to begin recovery procedures as soon as possible. Monitoring and alerting play a big part in enabling organizations to respond to disasters faster.
- **Personnel.** In many organizations, there is a team dedicated to disaster recovery planning, testing and implementation. They maintain the processes and documentation. In a disaster recovery scenario, the disaster recovery team should be contacted first so they can begin communicating to the required teams. In a real disaster, communicating with everybody will be difficult and, in some cases, not possible. Sometimes, companies use

communication services or software to facilitate emergency company-wide communications or mass communications with personnel involved in the disaster recovery operation.

- **Communications.** There are two primary forms of communication that occur during a disaster recovery operation, as well as a third form of communication that is sometimes required:
 - **Communications with the recovery personnel.** In many disaster scenarios, email is down, phones are down, and instant messaging services are down. If the disaster hasn't taken out cell service, you can rely on communications with smart phones (SMS messages, phone calls).
 - **Communications with the management team and the business.** As the recovery operation begins, the disaster recovery team must stay in regular contact with the business and the management team. The business and management team need to understand the severity of the disaster and the approximate time to recover. As things progress, they must be updated regularly.
 - **Communications with the public.** In some cases, a company experiencing a large-scale disaster must communicate with the public, for example, a service provider, a publicly traded company, or a provider of services to consumers. At a minimum, the communication must indicate the severity of the incident, when service is expected to resume, and any actions consumers need to take.
- **Assessment.** During the response phase, the teams verified that recovery procedures had to be initiated. In the assessment phase, the teams dive deeper to look at the specific technologies and services to find out details of the disaster. For example, if during the response phase, the team found email to be completely down, then they might check to find out if other technologies are impacted along with email.
- **Restoration.** During the restoration phase, the team performs the recovery operations to bring all services back to their normal state. In many situations, this means failing over to a secondary data center. In others, it might mean recovering from backups. After a successful failover to a secondary data center, it is common to start planning the failback to the primary data center once it is ready. For example, if the primary data center flooded, you would recover to the second data center, recover from the flood, then fail back to the primary data center.
- **Training and awareness.** To maximize the effectiveness of your disaster recovery procedures, you need to have a training and awareness campaign. Sometimes, technical teams will gain disaster recovery knowledge while attending training classes or conferences for their technology. But they also need training about your organization's disaster recovery procedures and policies. Performing routine tests of your disaster recovery plans can be part of such training. That topic is covered next, in section 7.13.

7.13 Test disaster recovery plans (DRP)

Testing your disaster recovery plans is an effective way to ensure your company is ready for a real disaster. It also helps minimize the amount of time it takes to recover from a real disaster, which can benefit a company financially. There are multiple ways of testing your plan:

- **Read-through/tabletop.** The disaster recovery teams (for example, server, network, security, database, email, etc.) gather and the disaster recovery plan is read. Each team validates that their technologies are present and the timing is appropriate to ensure that everything can be recovered. If not, changes are made. A read-through can

often help identify ordering issues (for example, trying to recover email before recovering DNS) or other high-level issues. In a read-through exercise, teams do not perform any recovery operations.

- **Walkthrough.** A walkthrough is a more detailed read-through — the same teams look at the details of the recovery operations to look for errors, omissions or other problems.
- **Simulation.** A simulation is a simulated disaster in which teams must go through their documented recovery operations. Simulations are very helpful to validate the detailed recovery plans and help the teams gain experience performing recovery operations.
- **Parallel.** In a parallel recovery effort, teams perform recovery operations on a separate network, sometimes in a separate facility. Some organizations use third-party providers that provide recovery data centers to perform parallel recovery tests. Companies sometimes use a parallel recovery method to minimize disruption to their internal networks and minimize the need to maintain the IT infrastructure necessary to support recovery efforts.
- **Full interruption.** In a full interruption recovery, the organizations halt regular operations on a separate network, sometimes in a separate facility. Many times, a full interruption operation involves failing over from the primary data center to the secondary data center. This type of recovery testing is the most expensive, takes the most time, and exposes the company to the most risk of something going wrong. While those drawbacks are serious, full interruption tests are a good practice for most organizations.

7.14 Participate in business continuity (BC) planning and exercises

Business continuity includes disaster recovery, but it covers other things as well. Disaster recovery is a very specific series of processes to recovery from a disaster. Business continuity focuses on ensuring the business experiences minimal or no downtime (with the hope that a disaster recovery process won't be needed). Think of business continuity as a strategy and disaster recovery as a tactic. The bullets below detail the steps required to plan business continuity. Note that these steps can be used to build a disaster recovery plan too.

- **Plan for an unexpected scenario.** Form a team, perform a business impact analysis for your technologies, identify a budget and figure out which business processes are mission-critical.
- **Review your technologies.** Set the recovery time objective and recovery point objective, develop a technology plan, review vendor support contracts, and create or review disaster recovery plans.
- **Build a communication plan.** Finalize who needs to be contacted, figure out primary and alternative contact methods, and ensure that everybody can work, possibly from a backup location.
- **Coordinate with external entities.** Work with relevant external entities, such as the police department, government agencies, partner companies and the community.

7.15 Implement and manage physical security

Physical security represents securing your physical assets such as land, buildings, computers and other company property.

- **Perimeter security controls.** The perimeter is the external facility surrounding your buildings or other areas, such as the space just outside of a data center. Two key considerations are access control and monitoring:
 - **Access control.** To maximize security, your facilities should restrict who can enter. This is often handled by key cards and card readers on doors. Other common methods are a visitor center or reception area with security guards and biometric scanners for entry (often required for data centers).
 - **Monitoring.** As part of your perimeter security, you should have a solution to monitor for anomalies. For example, if a door with a card reader is open for more than 60 seconds, it could indicate that it has been propped open. If a person scans a data center door with a badge but that badge wasn't used to enter any other exterior door on that day, it could be a scenario to investigate — for example, maybe the card was stolen by somebody who gained access to the building through the air vents. A monitoring system can alert you to unusual scenarios and provide a historical look at your perimeter activities.
- **Internal security controls.** Internal security focuses on limiting access to storage or supply rooms, filing cabinets, telephone closets, data centers and other sensitive areas. There are a couple of key methods to use:
 - **Escort requirements.** When a visitor checks in at your visitor center, you can require an employee escort. For example, maybe the visitor is required to always be with an employee and the guest badge does not open doors via the door card readers. Escort requirements are especially important for visitors who will be operating in sensitive areas (for example, an air conditioning company working on a problem in your data center).
 - **Key and locks.** Each employee should have the ability to secure company and personal belongings in their work space to help prevent theft. If they have an office, they should lock it when they aren't in the office. If the employee has a desk or cubicle, they should have lockable cabinets or drawers for storing sensitive information and other valuables.

7.16 Address personnel safety and security concerns

This section covers personnel safety — making sure employees can safely work and travel. While some of the techniques are common sense, others are less obvious.

- **Travel.** The laws and policies in other countries can sometimes be drastically different than your own country. Employees must be familiar with the differences prior to traveling. For example, something you see as benign might be illegal and punishable by jail in another country. Other laws could make it difficult to do business in another country or put your company at risk. When traveling to other countries, you should familiarize yourself with the local laws to minimize danger to yourself and your company. Another key concern when traveling is protecting company data. To protect company data during travel, encryption should be used for both data in transit and data at rest. It is also a good practice (although impractical) to limit connectivity via wireless networks while traveling. Take your computing devices with you, when possible, since devices left in a hotel are subject to

tampering. In some cases, such as when traveling to high-risk nations, consider having personnel leave their computing devices at home. While this isn't always feasible, it can drastically reduce the risk to personnel and company devices or data. In some organizations, employees are given a special travel laptop that has been scrubbed of sensitive data to use during a trip; the laptop is re-imaged upon return home.

- **Security training and awareness.** Employees should be trained about how to mitigate potential dangers in the home office, while traveling or at home. For example, campus safety includes closing doors behind you, not walking to your car alone after hours, and reporting suspicious persons. Travel safety includes not displaying your company badge in public places and taking only authorized ride hailing services. Safety outside of work includes using a secure home network and not inserting foreign media into devices. While the training and awareness campaigns will differ, a key element is to have a campaign that addresses your organization's particular dangers.
- **Emergency management.** Imagine a large earthquake strikes your primary office building. The power is out, and workers have evacuated the buildings; many go home to check on their families. Other employees might be flying to the office for meetings the next day. You need to be able to find out if all employees are safe and accounted for; notify employees, partners, customers, and visitors; and initiate business continuity and/or disaster recovery procedures. An effective emergency management system enables you to send out emergency alerts to employees (many solutions rely on TXT or SMS messages to cellular phones), track their responses and locations, and initiate emergency response measures, such as activating a secondary data center or a contingent workforce in an alternate site.
- **Duress.** Duress refers forcing somebody to perform an act that they normally wouldn't, due to a threat of harm, such as a bank teller giving money to a bank robber who brandishes a weapon. Training personnel about duress and implementing countermeasures can help. For example, at a retail store, the last twenty-dollar bill in the cash register can be attached to a silent alarm mechanism; when an employee removes it for a robber, the silent alarm alerts the authorities. Another example is a building alarm system that must be deactivated quickly once you enter the building. If the owner of a business is met at opening time by a crook who demands that she deactivates the alarm, instead of entering her regular disarm code, the owner can use a special code that deactivates the alarm and notifies the authorities that it was disarmed under duress. In many cases, to protect personnel safety, it is a good practice to have personnel fully comply with all reasonable demands, especially in situations where the loss is a laptop computer or something similar.

Domain 7 Review Questions

Read and answer the following questions. If you do not get one at least one of them correct, spend more time with the subject. Then move on to Domain 8.

1. You are conducting an analysis of a compromised computer. You figure out that the computer had all known security patches applied prior to the computer being compromised. Which two of the following statements are probably true about this incident?
 - e. The company has a zero-day vulnerability.
 - f. The company was compromised by a zero-day exploit.
 - g. The computer does not have a configuration management agent.
 - h. The computer does not have anti-malware.
2. You are investigating poor performance of a company's telephone system. The company uses IP-based phones and reports that in some scenarios, such as when there is heavy use, the call quality drops and there are sometimes lags or muffling. You need to maximize the performance of the telephone system. Which technology should you use?
 - a. System resilience
 - b. Quality of service
 - c. Fault tolerance
 - d. Whitelisting
 - e. Blacklisting
 - f. Configuration management
3. You are preparing your company for disaster recovery. The company issues the following requirements for disaster recovery testing:
 - The company must have the ability to restore and recover to an alternate data center.
 - Restore and recovery operations must not impact your data center.
 - IT teams must perform recovery steps during testing.

Which type of recovery should you use to meet the company's requirements?

- a. Partial interruption
- b. Tabletop
- c. Full interruption
- d. Parallel

Answers to Domain 7 Review Questions

1. Answer: A, B

Explanation: When a vulnerability exists but there is no patch to fix it, it is a zero-day vulnerability. When exploit code exists to take advantage of a zero-day vulnerability, it is called a zero-day exploit. In this scenario, because the computer was up to date on patches, we can conclude that there was a zero-day vulnerability and a zero-day exploit.

2. Answer: B

Explanation: Quality of service provides priority service to a specified application or type of communication. In this scenario, call quality is being impacted by other services on the network. By prioritizing the network communication for the IP phones, you can maximize their performance (though that might impact something else).

3. Answer: D

Explanation: The first key requirement in this scenario is that the data center must not be impacted by the testing. This eliminates the partial interruption and full interruption tests because those impact the data center. The other key requirement is that IT teams must perform recovery steps. This requirement eliminates the tabletop testing because tabletop testing involves walking through the plans, but not performing recovery operations.

Domain 8. Software Development Security

This domain focuses on managing the risk and security of software development. Security should be a focus of the development lifecycle, and not an add-on or afterthought to the process. The development methodology and lifecycle can have a big effect on how security is thought of and implemented in your organization. The methodology also ties into the environment that the software is being developed for. Organizations should ensure that access to code repositories is limited to protect their investment in software development. Access and protection should be audited on a regular basis. You must also take into consideration the process of acquiring a development lifecycle, whether from another company, or picking up a development project that is already in progress.

8.1 Understand and integrate security throughout the software development lifecycle (SDLC)

This section discusses the various methods and considerations when developing an application. The lifecycle of development does not typically have a final goal or destination. Instead, it is a continuous loop of efforts that must include steps at different phases of a project.

- **Development methodologies.** There are many different development methodologies that organizations can use as part of the development lifecycle. The following table lists the most common methodologies and the key related concepts.

Methodology	Key Concepts
Build and fix	• Lacks a key architecture design • Problems are fixed as they occur • Lacks a formal feedback cycle • Reactive instead of proactive
Waterfall	• Linear sequential lifecycle • Each phase is completed before continuing • Lacks a formal way to make changes during a cycle • Project is completed before collecting feedback and starting again
V-shaped	• Based on the waterfall model • Each phase is complete before continuing • Allows for verification and validation after each phase • Does not contain a risk analysis phase
Prototyping	• Three main models: • Rapid prototyping uses a quick sample to test the current project. • Evolutionary prototyping uses incremental improvements to a design. • Operational prototypes provide incremental improvements, but are intended to be used in production.
Incremental	• Uses multiple cycles for development (think multiple waterfalls) • The entire process can restart at any time as a different phase • Easy to introduce new requirements

	• Delivers incremental updates to software
Spiral	• Iterative approach to development • Performs risk analysis during development • Future information and requirements are funneled into the risk analysis • Allows for testing early in development
Rapid application development	• Uses rapid prototyping • Designed for quick development • Analysis and design are quickly demonstrated • Testing and requirements are often revisited
Agile	• Umbrella term for multiple methods • Highlights efficiency and iterative development • User stories describe what a user does and why • Prototypes are filtered down to individual features

- **Maturity models.** There are five maturity levels of the Capability Maturity Model Integration (CMMI):
 - 1. Initial.** The development process is ad hoc, inefficient, inconsistent and unpredictable.
 - 2. Repeatable.** A formal structure provides change control, quality assurance and testing.
 - 3. Defined.** Processes and procedures are designed and followed during the project.
 - 4. Managed.** Processes and procedures are used to collect data from the development cycle to make improvements.
 - 5. Optimizing.** There is a model of continuous improvement for the development cycle.
- **Operation and maintenance.** After a product has been developed, tested and released, the next phase of the process is to provide operational support and maintenance of the released product. This can include resolving unforeseen problems or developing new features to address new requirements.
- **Change management.** Changes can disrupt development, testing and release. An organization should have a change control process that includes documenting and understanding a change before attempting to implement it. This is especially true the later into the project the change is requested. Each change request must be evaluated for capability, risk and security concerns, impacts to the timeline, and more.
- **Integrated product team.** Software development and IT have typically been two separate departments or groups within an organization. Each group typically has different goals: developers want to distribute finished code, and IT wants to efficiently manage working systems. With DevOps, these teams work together to align their goals so that software releases are consistent and reliable.

8.2 Identify and apply security controls in development environments

The source code and repositories that make up an application can represent hundreds or thousands of hours of work and comprise important intellectual property for an organization. Organizations must be prepared to take multiple levels of risk mitigation to protect the code, as well as the applications.

- **Security of the software environments.** Historically, security has been an afterthought or a bolt-on after an application has been developed and deployed, instead of a part of the lifecycle. When developing an application, considerations must be made for the databases, external connections and sensitive data that are being handled by the application.
- **Security weaknesses and vulnerabilities at the source-code level.** The MITRE organization publishes a list of the 25 most dangerous software errors that can cause weaknesses and vulnerabilities in an application (<http://cwe.mitre.org/top25/#Listing>). For example, if an input field is not verified for content and length, then unexpected errors can occur. Additionally, if file access or encryption is lacking in an application, then users could potentially access information that they do not have permissions for. Code reviews, static analysis, testing and validation can all help mitigate risks in developing software.
- **Configuration management as an aspect of secure coding.** The change control process should be tightly integrated with development to ensure that security considerations are made for any new requirements, features or requests. A centralized code repository helps in managing changes and tracking when and where revisions to the code. The repository can track versions of an application so you can easily roll back to a previous version if necessary.
- **Security of code repositories.** The version control system that houses source code and intellectual property is the code repository. There might be different repositories for active development, testing and quality assurance. A best practice for securing code repositories is to ensure that they are as far away from the internet as possible, even if that means that they are on a separate internal network that does not have internet access. Any remote access to a repository should use a VPN or another secure connection method.
- **Security of application programming interfaces.** There are five generations of programming languages. The higher the generation, the more abstract the language is and the less a developer needs to know about the details of the operating system or hardware behind the code. The five generations are:
 - 1: Machine language.** This is the binary representation that is understood and used by the computer processor.
 - 2: Assembly language.** Assembly is a symbolic representation of the machine-level instructions. Mnemonics represent the binary code, and commands such as ADD, PUSH and POP are used. The assemblers translate the code into machine language.
 - 3: High-level language.** High-level languages introduce the ability to use IF, THEN and ELSE statements as part of the code logic. The low-level system architecture is handled by the programming language. FORTRAN and COLBOL are examples of generation 3 programming languages.

4: Very high-level language. Generation 4 languages further reduce the amount of code that is required, so programmers can focus on algorithms. Python, C++, C# and Java are examples of generation 4 programming languages.

5: Natural language. Generation 5 languages enable a system to learn and change on its own, as with artificial intelligence. Instead of developing code with a specific purpose or goal, programmers only define the constraints and goal; the application then solves the problem on its own based on this information. Prolog and Mercury are examples of generation 5 programming languages.

8.3 Assess the effectiveness of software security

Putting protections in place is not enough security to give you peace of mind. To know that those protections are working as designed, organizations should routinely audit their access protections. You should also revisit your implementations to identify new risks that might need to be mitigated, and to ensure that the project is meeting the requirements that were agreed upon.

- **Auditing and logging of changes.** The processes and procedures for change control should be evaluated during an audit. Changes that are introduced in the middle of the development phase can cause problems that might not yet be discovered or caused in testing. The effectiveness of the change control methods should be an aspect of auditing the development phase.
- **Risk analysis and mitigation.** Most of the development methodologies discussed in section 8.1 include a process to perform a risk analysis of the current development cycle. When a risk has been identified, a mitigation strategy should be created to avoid that risk. Additionally, you can document why a risk might be ignored or not addressed during a certain phase of the development process.

8.4 Assess security impact of acquired software

When an organization merges with or purchases another organization, the acquired source code, repository access and design, and intellectual property should be analyzed and reviewed. The phases of the development cycle should also be reviewed. You should try to identify any new risks that have appeared by acquiring the new software development process.

8.5 Define and apply secure coding guidelines and standards

Many organizations have a security strategy that is focused at the infrastructure level; it deals with things like network security, identity and access management, and endpoint security. Organizations that develop code internally should also include coding in their security strategy. They should specify practices for ensuring the security of the code as well as coding standards, such as the preferred programming language for particular use cases.

- **Security weaknesses and vulnerabilities at the source-code level.** Just about every application (or chunk of source code) has bugs. While not all bugs are specifically related to security, they can sometimes lead to a security vulnerability. One effective way of finding and fixing bugs is to use source code analysis tools, which are also called static application security testing (SAST) tools. These tools are most effective during the software development

process, since it's more difficult to rework code after it is in production. However, be aware that these tools can't find many weaknesses and they introduce extra work for the teams, especially if they generate a lot of false positives. Today, with security being of paramount concern, the expectation is that all source code is scanned during development and after release into production.

- **Security of application programming interfaces.** Application programming interfaces (APIs) enable applications to make calls to other applications. Without proper security, APIs are a perfect way for malicious individuals to compromise your environment or application. The security of APIs starts with requiring authentication using a method such as OAuth or API keys. Authorization should also be used and enforced. For example, one API key might enable you to read information but you need a separate API key to alter or write information. Many companies use an API security gateway to centralize API calls and perform checks on the calls (checking tokens, parameters, messages, etc.) to ensure they meet the organization's requirements. Other common methods to secure your APIs is to use throttling (which protects against DoS or similar misuse), scan your APIs for weaknesses, and use encryption (such as with an API gateway).
- **Secure coding practices.** There are established practices you should follow to maximize the security of your code. Some of the most common ones are:
 - **Input validation.** Validate input, especially from untrusted sources, and reject invalid input.
 - **Don't ignore compiler warnings.** When compiling code, use the highest warning level available and address all warnings that are generated.
 - **Deny by default.** By default, everybody should be denied access. Grant access as needed.
 - **Authentication and password management.** Require authentication for everything that is not meant to be available to the public. Hash passwords and salt the hashes.
 - **Access control.** Restrict access using the principle of least privilege, and deny access if there are issues checking access control systems.
 - **Cryptographic practices.** Protect secrets and master keys by establishing and enforcing cryptographic standards for your organization.
 - **Error handling and logging.** Avoid exposing sensitive information in log files or error messages. Restrict access to logs.
 - **Data protection.** Encrypt sensitive information, everywhere.
 - **Communication security.** Use Transport Layer Security (TLS) everywhere possible.
 - **System configuration.** Lock down servers and devices. Keep software versions up to date with fast turnaround for security fixes. You can find good information for securing your servers and devices from NIST. Visit <https://www.nist.gov> to search for standards and guides related to your environment.
 - **Memory management.** Use input and output control, especially for untrusted data, and watch for buffer size issues (use static buffers). Free memory when it is no longer required.

Domain 8 Review Questions

Read and answer the following questions. If you do not get one at least one of them correct, spend more time with the subject.

1. You are a software development manager starting a new development project. You want to focus the development process around user stories. The development process must be efficient and have multiple iterations as changes and requirements are discovered. Which development methodology should you use?
 - a. Agile
 - b. Waterfall
 - c. Spiral
 - d. Rapid application development
2. You are in the early stages of the development lifecycle and creating design requirements. The application will contain several forms that allow users to enter information to be saved in a database. The forms should require users to submit up to 200 alphanumeric characters, but should prevent certain strings. What should you perform on the text fields?
 - a. Input validation
 - b. Unit testing
 - c. Prototyping
 - d. Buffer regression
3. You plan on creating an artificial intelligence application that is based on constraints and an end goal. What generation language should you use for the development process?
 - a. Generation 2
 - b. Generation 3
 - c. Generation 4
 - d. Generation 5

Answers to Domain 8 Review Questions

1. Answer: A

Explanation: Agile development emphasizes efficiency and iterations during the development process. Agile focuses on user stories to work through the development process.

2. Answer: A

Explanation: The text fields that the users interact with should have input validation to ensure that the character limit has not been exceeded and that no special characters that might cause database inconsistencies are used.

3. Answer: D

Explanation: Generation 5 languages are associated with artificial intelligence. The constraints of the application and its goal are defined; then the program learns more on its own to achieve the goal.

Useful References

Supplement your exam preparation or polish the attained certification by reviewing relevant resources.

- [Top Webcasts On IT Security And Compliance To Attend Or Watch Recorded While Earning CPE Credits](#)
- [How to Build an Effective Data Classification Policy for Better Information Security](#)
- [Five Reasons to Ditch Manual Data Classification Methods](#)
- [10 Ways to Improve Automated Data Classification](#)
- [SysAdmin Magazine: “My Precious!” — Finding & Securing Sensitive Data](#)
- [How to Identify and Prioritize Information Security Risks](#)
- [How to Perform IT Risk Assessment](#)
- [SysAdmin Magazine: Assess Your Risks or Die Tryin’](#)
- [Insider Threat Playbook: How to Deter Data Theft by Departing Employees](#)
- [Five Ways to Reduce the Risk of Employee Theft of Confidential Information](#)
- [Best Practices: How to Off-Board an Employee for Good](#)
- [Best Practices: How to Minimize the Risk of Insider Threats](#)
- [Insider Threat Detection: 10 Techniques for Top-to-Bottom Defense](#)
- [Cyber Chief Magazine: Insider Threat — Uncover and Neutralize](#)
- [Top 20 Critical Security Controls for Effective Cyber Defense](#)
- [White Paper: Compliance Demystified](#)
- [GDPR Rules of Data Breach Notification: How to Report Personal Data Loss and Avoid Fines?](#)
- [SysAdmin Magazine: Why GDPR Won’t Kill You](#)
- [Brace Yourself — HIPAA Security Risk Assessment Is at Your Door](#)
- [White Paper: Seven Questions for Assessing Data Security in the Enterprise](#)
- [Windows Server Hardening Checklist](#)
- [Four Things You Should Never Store in Your Database](#)
- [What to Know about the Threat of Privileged Users](#)
- [Key Lessons Learned from Data Breaches Caused by Privilege Abuse](#)
- [Best Practices: How to Harden Privileged Account Security](#)
- [How to Collect Server Inventory](#)
- [Best Practices: How to Implement Audit Policy](#)

About the Author



Brian Svidergol is focused on infrastructure and cloud-based solutions with a special emphasis on identity and access management. Brian has worked with numerous companies, from startups to Fortune 500 companies, helping them design and implement secure infrastructure and cloud solutions. He holds numerous industry certifications from Microsoft, Red Hat, VMware, Novell and NetApp. Brian has authored six books published through O'Reilly, Sybex, and Microsoft Press and self-published one book.

About Netwrix

Netwrix Corporation is a software company focused exclusively on providing IT security and operations teams with pervasive visibility into user behavior, system configurations and data sensitivity across hybrid IT infrastructures to protect data regardless of its location. Over 9,000 organizations worldwide rely on Netwrix to detect and proactively mitigate data security threats, pass compliance audits with less effort and expense, and increase the productivity of their IT teams. Founded in 2006, Netwrix has earned more than 140 industry awards and been named to both the Inc. 5000 and Deloitte Technology Fast 500 lists of the fastest growing companies in the U.S.

For more information about Netwrix, visit www.netwrix.com.

Netwrix Auditor is a visibility platform for user behavior analysis and risk mitigation that enables control over changes, configurations and access in hybrid IT environments to protect data regardless of its location. The platform delivers security intelligence to identify security holes, detect anomalies in user behavior and investigate threat patterns in time to prevent real damage.

Netwrix Auditor includes applications for Active Directory, Azure AD, Exchange, Office 365, Windows file servers, EMC storage devices, NetApp filer appliances, SharePoint, Oracle Database, SQL Server, VMware and Windows Server. Empowered with a RESTful API and user activity video recording, the platform delivers visibility and control across all on-premises or cloud-based IT systems in a unified way.

Learn more about Netwrix Auditor and download its free 20-day trial by visiting netwrix.com/go/freetrial.

Corporate Headquarters:

300 Spectrum Center Drive, Suite 200, Irvine, CA 92618

Phone: 1-949-407-5125 **Toll-free:** 888-638-9749 **EMEA:** +44 (0) 203-588-3023



netwrix.com/social