Chauster
UpSkilling Solutions

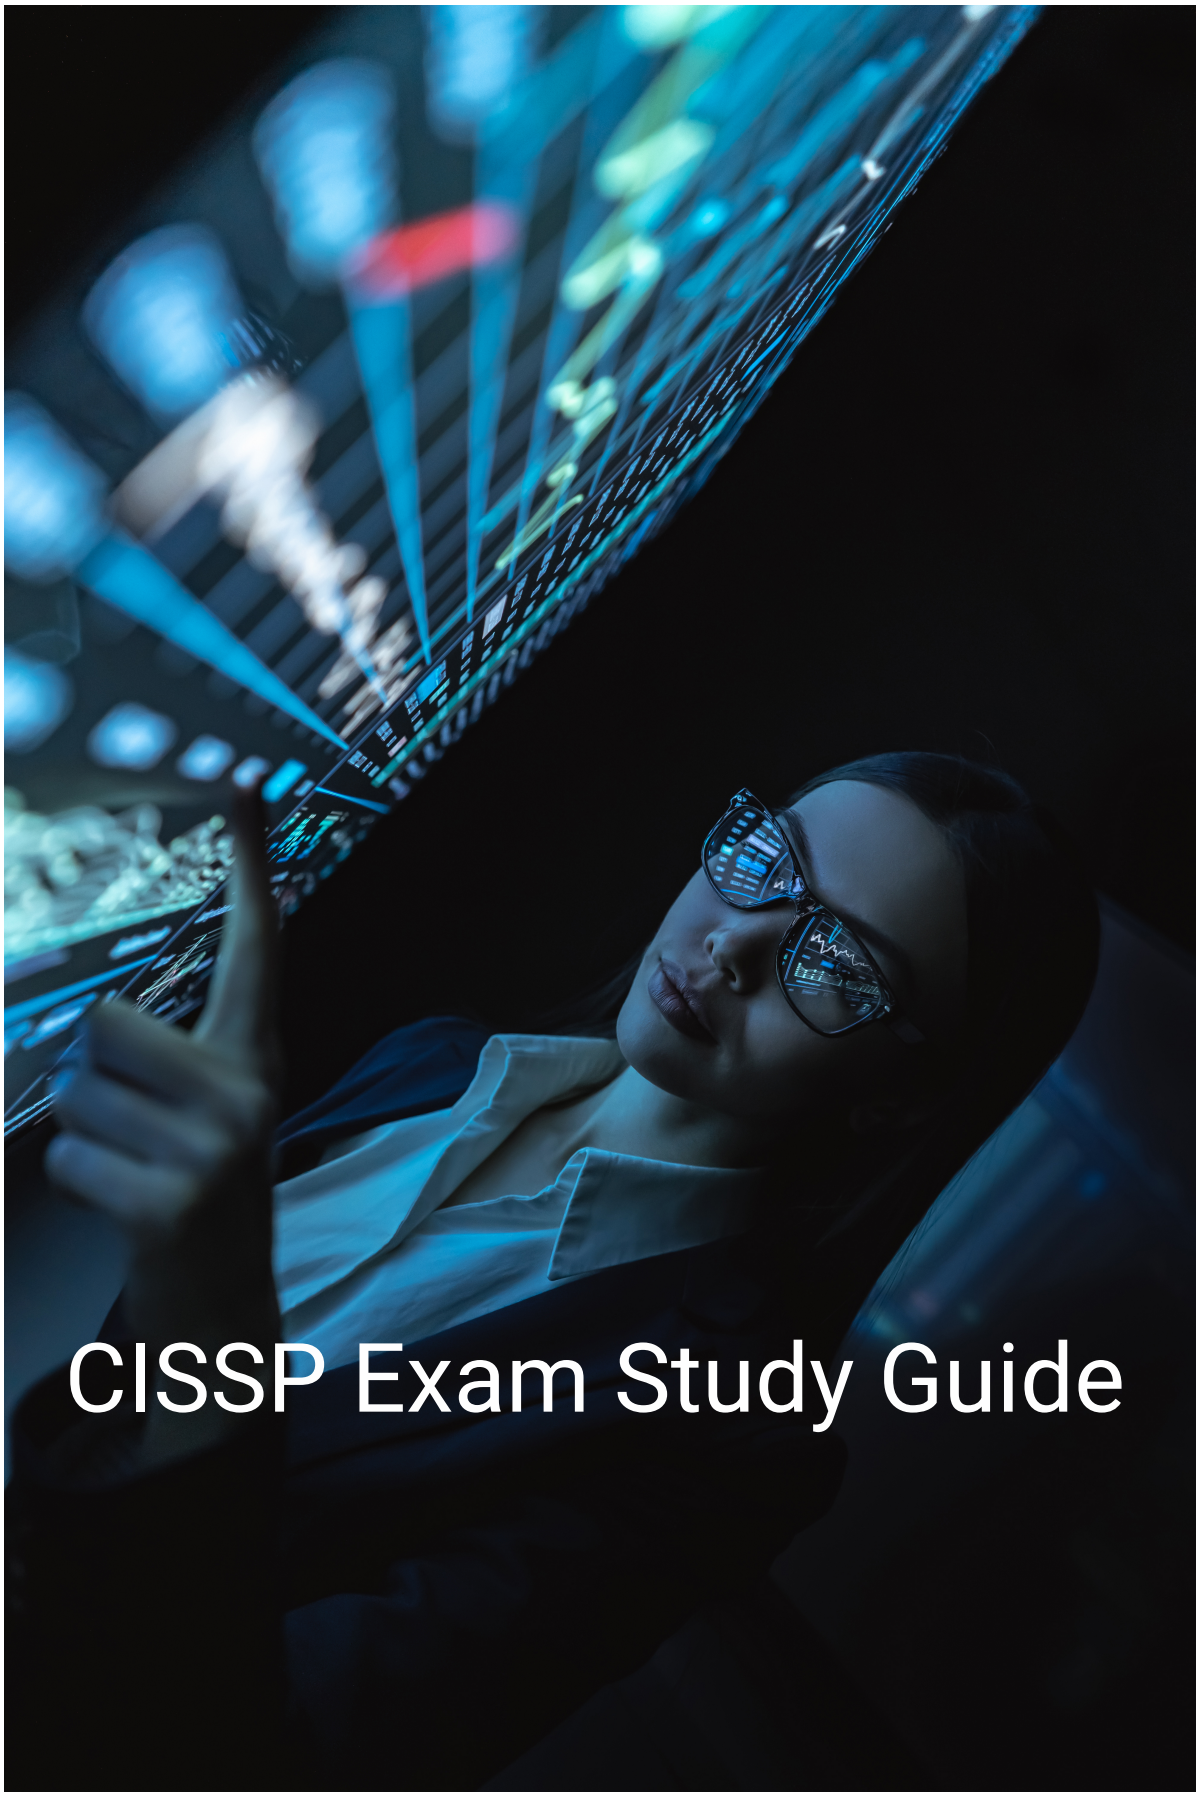

CISSP Exam Study Guide

Table of contents

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. Communications and Network Security
5. Identity and Access Management
6. Security and Assessment 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 itself is long, especially compared with other industry certifications. You can take it in English or another language.The English language exam is now a computerized adaptive testing (CAT) exam, so it changes based on your answers. You get up to 3 hours to complete a maximum of 150 questions (and a minimum of 100 questions). Its available in Chinese, English, German, Japanese, and Spanish.

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. Combine the mediums you use. Reading something, hearing something and doing something helps your brain process and retain information. If your plan is to just 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 or 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 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, 2024, 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/certifications/cissp/cissp-certification-exam-outline.

While most of the exam topics remain the same, there are some minor changes to reflect the latest industry trends and information. 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.

Domain 1. Security and Risk Management

1.1 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 over the code of ethics available at www.isc2.org/Ethics. At a minimum, understand the four canons:

  - **Protect society, the common good, necessary public trust and confidence, and the infrastructure.** Put the common good ahead of yourself. Ensure that the public can have faith in IT infrastructure and security.
  - **Act honorably, honestly, justly, responsibly, and legally.** Always follow the laws. If you find yourself working on a project with conflicting laws from different countries or jurisdictions, prioritize the local jurisdiction from which you are performing the services.
  - **Provide diligent and competent service to principals.** Never pass yourself off 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. If you follow the first three canons, you automatically comply with this one.

- **Organizational code of ethics.** You must also support ethics at your organization. Examples include evangelizing ethics throughout the organization, providing documentation and training around ethics, and looking for ways to enhance the existing organizational ethics. Ethical codes differ, so be sure to familiarize yourself with your organization's guidelines.

1.2 Understand and apply security concepts

Be sure to understand the following security concepts. Note that the first three — confidentiality, integrity and availability — make up what's known as the "CIA triad." Together, confidentiality, integrity, availability, authenticity, and nonrepudiation make up what's known as the "5 Pillars of Information Security". The 2024 exam update added "5 Pillars of Information Security" to the title of the topic but the topic content remains the same.

- **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 protection. 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 ensure 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 that information. At a high level, you use access controls — locked doors, folder permissions and multifactor 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.** Integrity is ensuring that data isn't changed improperly. Some of the same strategies that protect confidentiality also help maintain integrity. 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 it a numeric value, which is calculated at the source before the transfer and then again by the recipient after the transfer; a match ensures 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.

- **Authenticity.** Authenticity involves ensuring a transmission, message or sender is legitimate. See the NIST glossary for examples: https://csrc.nist.gov/glossary/term/authenticity.

- **Nonrepudiation.** If I send you a digitally signed email message and you receive it and read it, then I cannot deny that I sent the message (the premise being that I used my private key to sign the message and you can be assured of that by using my public key to validate the signature). If I used some type of proof of delivery (such as a registered email service), then you cannot deny that you received the message. Read receipts and tracking pixels are generally not recognized as valid ways of proving something was received and read. While this example focuses on email, nonrepudiation applies to other mediums as well.

1.3 Evaluate and apply security governance principles

Adopt a framework, such as the National Institute of Standards and Technology (NIST) framework, to establish security governance principles. New for the 2024 exam update is the addition of the term "sustain" which alludes to keeping your governance principles fresh as changes are seen in technology. When adopting a framework, be sure the framework includes the following:

- **Alignment of security function to strategy, goals, missions and objectives.** 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. All of these — strategy, goals, missions, and objectives — flow down, with each one supporting the others. Objectives are the closest to the ground and represent small efforts to help you achieve a mission. Missions represent a collection of objectives, and one or more missions leads to goals. When you reach your goals, you are achieving the strategy! A business has a mission and uses objectives to try to meet the 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 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 closely aligned with the mission and objectives. Your organization's strategy, goals, missions, and objectives will periodically change. As such, you need to ensure that you review the changes on a periodic basis (example, once per year) and align your efforts to the latest updates (as you and your team might not participate in the strategy setting).

- **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 key) and divestitures (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. When it comes to governance committees, ensure your organization has a rotation program to have new representation on governance committees. When it comes to risks with organization processes, ensure that you are reevaluating the processes at least yearly. For example, many organizations didn't account for artificial intelligence for acquisitions a few years ago. Today, accounting for artificial intelligence is an important requirement.

- **Security control frameworks (e.g., International Organization for Standardization (ISO), National Institute of Standards and Technology (NIST), Control Objectives for Information and Related Technology (COBIT), Sherwood Applied Business Security Architecture (SABSA), Payment Card Industry (PCI), Federal Risk and Authorization Management Program (FedRAMP)).** 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.). For the 2024 exam update, specific frameworks have been added to the title. All of 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. 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 re-certify on it regularly (typically yearly).

When it comes to sustaining your security governance principles, you should consider the following:

1. Perform periodic risk assessments. Are there new risks due to the technology landscape changing? Update your governance principles based on the periodic assessments. Risk assessments can be performed by an internal team or an external team. It is a good practice to use both, especially at large organizations.
2. Performing periodic governance audits (internal and external). Any new findings that point to requiring updates to your security governance principles?
3. Monitor your adopted framework (such as NIST) for changes and updates. As changes come in, you will need to update your security governance principles.
4. Update your security awareness training to include updates to your security governance principles.

1.4 Understand legal, regulatory, and compliance issues that pertain to information security in a holistic context

Alright, let's talk about privacy — basically, how people and companies should handle your personal info, like your name, address, phone number, or even what websites you visit. There are rules in different places around the world to make sure this stuff stays safe. For example, in Europe, there's a big rule called GDPR (General Data Protection Regulation). It says companies have to be super careful with your data, tell you what they're doing with it, and even delete it if you ask. Over in California, there's a law called CCPA (California Consumer Privacy Act), which gives people the right to see what info companies have on them and say, "Hey, don't sell my stuff!"

Other countries have their own rules, too. In China, they have the Personal Information Protection Law (PIPL), which works a lot like GDPR, making sure companies ask for permission before using your personal data. In South Africa, there's something called POPIA (Protection of Personal Information Act), which also protects people's info from being misused. All of these laws are about the same big idea: your personal info belongs to you, not companies, and they should treat it with care. So, when you hear about "privacy issues," it's really about making sure people's personal data isn't being shared or used in ways they don't want.

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. For the 2024 exam update, compliance was added to the section. For all of the areas, the compliance part signifies monitoring and audits to ensure compliance with industry standards and frameworks. Often, external audits are important to independently evaluate compliance. Finally, don't forget about documentation which is a key part of compliance. Your organization needs policies, procedures, evidence (log data, as one example), and business continuity / disaster recovery plans.

- **Cybercrimes 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 personal health information was not protected as required by the Health Insurance Portability and Accountability Act (HIPAA) law. 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:

  - **Trademark** —A logo, symbol or mascot used for marketing a brand
  - **Patent** — 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 which prevents unauthorized duplication, distribution or modification
  - **Licensing** — A contract between the software producer and the consumer which 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 is subject 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.

- **Privacy.** Many laws include privacy protections for personal data. The EU's General Data Protection Regulation (GDPR) has strong privacy rules that apply to any organization anywhere that stores or processes the personal data of EU residents; these 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, set limits to personal data collection, protect personal data with reasonable security and more. 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 Privacy Shield, personal data flow from the EU to the United States would not be allowed.

- **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. 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 (which is 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.

1.5 Understand requirements for investigation types (i.e., administrative, criminal, civil, regulatory, industry standards)

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.** An administrative investigation has a primary purpose of providing the appropriate authorities with incident information. Thereafter, the authorities will determine the proper action, if any, 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 the court case. 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 is typically about monetary damages, not incarceration or a criminal record. In a civil case, a preponderance of evidence is required to secure a victory. This differs from criminal cases, where a suspect is innocent until proven guilty beyond a reasonable doubt.

- **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.

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

Develop clear security policy documentation, including the following:

- **Policies.** These are high-level documents, often written by the management team. Policies are mandatory. They are purposely vague. For example, a policy might require you to ensure the confidentiality of company data but not document 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 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. For example, a Group Policy object (GPO) in a Windows network is sometimes used to comply with standards. Configuration management solutions can also help you establish baselines and spot configurations that drift away from them.

1.7 Identify, analyze, assess, prioritize, and implement 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. 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). New for the 2024 exam update is the addition of "assess". This is an entirely new facet. The existing content remains the same but you now have to understand how you assess how an organization is handling business continuity.

- **Develop and document the scope and the 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 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 business continuity plan, disaster recovery 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).Do not forget to look at external dependencies, like vendors of digital assets – the contracts you have with them - or any regulatory framework that influences your organization. When all of these are already in place, you assess them on a periodic basis to ensure that the organization is maintaining the documentation, that the content remains accurate, and that the organization is periodically reviewing and updating all of the business continuity items.

- **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). Last, you need 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. From an assessment perspective, you will review the BIA to ensure that critical business functions and their supporting systems are identified. You will compare the RTO and RPO to the organization's capabilities and systems (are they able to meet the RTO and RPO?).

- **Test and validate.** As part of assessing an organization's business continuity stance, one of the most important tasks is testing. Without testing, you cannot be certain that the business continuity plans are ready for use in a real-world scenario.

  - **Recovery tests.** These can be data center failovers, failover of cloud regions, failover of highly available components (hardware or similar) and apps.
  - **Tabletop exercise.** These are virtual exercises where teams are presented with one or more scenarios and the participating teams walk through each step of the business continuity and disaster recovery paths and options. The goal of the exercise is to test the BC/DR plans and identify any gaps along the way. Common scenarios in a tabletop exercise include walking through a cyberattack, a natural disaster, or a cloud provider failure.
  - **Update and retest.** After recovery testing and tabletop exercises, teams should update their plans, remediate any gaps, then re-test. Many organizations opt to test yearly to ensure readiness.

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 are routinely targeting 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 employment candidates thoroughly is a key 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, all references should be contacted.

- **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 example, 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, transfers, and termination processes.** Onboarding is the processes tied to a new employee starting at your organization. Having documented processes in place enables new employees to be integrated

as quickly and as consistently as possible, while reducing risk. For example, if you have 5 IT admins that work in the onboarding processes, you might get different results if you don't have the processes standardized and documented. A new hire might end up with more access than required for their job.

When somebody moves from one job to a different job in the same organization, it is called a transfer. For example, when a person in the Human Resources (HR) department moves to the Facilities team, it is a transfer. Transfers are important because a person transferring will likely need to have their access adjusted (access for previous job role removed and access required for new job role added).

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 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 organizational 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, they sit in a designated area, or their display names in the directory call out their status.

The following topics were removed from the exam blueprint for the 2024 exam update but remain in the study guide as supplemental information:

- **Compliance policy requirements.** Organizations have various compliance mandates to adhere to, based on their industry, country and operating methodologies. Common to all such companies is the need to maintain documentation about the company's compliance. Employees should be required to understand the company's high-level compliance mandates 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 which outlines the type of data covered by the policy and who the policy applies to. Employees and contractors 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.

- **Threat and vulnerability identification.** 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.

  - **Risk analysis, assessment, and scope.** You have a risk when you have a threat and a vulnerability. Next, you need to figure out the chances of it happening and the consequences if it does happen. In 2024, "scope" was added. To properly scope an assessment (or any project), you need to identify the requirements and the stakeholders. You need to engage with the stakeholders to gather their individual needs. Scoping is important for ensuring proper completion, scheduling, and proper costs. Be familiar with the assessment/ analysis 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 high.
    - **Quantitative.** This method is more objective than the qualitative method; it uses dollars or other metrics to calculate 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.

- **Risk response and treatment (e.g. cybersecurity insurance).** Next you 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.** Note that for the 2024 version of the exam, this specific topic was removed. However, it might provide with you additional context so it remains in this study guide. When you have risk, you can use a software or hardware solution to reduce that risk. A countermeasure, sometimes referred to as a "control" or a "safeguard," can help reduce the risk. 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). You can implement any of several different countermeasures to reduce the likelihood of that password being compromised: 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 put the concepts into your environment. This is not a step-by-step technical configuration, but the process of the implementation — where you start, in which order it occurs and how you finish.

- **Applicable types of controls (e.g., preventive, detectiion, corrective).** 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 computer.
  - 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 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.

- **Control assessments (e.g. security and privacy).** You need to periodically assess your security and privacy controls. What's working? What isn't working? As part of this assessment, the existing documents 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. Often, security and privacy control assessment are performed and/or validated by different teams, with the privacy team handling the privacy aspects.
- **Continuous 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. Note that for 2024, the "Continuous" wording was added. This indicates an ongoing need to monitor, measure, and adjust as technologies change, new systems are introduced, and software/hardware capabilities change.

- **Reporting (e.g., internal, external).** 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 (e.g., risk maturity modeling).** Continuous improvement is an 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). There are defined risk maturity models that use categories or levels. For example, imagine a 5-level risk maturity model, with 1 being minimal maturity and 5 being the highest level of maturity. Continuous improvement with risk maturity modeling is really about an organization advancing from a lower level (such as 3) to a higher level (such as 4). Which additional improvements need to be made to move up a level? That's the start of continuous improvement related to risk maturity modeling.

- **Risk frameworks (e.g., International Organization for Standardization (ISO), National Institute of Standards and Technology (NIST), Control Objectives for Information and Related Technology (COBIT), Sherwood Applied Business Security Architecture (SABSA), Payment Card Industry (PCI).** A risk framework documents how your organization handles risk assessment, risk resolution and ongoing monitoring. For an example of a risk framework, see http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-37r1.pdf. There are other risk frameworks, such as the British Standard BS 31100. Be familiar with 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 them (often by putting yourself in an attacker's shoes or mindset). There are 4 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.
- VAST (visual, agile and simple threat modeling) applies across IT infrastructure and software development without requiring security experts.
- DREAD (damage, reproducibility, exploitability, affected users, discoverability) provides a flexible rating approach to gauge an identified threat along five main questions.

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 know now this person is a threat and can focus on what they might want to achieve. However, there are many times when an organization is 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 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.

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 supply chain risk management (SCRM) concepts

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 it
  - 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 based on the company's 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 and does that introduce any conflicts of interest?

- **Risk mitigations (e.g., 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 one company is the primary supplier of a critical hardware component, then an 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

  Note that for the 2024 version of the exam, the following 2 topics were removed. However, they might provide with you with additional context and information useful for the exam.

- **Minimum security requirements.** As part of 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 should be reviewed. Your company sometimes has an SLA standard that should be applied, when possible, to the service level agreements 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 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 increase awareness and training (e.g., social engineering, phishing, security champions, gamification).** 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 and free USB sticks), 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 phishing test campaigns, tailgating at company doors and social engineering tests). The following techniques and methods are important to know for the exam:

  - **Social engineering.** While many organizations don't perform social engineering campaigns (testing employees using benign social engineering attempts) as part of security awareness, it is likely to gain traction. Outside of campaigns, presenting social engineering scenarios and information is a common way to educate.
  - **Phishing.** Phishing campaigns are very popular. Many organizations use third-party services to routinely test their employees with fake phishing emails. Such campaigns produce valuable data, such as the percentage of employees who open the phishing email, the percentage who open attachments or click links, and the percentage who report the fake phishing email as malicious.
  - **Security champions.** The term "champion" has been gaining ground. Organizations often use it to designate a person on a team who is a subject matter expert in a particular area or responsible for a specific area. For

example, somebody on your team could be a monitoring champion — they have deep knowledge around monitoring and evangelize the benefits of monitoring to the team or other teams. A security champion is a person responsible for evangelizing security, helping bring security to areas that require attention, and helping the team enhance their skills.

- **Gamification.** Legacy training and education are typically based on reading and then answering multiple-choice questions to prove one's knowledge. Gamification aims to make training and education more fun and engaging by packing educational material into a game. That might mean playing an actual game, but it might also mean keeping track of scores, having leader boards, and enabling people to earn something based on their scores or progress (kudos, special avatars or similar). Gamification has enabled organizations to get more out of the typical employee training.

New for 2024 is changing of "present" to "increase". This is a subtle switch from the view that you are putting in all of this stuff in a new environment vs. taking an existing environment (where users have had some training and awareness) and increasing awareness. It also plays into a need that training and awareness isn't a one-time thing, but an ongoing challenge that teams work on regularly (such as quarterly or semi-annually).

- **Periodic content reviews to include emerging technologies and trends (e.g., cryptocurrency, artificial intelligence (AI), blockchain).** Threats are complex, so training needs to be relevant and interesting to be effective. This means updating training materials 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, it isn't effective. The same applies to other materials. Instead of relying on long and detailed security documentation for training and awareness, consider using internal social media tools, videos and interactive campaigns. New for 2024 is the addition of emerging technologies and trends and the reference to crypto, AI, and blockchain. As with many areas within information security, content needs to be updated routinely. For example, phishing emails historically contained a malicious link or a malicious attachment. Users were heavily trained on being wary of links and attachments. Attackers began to use QR codes. Organizations that didn't periodically review and update their training material to include QR codes were at greater risk vs. organizations using periodic content reviews.

- **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 who click on a fake phishing campaign email link. Is the awareness and training bringing that number clicks down over time? If so, the program is effective. If not, you need to re-evaluate.

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. Using that data, you need to figure out the likelihood of additional customers experiencing a security incident. 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)

Domain 1. Answers to 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, the focus is on the data loss, so 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 like intellectual property and data. Domain 3 includes some physical security topics. Note that for 2024, this domain remains unchanged (titles remain the same, content remains the same, nothing added or removed).

2.1 Identify and classify information and assets

To improve security, you need to identify both your data and your physical assets, and then 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 and 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 (DLP) 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. Other organizations use it for special scenarios, such as during a merger or acquisition. When studying for this section, concentrate on knowing the following items:

- **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 or asset owner, who should be provided with details about the access being requested. Before a user is granted access to the data or asset, 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 Establish information and asset handling requirements

This section covers how people and systems work with data and assets. This includes any action you can take with the data, such as read, copy, edit or delete. It also includes establishing asset handling requirements. The following 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. What else require markings and labels? Assets. For example, laptops should be identifiable as the organization's devices. Many organizations opt to use bar codes and scan them in for service, which eases the burden of handling physical assets.

- **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 that 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. You also need a way to store your

assets. Servers, spare computers and other assets need to be securely stored to minimize the risk of data loss or misuse of assets.

- **Destruction.** Your organization should have a list of requirements and 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 only require 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 was classified as the most sensitive data at your organization. When it comes to your assets such as computers, you also need a destruction method. Some organizations opt to remove all storage devices and recycle the rest of the hardware through third-party recycling services. The storage devices are destroyed using the organization's method, such as physical destruction (shredding or similar).

2.3 Provision information and assets securely

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:

- **Information and asset ownership.** 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). If you don't have an assigned owner for data, how can you go through a formal access approval? You can't, at least not as effectively. The same applies to assets. How can you properly account for assets if you don't know which department owns them? And how can you assign the right asset type for high-risk travel if you don't have assets classified? The data owner is also responsible for classification labels. In larger companies, an asset management department handles asset classification.

- **Asset inventory (tangible, intangible).** This topic covers asset inventory, which is the act of keeping track of your assets, such as computers or software. Without a proper inventory, it is difficult to ascertain the health and whereabouts of your assets. Some assets, such as computing devices, might provide access to your data, so there are security implications tied to asset inventory.

  - **Tangible.** A tangible asset is something you can hold, such as a physical server.
  - **Intangible.** An intangible asset is a non-physical asset, such as an application, trademark, domain name or patent.

- **Asset management.** Managing your assets, whether tangible or intangible, is a day-to-day task that you should take seriously. Many experts agree that having a solid asset management program is a key part of keeping your environment secure. Some of the key aspects of asset management include:

- **Firmware and BIOS updates.** For each physical computing device, you need to keep the firmware and the BIOS up to date on security fixes.

- **Software updates and patches.** While software updates and patches sometimes offer new features, usually they provide security enhancements and bug fixes. Experts often claim that installing security updates and patches is the single biggest thing an organization can do to enhance their overall security.

- **Unauthorized software removal.** Even if you have an excellent asset inventory and keep your firmware, BIOS and software up to date, you might still have a serious issue with unauthorized software. It is common for users and administrators to install software that your organization hasn't approved or vetted. Many software vendors and independent developers are releasing applications that don't require a traditional install, which enables users and administrators to bypass software installation restrictions commonly used in enterprise networks. You need an automated solution to identify and remove unauthorized software.

- **Configuration management.** To reduce support costs and maximize security, you need a solution to enforce your organization's desired configuration on computing devices. Configuration items such as the default browser, browser security settings and data encryption settings provide important endpoint security. While an asset management team might not be responsible for directly dictating a configuration, they often help enforce configurations with their software management and inventory solutions.

2.4 Manage data lifecycle

Often, administrators focus on retaining data and protecting data, and neglect removing data, but all parts of the data lifecycle are important.

- **Data roles (i.e. owners, controllers, custodians, processors, users/subjects).** You should be able to identify the roles associated with data, especially the following roles called out as examples:

  - **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). They also own classification labeling.

  - **Controllers.** A data controller is the person ultimately responsible for the data and decides how it will be used, processed and protected. Some countries or regulations define controllers and specific controller tasks.

  - **Custodians.** A custodian is a hands-on role that implements and operates solutions for data (e.g., backups and restores). While a system owner is responsible for the computer environment (hardware, software) that houses data, this is typically a management role; operational tasks are handed off to the custodian.

- **Processors.** Data processors are the users who read and edit data regularly. They 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?

- **Users/subjects.** An individual who can be identified in data. This could be because the data contains a name, ID number or other identifying information. Some countries regulate the rights of data subjects with regard to data identifying them.

- **Data collection.** Many organizations collect data. Sometimes, data is required for transactions (such as buying something). Other times, data is required to improve the user experience (for example, by knowing that a user is viewing the website on a smartphone). Countries are enacting laws around data collection. For example, the EU's General Data Protection Regulation (GDPR) calls out requirements around users knowing, understanding and consenting to data collected about them. Here are some key practices around data collection:

  - **Document what you collect.** Many organizations have been caught off guard when they discover exactly what they are collecting and storing. Mobile apps often collect enormous amounts of data, and it can be difficult to track down exactly what's being collected. Organizations are responsible for securely storing collected data.
  - **Tell people what you collect and how to manage it.** It is a good practice to tell users what you collect and how you collect it. It is also a good practice to ask users for approval to collect and use data. It is difficult for users to give consent when organizations don't tell them the details.
  - **Limit what you collect.** 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 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.

- **Data location.** If you go back a few years, the location of data wasn't a big talking point. Many administrators rarely thought about it beyond the disaster recovery or business continuity aspects. But laws are quickly changing. Some countries have laws that require certain types of data to stay within the country. Some organizations have enacted policies and standards covering the location of data. Many cloud providers are offering flexible data location options, too.

- **Data maintenance.** Data maintenance refers to managing the data as it makes its way through the data lifecycle (simplified: creation, usage, retirement). Data maintenance is the process (often automated) of making sure the data is available (or not available) based on where it is in the lifecycle. For example, new data that is widely used should be available on high performance storage. Old data that is rarely used should be available but without the need of high performance. Expired data or data no longer of use should be archived and/or permanently deleted, based on your organization's data retention policy.

- **Data retention.** Some data retention requirements specify the minimum length of time data is kept. For example, your organization might have a policy to maintain email data for 7 years. But there is another

aspect to retention — not keeping data longer than the policy mandates. You should ensure that your organization holds data for the required period and that it securely deletes data that it no longer requires to reduce the risk of its exposure.

- A company often has multiple reasons for keeping data, including regulatory or legal reasons and company policies. In many cases, a company must keep data for longer than the data provides value (to satisfy regulatory reasons, for example).
- If your company maintains data longer than required and that data is compromised in a breach, it can have a big impact on the company. As part of your comprehensive security policies, you should account for the destruction of unneeded data.

- **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.

- Data destruction. When you need to delete old or unused data, you need a secure method to destroy it so that it doesn't get into the wrong hands. 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.

2.5 Ensure appropriate asset retention (e.g., end of life (EOL), end of support (EOS))

This section covers hardware and personnel from an asset retention standpoint.

- **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 account for the hardware needed to get to the data that you are saving. This includes tape drives, various media readers, and any other hardware or software necessary.

- **Personnel.** 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.

The terms "end of life" and "end of support." often mean the same thing: The company is no longer actively working on the products, there won't be any further updates or bug fixes, and you can't call them for help if something goes wrong. Organizations should watch EOL and EOS dates carefully and plan to upgrade or replace the assets prior to the EOL or EOS date. Assets (hardware or software) that are past their EOL or EOS date could present security risks for your organization. And because the vendor is no longer supporting them, there often isn't an easy fix if a security issue is discovered.

2.6 Determine data security controls and compliance requirements

You need data security controls that protect data in each possible state: at rest, in transit or in use. Each state requires a different approach to security. There aren't as many security options for data in use as there are for data at rest or data in transit. Keeping the latest patches deployed to all computing devices, maintaining a standard computer build process, and running anti-virus and anti-malware software are typically the primary protections for data in use.

- **Data states (e.g., in use, in transit, at rest).** The industry identifies 3 data states:

  - **Data in use.** Data that is actively being worked on (for example, a spreadsheet that a person is currently editing)
  - **Data in transit.** Data that is moving from a source (such as a computer) to a destination (such as another computer)
  - **Data at rest.** Data stored on a storage medium (disk, tape, etc.)

- **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 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; then you would know that your selection process is working as expected.

- **Data protection methods (e.g., Digital Rights Management (DRM), Data Loss Prevention (DLP), Cloud Access Security Broker (CASB)).** Options for protecting your data vary depending on its state:

  - **For data at rest.** You can encrypt data at rest. You should consider encryption for operating system volumes and data volumes. 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. Digital rights management (DRM) is useful for data at rest. That's because DRM travels with the data and remains with the data regardless of the data state. DRM is especially useful when you can't encrypt data volumes.

    A cloud access security broker (CASB) solution often combines DLP, a web application firewall, some type of authentication and authorization, and a network firewall in a single solution. A CASB solution is helpful for protecting data in use (and data in transit).

- **For data in transit.** When data is in transit, it is often being transferred from one computer to another (for example, from a server to a client computer). Data in transit is not just data moving from your local area network to or through the internet; it is also data moving from anywhere to anywhere. You can use encryption for data in transit. For example, a web server uses a certificate to encrypt data being viewed by a user. You can use IPsec to encrypt communications, too. 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 loss prevention (DLP) solutions are useful for data in transit. DLP solution can scan data in transit (such as data leaving your organization network on the way to the internet) and stop the transfer of data based on a set of DLP rules. For example, if outbound data contains numbers in a structure that matches a U.S. Social Security number, and you have a rule to protect against that, such communication can be blocked. A CASB solution, as described above, is helpful for protecting data in transit.

- **For data in use.** Encryption is suited for data at rest and in motion, but there are limitations for data in use. Data in use is often in memory because it must be available to applications and/or operating systems. There are some third-party memory encryption solutions, but the selection is limited. Organizations often use strong authentication, monitoring, and logging to protect data in use. In the future, encryption of data in use will be as commonplace as it is for data at rest

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 each 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

Domain 2 Answers to 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 person ultimately 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 Research, implement, and manage engineering processes using secure design principles

When managing engineering processes from a security perspective, you need to use proven principles to ensure you end up with a secure solution that meets or exceed the security requirements. Research plays a big role in a couple of phases, such as the idea or concept phase and the design phase. Also new are all of the sub-topics for this section, all of which center around the security aspects of engineering.

- **Threat modeling.** Threat modeling is a process to identify, catalog and organize threats that pertain to your environment (often a single application or service at a time). Implementers can rely on the results of threat modeling to enhance security. There are different models that organizations use. One very popular model is the STRIDE model (spoofing, tampering, repudiation, information disclosure, denial of service, elevation of privilege). Others include PASTA, Trike, Dread and VAST. Don't spend time researching every available model and how they different from each other. Instead, understand the process of threat modeling, understand STRIDE, and know a few of the acronyms so that you can identify them as threat modeling methods.

- **Least privilege.** The principle of least privilege dictates that a person has the least amount of privilege to perform their job tasks. For example, an administrator who manages Windows servers doesn't require and shouldn't have administrative privileges to Linux hosts. In many cases, it is faster and easier to give administrators and developers more access than they need. To minimize risk, capture the job and task requirements and come up with the least amount of privilege required.

- **Defense in depth.** To adequately protect your environment, you need to have multiple layers of defense. That means combining multiple products and services together. Sometimes, these products and services work together and sometimes they work independently. Let's look at one real-world example: email. The first line of defense is protecting your domain names through private registration and locking changes. On top of that, you might incorporate Send Policy Framework (SPF) to limit who can send email from your domain.

Then, you will have an anti-spam/anti-virus scanner on the internet, prior to allowing email into your network. Thereafter, you will have host-based anti-spam/anti-virus on all computers. Finally, you might have email client security (attachment filtering, previews turned off, additional spam filtering, etc.). Working together, these comprise defense in depth. This concept applies across all aspects of your environment security.

- **Secure defaults.** When you think about defaults, think about how something operates brand new, just out of the box. For example, a wireless router comes with a default administrative password, and software that performs lookups in a database might, by default, use a legacy protocol or method. Securing default configurations is important not just for pre-packaged hardware and software; if you are developing a product or service, you should build in secure defaults. For example, if your service operates over the internet, you should use the latest version of TLS by default. Or, if you are creating software for a firewall, you can make the administrative portal, by default, available only on the LAN interface and not the WAN interface.

- **Fail securely.** When something fails, it must not reveal sensitive information. For example, if a web application has a failure, it shouldn't reveal the directory structure, lines of code and other internal IT information. Additionally, when something fails, it should not end up in a less secure state. In some cases, failing securely is reverting to defaults. And that's when secure defaults come into the mix.

- **Segregation of duties (SoD).** Segregation of duties is a process to minimize the opportunity to misuse data or damage your environment. It is widely used in organizations. For example, developers do not have access to the production environment to make code changes.

- **Keep it simple and small.** The term "keep it simple" refers to reducing complexity and keeping things easy to understand. Reducing complexity in your system, service or environment helps administrators keep it healthy, simplifies troubleshooting and makes it easier to secure. When designing a secure solution, you should try to keep it simple to take advantage of the benefits. New for 2024 is the addition of "and small" to the topic. The less "stuff" you have, the easier it is to keep it simple. For many organizations, reducing the total number of apps, reducing the total amount of vendors, and reducing overlap help to keep the environment smaller and simpler.

- **Zero Trust or trust but verify.** Zero Trust is a methodology that authenticates and authorizes all users and devices as though they are untrusted, and authentication and authorization continue to occur throughout the day. Historically, users with devices on a corporate network were automatically treated as trusted and secure. With Zero Trust, even users and devices on a corporate network go through the same validation as users and devices outside the network. Some organizations describe Zero Trust as a model that forgoes the traditional network edge (where the corporate network meets the internet). Zero Trust is a way to enable people to securely work from anywhere and from any device. The expression "Trust but verify" is an old Russian proverb that has made its way into information technology as well as many other areas. For example, let's look at email. An email sent by a user on your corporate network would normally be considered trusted (it comes from one of your users, on one of your devices, on your network). But, just in case, your anti-virus systems will scan it to verify it's not malicious before it goes out to the internet. New for 2024 is the addition of "or trust but verify" to the title (it previously was called out in a bullet in the previous study guide). Some consider "Zero Trust" as users and devices are never trusted and instead validated every time. While "trust but verify" is a

less strict way (and traditional way) to authenticate users and devices. For example, with "trust but verify", you might allow a user to connect to your organization's VPN (when you verify their identity) and thereafter, they are trusted. Zero Trust sometimes incorporates continuous authentication and/or continuous evaluation.

- **Privacy by design.** To maximize your success with privacy, you need to take a proactive approach and build privacy into your apps and services. Trying to add it later is difficult, and some privacy is likely lost beforehand. There are 7 recognized principles to achieve privacy by design:

  - **Proactive, preventative.** Think ahead and design for things that you anticipate might happen.
  - **Default setting.** Make private the default. For example, if you develop a social media app, don't share all of the user data with everybody by default.
  - **Embedded.** Build privacy in; don't add it later.
  - **Positive-sum.** Achieve security and privacy, not just one or the other. (And don't focus more on one than the other.)
  - **Full lifecycle protection.** Privacy should be achieved before, during and after a transaction. Part of achieving this is securely disposing of data when it is no longer needed.
  - **Visibility,** transparency, open. Publish the requirements and goals. Audit them and publish the findings.
  - **Respect, user-centric.** Involve end users, providing the right amount of information for them to make informed decisions about their data.

- **Shared responsibility.** In a shared responsibility model, multiple parties team up to meet a common goal. For example, a cloud service provider that offers virtual machines in the cloud is responsible for securing the virtual infrastructure, network and storage; their customers are responsible for securing the operating systems of the virtual machines.

- **Secure access service edge.** This is a new topic for the 2024 exam update. Secure access service edge, or SASE (people pronounce it "sassy"), is a new way to help keep computers and networks safe, especially when people are working from anywhere — not just in one office. It mixes two big ideas: network connectivity and security and how those combine to protect them). Instead of having all the protection in one spot (such as a firewall or a big wall around a castle, SASE spreads protection out so people can safely connect from anywhere — even from home or a coffee shop. It checks to make sure the person is allowed in, keeps an eye on their traffic, and blocks anything bad. It's like having a smart security guard in the cloud, always watching and ready to help.

3.2 Understand the fundamental concepts of security models (e.g., Biba, Star Model, Bell-LaPadula)

Security models enable people to access only the data classified for their clearance level. There are many models. Understand three specific security models: Biba, Star Model and Bell-LaPadula. All of them 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.

- **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.

- **Star Model.** This is not an official model, but its name refers to using asterisks (stars) to dictate whether a person at a specific level of confidentiality is allowed to write data to a lower level of confidentiality. It also determines whether a person can read or write to a higher or lower level of confidentiality. The principles of the stars are built into the Bell-LaPadula model (with the star property indicating that a subject can't write information down to a lower classification document) and the Biba model (with the star property indicating that a subject can't write information up to a higher classification document).

- **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, which are ways 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.

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 security capabilities of Information Systems (IS) (e.g., memory protection, Trusted Platform Module (TPM), encryption/decryption)

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/decryption.** 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. Whichever method is used for encryption, a corresponding decryption process is used on the receiving side.

- **Signing.** You can also sign communications, 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, etc.) 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 could 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. Like 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 avoid 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 themselves are not secret, but instead well-known. You can research their history, how they work and find extensive details about them.
  - **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.

- **Operational Technology / 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. New for 2024 is the addition of "Operational Technology" to the title. Thus, this topic expands from SCADA and similar systems (such as Distributed Control Systems or DCS and Programmable Logic Controllers or PCLs) to also include any systems that control physical devices and infrastructure. For example, operational technology includes building automation (heating and air, lighting, cameras) and transportation systems. It also includes emerging technologies such as robots and drones. Note that ICS is a subset of OT.

- **Cloud-based systems (e.g., software as a service (SaaS),** infrastructure as a service (IaaS), platform as a service (PaaS)). Unlike on-premises systems, 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 of firewalls and similar security solutions), encryption (use for all network communication and data at rest), and access control (use of centralized identity access and management system with

multifactor 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 updates. 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. Like cloud-based systems, review the IoT vendor to understand their history with reported vulnerabilities, response time to vulnerabilities and their overall approach to security. Not all IoT devices are suitable for enterprise networks!

- **Microservices (e.g., application programming interface (API)).** A microservice is a small service that often handles a single function. A single application could be made up of 25 microservices. For example, a website for shopping might have a service for ads, a service for search, a service for reviews and a bunch of other services. Securing microservices involves some of the same techniques as other areas: using HTTPS only, encrypting everything whenever possible and routine scanning. More closely aligned with microservices is the concept of shifting security left by integrating it into the CI/CD pipeline. Another area to help with security is looking at dependencies and dependencies of dependencies. For example, your service might depend on a third-party library. And that third-party library might use (depend on) another library. Understanding all of the dependencies will help ensure you can address security issues (through updates, patching, switching to other libraries, writing your own libraries, etc.). New for the 2024 exam update is the example of API being a microservice. While the section isn't intended to cover API security, you should have familiarity with the basics of API security including API authentication (strong authentication, multi-factor authentication, the downside of using static API keys (they get exposed or stolen), authorization (once you authenticate to the API, there should be limits to what you can do, such as by using scopes), validating input, rate limiting and/or throttling (to avoid denial of service or performance issues), logging and monitoring (to adhere to your organization's security policies), and the value of an API gateway (granular control over the entirety of API communications).

- **Containerization.** Many vendors have security benchmarks and hardening guidelines that you can follow to enhance container security. One challenge with containers is the lack of isolation compared to a traditional infrastructure of physical servers and virtual machines.

- **Serverless.** A serverless architecture is one where you rely on a vendor for many traditional operational responsibilities around managing servers. In a serverless model, you are dealing with a shared security model, whereby your organization and your vendor share the responsibility of security. Your organization will be focused mostly on the code security while the vendor is focused mostly on the cloud infrastructure security.

- **Embedded systems.** Embedded systems suffer from the same flaws as many computing devices. Malware, buffer overflows and web service exploits are not uncommon on embedded systems. Commonly accepted security practices for embedded systems include a secure boot feature and physically protecting the hardware (because access to the hardware opens up new attack vectors). Keeping the software and firmware up to date is also recommended.

- **High-performance computing (HPC) systems.** HPC systems are sometimes open to use for many organizations, universities or government agencies. These systems are high performance and often rented, leased or shared. This can limit the effectiveness of firewalls and invalidate the idea of air gapping the system. There are some established guidelines, such as deploying head nodes and routing all outside traffic through them, as well as isolating parts of a system. Another technique is fingerprinting how people are using HPC systems to facilitate detecting anomalous behavior.

- **Edge computing systems.** To protect edge computing systems, you can use intelligence from side-channel signals that can pick up hardware trojans and malicious firmware. You can also look at physical security, similar to how you would approach security in a data center. Finally, you can deploy IDS on the network side to monitor for malicious traffic. In many scenarios, you are an edge computing system customer and must rely on a vendor for some of the security and vulnerability remediation.

- **Virtualized systems.** In virtualized environments, you need to protect both the virtual machines and the virtual infrastructures (such as the hypervisors). Attackers often look to the hypervisor as a potential way to gain access to virtual machines (which might be better protected). Often, virtualization administrators have access to all virtualmachines and the data storage behind them, so attackers can also target virtualization admins, credentials and service accounts.

3.6 Select and determine cryptographic solutions

This section covers cryptography at a high level.

- **Cryptographic lifecycle (e.g., keys, algorithm selection).** When we think about the lifecycle of technologies, we often think about hardware and software support, performance, and reliability. When it comes to cryptography, things are a bit different: The lifecycle is focused squarely on 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 (e.g., symmetric, asymmetric, elliptic curves, quantum).** 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 enhancements performance.
  - **Quantum.** The newest method (and newly mentioned on the CISSP exam blueprint), quantum cryptography (often referred to as “post-quantum cryptography”) is still in development. There is an expectation of larger key sizes which will require more overhead. The National Institute of Standards and Technology (NIST) is currently standardizing post-quantum cryptography and expects to publish some initial standards in the next couple of years.

- **Public Key Infrastructure (PKI) (e.g., quantum key distribution).** 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 websites 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 responsible only 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 be 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.
- Quantum Key Distribution (QKD) is a new way to distribute keys vs. the traditional math-based methods. With QKD, there is a reliance on the laws of quantum physics and hardware. It is new and remains expensive and in limited use. Any attempts to eavesdrop on a quantum key exchange because the eavesdropping changes the quantum state.

- Key management practices (e.g. rotation). 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 and digital certificates.** 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. Digital certificates are typically used to secure web transactions. Certification authorities issue digital certificates to organizations and tie the certificates to domain names. Subsequently, organizations use those certificates for their web-based services to establish legitimacy and to protect the communications with encryption.

The following topics were removed from the 2024 version of the exam. We are leaving the content in the study guide to provide more context and information.

- **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 (e.g., hashing).** A hash function implements encryption with a specified algorithm but without a key. It is a one-way function. Unlikely 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.

3.7 Understand methods of cryptanalytic attacks

There are several methods to attack cryptography. Each has strengths and weaknesses and specific use cases. 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.

- **Frequency analysis.** With frequency analysis, attackers use the personality of a language to defeat substitution ciphers. For example, you can look at the frequency of a letter and compare that with the expected frequency in a language. In English, E is the most common letter, so the most common letter in an encrypted ciphertext would be expected to be a substitution for the letter E. You can also look at other areas, like the same letter appearing twice in a row, single-letter words, the most common words in a language, etc. Using a computer, you can quickly break down the ciphertext.

- **Chosen ciphertext.** In this type of attack, an attacker has access to one or more ciphertexts and their corresponding plaintext. The goal is to obtain the key using the ciphertext and plaintext. In some cases, attackers will try to trick a receiver into decrypting their modified ciphertexts to help in the attack. If an encryption scheme enables a modification to the ciphertext to cause a predictable change to the plaintext, then it is malleable. For attackers, this is desirable.

- **Implementation attacks.** In this type of attack, attackers look for weaknesses in the implementation, such as a bug in the software or outdated firmware in hardware.

- **Side-channel.** Similar to an implementation attack, side-channel attacks look for weaknesses outside of the core cryptography functions themselves. For example, a side-channel attack could target a computer's CPU. Or it could be used to gain key information about the environment during encryption or decryption. For example, an attacker could look for electromagnetic emissions or the amount of execution time required during decryption. Such information is often combined together to try to break down the cryptography.

3.8 Apply security principles to site and facility design

This section focuses on 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 unattended 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.

3.9 Design 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/intermediate distribution frame.** A wiring closet is typically the smallest 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 closest 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 interiors; people must pass through a visitor center and a controlled doorway prior to be able to enter a wiring closet. New for 2024 is the change from "intermediate distribution facilities" or "intermediate distribution frame" (more commonly referred to as an IDF. The update makes the title a bit more granular and specific to a wiring closet rack or panel.

- **Server rooms/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. In some cases, an evidence storage area should employ the most stringent security for an organization. This is to support the chain of custody. An evidence storage room can contain physical evidence (such as a smartphone) or digital evidence (such as a database).

- **Restricted and work area security.** 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 heating, ventilation, and air conditioning (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 (e.g., natural disasters, man-made).** Some buildings use water-based sprinklers for fire suppression. Water damage is possible. In a fire, shut down the electricity before turning on the water sprinklers (this can be automated). 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 (natural disasters). These issues should be considered before deciding on a data center site or a backup site. For example, you should avoid building your primary data center on the same earthquake fault line as your backup data center, even if they are hundreds of miles away from each other. 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. New for the 2024 exam update is the addition of the phrase “natural disasters” and “man-made”. The content already covered natural disasters. For the “man-made” reference, this refers to human elements such as sabotage and vandalism that might impact your environment. It could also include man-made choices, such as planting inappropriate foliage directly next to wood buildings.

- **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 that 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 a 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.

- **Power (e.g., redundant, backup).** We talked a little bit about power in the environmental issues section above. But you also need to consider designing your power to provide for high availability. Most power systems have to be tested at regular intervals. As part of your design, you should mandate redundant power systems to accommodate testing, upgrades and other maintenance. Additionally, you need to test the failover to a redundant power system and ensure it is fully functional. The International Electrical Testing Association (NETA) has developed standards around testing your power systems.

3.10 Manage the information system lifecycle

This is a brand-new topic for the 2024 exam update. Managing the information system lifecycle means taking care of a computer system from the very beginning to the very end. It starts when someone gets the idea for a new system, like a new app or tool for a company. Then it moves through planning, building, testing, and using it every day. But it doesn't stop there — the system needs updates, fixes, and security checks while it's being used. Eventually, when it's too old or no longer needed, it gets retired or replaced, and all the data needs to be handled safely. Managing this full lifecycle helps make sure the system stays useful, secure, and doesn't create problems in the future.

- **Stakeholder needs and requirements.** Stakeholders are the people who care about or are affected by the system — like customers, employees, managers, or even IT staff. If their needs and requirements aren't understood early on, the system might not work the way people expect, or it might miss key features. That's why good communication at the start (and throughout the project) really matters. By listening to stakeholders and including their feedback, the system is more likely to be useful, user-friendly, and successful from beginning to end.

- **Requirements analysis.** Requirements analysis is a key step in managing the information system lifecycle because it helps figure out exactly what the system needs to do. Think of it like making a checklist before building something — you need to know what features people want, what problems the system should solve, and what rules it has to follow (like security or privacy rules). This step usually involves talking with stakeholders, asking questions, and writing down clear goals. If you skip or rush this part, the system might be missing important stuff or cause confusion later. Getting the requirements right from the start helps make sure the system is built the right way and doesn't need a bunch of fixes later.

- **Architectural design.** Architectural design is the step in the information system lifecycle where you figure out how everything will fit and work together — kind of like drawing up blueprints before building a house. After you understand what the system needs to do (from the requirements analysis), this step is about planning the best way to build it. This includes deciding how users will connect, where the data will be stored, and how to keep things secure. A good design helps make sure the system is strong, easy to use, and can grow or change if needed. If this part is done well, it saves a lot of time and trouble later on in the project.

- **Development/implementation.** Development and implementation is the part of the information system lifecycle where all the planning and designing turns into a real, working system. During development, the system is built — this could mean writing code, setting up servers, or creating screens users will see. Then comes implementation, where the system gets tested, installed, and rolled out to the people who will actually use it. This step also includes training users and making sure everything works smoothly. If problems pop up, they're fixed before the system goes fully live. A smooth development and implementation phase helps the system start off strong and ready to do its job. The implementation is often a phased-approach where your new app or service rolls out to a pilot group, then to a small department, then to a large department, and then to the entire organization. A phased-rollout helps teams fix issues while the impacts are smaller.

- **Integration.** Integration is the step in the information system lifecycle where the new system is connected with other systems or tools the organization already uses. Think of it like adding a new player to a basketball team — they need to work well with everyone else. During integration, things like data sharing, user access, and communication between systems are set up and tested. The goal is to make sure everything works together smoothly without breaking anything. If integration is done right, the new system feels like a natural part of the environment, not something separate or confusing.

- **Verification and validation.** Verification and validation are important steps in the information system lifecycle that help make sure the system does what it's supposed to do. Verification is like double-checking the work — making sure the system was built the right way, following the design and requirements. Validation is about making sure the system actually meets the users' needs and works in the real world. It answers the question, "Did we build the right thing?" Both steps usually involve lots of testing to catch any mistakes or bugs before the system goes live. This helps avoid problems later and makes sure the system is safe, reliable, and ready to use.

- **Transition/deployment.** Transition and deployment is the part of the information system lifecycle where the system officially goes live and starts being used in the real world. It's like opening day after building and testing something new. This step involves moving the system from a test environment into the actual work environment, making sure everything runs smoothly. It also includes training users, setting up support teams, and sometimes slowly rolling out the system in stages to avoid problems. A well-planned transition makes sure the switch is smooth and doesn't mess up daily work. It's the moment when all the planning, building, and testing finally pays off.

- **Operations and maintenance/sustainment.** Operations and maintenance, also called sustainment, is the longest part of the information system lifecycle. This is when the system is up and running, and people are using it every day to do their jobs. During this time, the system needs to be taken care of — like fixing bugs, adding updates, keeping it secure, and making sure everything keeps working smoothly. Think of it like taking care of a car: you need regular check-ups, oil changes, and maybe even some upgrades over time. Good operations and maintenance keep the system useful, safe, and ready to handle whatever comes next.

- **Retirement/disposal.** Retirement and disposal is the final step in the information system lifecycle, and it happens when a system is no longer needed or is being replaced by something better. Just like cleaning out an old phone or computer before getting rid of it, this step is about shutting things down the right way. It's important to safely remove or destroy any sensitive data so it doesn't fall into the wrong hands. The system might also need to be disconnected from the network and wiped clean. A proper retirement plan helps protect the organization's information and makes sure nothing important is lost or left behind.

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 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?

a. Overloading a system or network
b. Plugging in infected portable hard drives
c. Capturing packets on a network
d. Providing invalid input
e. 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

Domain 3 Answers to 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 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 Apply secure design principles in network architecture

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

New for the 2024 CISSP exam update is the change from "Assess and implement…" to " "Apply…". There is less focus on the consulting view ("assess") and more focus on engineering..

- **Open System Interconnection (OSI) and Transmission Control Protocol/Internet Protocol (TCP/IP) models.** The 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."

- **Secure protocols (e.g., Internet Protocol Security (IPSec), Secure Shell (SSH), Secure Sockets Layer (SSL)/ Transport Layer Security (TLS)).** For the 2024 exam update, the addition of specific protocol examples was added. Be prepared to know about the most prevalent secure protocols in use today.

  - **Kerberos.** A standards-based network authentication protocol, Kerberos is used in many products. The most popular is probably Microsoft Active Directory (Active Directory Domain Services or AD DS), which is a widely deployed directory solution found in most enterprise organizations. Kerberos is mostly used on LANs for organization-wide authentication, single sign-on (SSO) and authorization.
  - **SSL and TLS.** Data protection solutions most popular for protecting website transactions such as banking and e-commerce, SSL and TLS both offer data encryption, integrity and authentication. SSL is the original protocol, but it is considered a legacy and insecure protocol today. Instead, TLS is mostly used. TLS was initially introduced in 1999 but didn't gain widespread use until years later. The original versions of TLS (1.0 and 1.1) are considered deprecated and organizations should be relying on TLS 1.2 or TLS 1.3.
  - **SFTP.** This is a version of FTP that includes encryption and is used for transferring files between two devices (often a client and a server).
  - **SSH.** This remote management protocol, run over TCP/IP, is mostly used by IT administrators to manage devices such as servers and network devices. All communications are encrypted.
  - **IPSec.** IPSec is an IETF standard suite of protocols that is used to connect 2 points (such as computers or office locations) together. It is widely known for its use in virtual private networks (VPNs). IPSec provides a variety of security over IP networks, such as encryption, authentication and data integrity.

- **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 theses layers is called encapsulation. 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 (e.g., Internet Small Computer Systems Interface (iSCSI), Voice over Internet Protocol (VoIP), InfiniBand over Ethernet, Compute Express Link).** For the 2024 exam update, Fiber Channel Over Ethernet (FCoE) has been removed as an example protocol and there is a focus on more cloud-friendly protocols. Like encapsulation, converged protocols enable communication over different mediums.

  - **Fiber Channel Over Ethernet (FCoE).** Sends fiber channel control commands over Ethernet. As of the 2024 exam update, this example is no longer part of the title but leaving this in the study guide.
  - **Internet Small Computer Systems Interface (iSCSI).** Used in storage area networks, iSCSI sends SCSI commands over IP networks.
  - **Voice over Internet Protocol (VoIP).** 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.
  - **InfiniBand over Ethernet.** InfiniBand over Ethernet is a networking approach that combines the high-performance features of InfiniBand with the widespread availability of Ethernet. InfiniBand is known for

its low latency and high throughput, often used in high-performance computing (HPC) environments and data centers. By running InfiniBand protocols over Ethernet—typically using technologies like RoCE (RDMA over Converged Ethernet)—organizations can achieve similar performance benefits without needing a completely separate InfiniBand network. This allows for Remote Direct Memory Access (RDMA), which lets systems transfer data directly between memory without involving the CPU, improving efficiency and reducing delays.

- **Compute Express Link.** Compute Express Link (CXL) is a new type of high-speed connection that helps computers share memory and data more quickly and efficiently. It's similar to giving different parts of a computer — like the processor and memory — a faster, smarter way to talk to each other. CXL is especially helpful in big data centers or cloud environments where systems need to work together in a high performing environment. It reduces delays and helps computers use memory more flexibly, which can boost performance.

- **Transport architecture (e.g., topology, data/control/management plane, cut-through/store-and-forward).** Transport architecture is about how data moves across a network — kind of like planning roads and traffic lights so everything flows smoothly and safely. It includes the topology, which is the shape or layout of the network, like a star, ring, or mesh. There are also three main parts of how a network functions: the data plane (moves the data), the control plane (decides where the data should go), and the management plane (handles settings and updates). Two ways to move data are cut-through (data is sent before it's fully received, which is fast but adds risk) and store-and-forward (waits for the whole message before sending, which is slower but reduces risk).

- **Performance metrics (e.g., bandwidth, latency, jitter, throughput, signal-to-noise ratio).** Performance metrics are ways to measure how well a network is performing (speed and latency, as an example). Bandwidth is the amount of data that can move through the network at one time, like the width of a highway. Latency, often measured in milliseconds, is the delay or how long it takes for data to travel, like how fast a text message gets delivered. Jitter is when there's an uneven delay — some messages are fast, others are slow — which can mess up things like video calls. Throughput is the actual amount of data that successfully moves through at a given time. Signal-to-noise ratio compares good signals to background interference — higher is better.

- **Traffic flows (e.g., north-south, east-west).** Traffic flows describe the direction that data moves in a network. North-south traffic is data that goes in and out of the network, like from a user's computer to the internet or a cloud service — think of it as going up and down. East-west traffic is data that moves inside the network, like between servers, applications, or virtual machines — this goes side to side.

- **Physical segmentation (e.g., in-band, out-of-band, air-gapped).** Physical segmentation is a way to keep parts of a network physically separated to increase security and reduce risk. There are a few ways to do this. In-band means control and data travel on the same network — it's easier to manage but riskier if something goes wrong. Out-of-band means control traffic goes on a different, separate path such as a dedicated network. Then there's air-gapped, which is the most secure — it means the system isn't connected to any network at all. It's completely offline, like locking a computer in a room with no internet. As networks become more physically segmented, they become more difficult to use so there is a balance between usability and security to account for.

- **Logical segmentation (e.g., virtual local are networks (VLANs), virtual private networks (VPNs), virtual routing and forwarding, virtual domain).** Logical segmentation is a way to separate parts of a network using software instead of physical cables or hardware. This helps keep data organized and adds extra layers of security. One common way to do this is using a VLAN (Virtual Local Area Network), which groups devices together even if they aren't physically close. A VPN (Virtual Private Network) creates a safe tunnel for your data to travel through, even on public networks. Virtual Routing and Forwarding (VRF) lets one router act like many, keeping traffic from different users separate. Virtual domains split up tasks or users within devices like firewalls, so each group has its own space. Logical segmentation is often combined with physical segmentation in a large network.

- **Micro-segmentation (e.g., network overlays/encapsulation; distributed firewalls, routers, intrusion detection system (IDS)/intrusion prevention system (IPS), zero trust).** This is a new topic for the 2024 exam update, although it includes some sub-topics that were part of the exam already.

  - **Network overlays/encapsulation.** Network overlays and encapsulation are tools that help make micro-segmentation work better, especially in complex or cloud-based networks. Micro-segmentation means breaking a network into smaller parts to control traffic and stop threats from spreading. A network overlay is like a virtual map built on top of the real network — it helps create these smaller segments without changing the physical hardware. Encapsulation wraps data in a special "package" so it can travel safely across this virtual network. It's kind of like putting a toy inside a box so it doesn't get damaged on the way. These tools help keep different parts of the network isolated and secure, even if they're using the same hardware underneath.
  - **Distributed firewalls.** Distributed firewalls are an important part of micro-segmentation, which is a way to divide a network into smaller, safer zones. Instead of having just one big firewall at the edge of the network, distributed firewalls put security checks on each device or virtual machine. It's kind of like giving every door in a building its own lock, not just locking the front door. This way, even if an attacker gets into the network, they can't move around easily. Each segment has its own rules and protection.
  - **Routers.** Routers play a big role in micro-segmentation because they help control how data moves between different parts of the network. In micro-segmentation, the network is split into smaller zones to make it more secure. Routers help keep these zones separate by deciding which traffic is allowed to go from one zone to another. Think of a router like a traffic cop at an intersection, making sure only the right cars go the right way. This helps stop bad traffic from spreading across the network. When set up with the right rules, routers can make sure that each segment only talks to the parts it's supposed to — keeping everything more organized and protected.
  - **Intrusion detection system (IDS)/intrusion prevention system (IPS).** Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS) are tools that help protect networks, and they work really well with micro-segmentation. IDS watches network traffic and alerts you if something suspicious is happening, kind of like a security camera. IPS takes it a step further and can actually block bad traffic, like a security guard stepping in. When you use micro-segmentation to break the network into smaller pieces, IDS and IPS can be placed closer to each segment. This means they can watch traffic more closely and stop threats faster before they spread. Together, they help make each small part of the network safer and easier to manage.

- **Zero trust.** Zero Trust is a security idea that works really well with micro-segmentation. The main rule of Zero Trust is "never trust, always verify." That means no one and nothing is trusted automatically, even if they're already inside the network. Micro-segmentation helps make Zero Trust possible by breaking the network into small, secure pieces and setting up strict rules about who can access what. Each piece has to prove it's allowed to talk to another — just like showing your ID every time you enter a new room. This helps stop attackers from moving around the network if they get in. With Zero Trust and micro-segmentation working together, it's a lot harder for bad guys to do damage.

- **Edge networks (e.g., ingress/egress, peering).** This is a new topic added during the 2024 exam update. As the name implies, edge networking is networking at the edge of your primary network. For example, a corporate network edge is the very edge of the network, the last hop before the internet. For traffic sourced from the internet and going to your corporate network, the edge is the first device when leaving the internet and connecting to your corporate network.

  - **Ingress. Network traffic coming into a network from an outside network is considered ingress traffic.** For example, traffic coming from the internet to your corporate network is considered ingress traffic.
  - **Egress.** Network traffic exiting a network is considered egress traffic. For example, traffic from a computer on your corporate network going to an internet destination is considered egress traffic.
  - **Peering.** When two networks exchange traffic and/or are interconnected directly, it is called peering. For example, two providers in a co-location facility could opt to peer to facilitate optimized traffic exchange without utilizing the internet.

- Wireless networks (e.g., Bluetooth, Wi-Fi, Zigbee, satellite). While consumers generally see a wireless network as a single type of network that their phones can connect to, there are many other types you need to be aware of for the exam. Regardless of the security method you use, you should also use TLS or IPsec for network communication. Remember that wireless networks use collision avoidance, instead of the collision detection used on wired networks. This topic was updated during the 2024 exam update by changing Li-Fi to Bluetooth.

  - **Bluetooth.** Bluetooth is a wireless technology that lets devices talk to each other over short distances. It's often used to connect things like headphones, keyboards, speakers, or even smartwatches to your phone or computer without needing any cables. Bluetooth uses the 2.4 GHz spectrum. It has lower performance than Wi-Fi but consumes less power, which is why you often see it used in small mobile devices (such as pencils for tablets). From a security perspective, be aware of Bluejacking (unsolicited messages using Bluetooth), bluesnarfing (unauthorized access and/or information theft over Bluetooth), and bluebugging (can obtain full control of a target device over Bluetooth). A common way to minimize security issues with Bluetooth is to disable the protocol when not in use and be wary of Bluetooth connections, especially in public.
  - **Wi-Fi.** 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 3,466 Mbps
Ax (Wi-Fi 6)	2.4 and 5GHz	Up to 9.5 Gbps
Be (Wi-Fi 7)	2.4, 5, and 6 GHz	Up to 46 Gbps

- Wired Equivalent Privacy (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) 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) 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).

- **Zigbee.** Zigbee is a wireless protocol based on IEEE 802.15.4 and designed for IoT devices and other devices operating in a small area, such as a home or hospital room (sometimes called personal area networks). It has distance limitations of about 100 meters, although that can be extended with a mesh network. Some popular devices that use and/or support Zigbee include the Amazon Echo Plus, Samsung SmartThings and Philips Hue smart lights. Zigbee is very popular for home automation products.
- **Satellite.** Satellite is a type of wireless networking. It uses a satellite on both ends of a connection, such as a satellite on a house and one or more satellites in space. Satellite is widely known for providing broadband-like internet access to rural areas that don't have access to traditional internet connections such as fiber or cable. The primary downside of satellite networking is latency: Because data has to go into space and back, latency is higher than with fiber or cable. Satellite typically doesn't compete with fiber or cable or other high-speed connections.

- **Cellular/mobile networks (e.g., 4G, 5G).** Smartphones rely on cellular networks for voice and data. With the advent of 3G (third generation), smartphones got reliable internet connections for the first time. 4G (fourth generation), the most prevalent solution today, greatly expanded cellular performance. 5G (fifth generation) offers even greater performance and lower latency. However, 5G has limited coverage. This topic's title was changed from "Cellular networks…" to "Cellular/mobile networks…" during the 2024 exam update.

- **Content distribution network (CDN).** 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.

- **Software defined networks (SDN), (e.g., application programming interface (API), Software-Defined Wide Area Network, network functions virtualization).** Software-Defined Networking (SDN) is a newer way to control networks using software instead of just hardware. In a regular network, you'd have to change settings on physical devices like switches or routers, which can be slow and hard to manage. With SDN, you can use software and APIs (application programming interfaces) to control how data moves across the network — kind of like using a remote control instead of walking over to press buttons. SD-WAN (Software-Defined Wide Area Network) is a type of SDN that helps manage traffic between different locations, like between office branches. And Network Functions Virtualization (NFV) lets you run things like firewalls or load balancers as software instead of needing special hardware.

- **Virtual Private Cloud (VPC). A Virtual Private Cloud (VPC) is like having your own private space inside a public cloud, such as Amazon Web Services (AWS) or Microsoft Azure.** Even though you're using the same public cloud as other people, your VPC is separated and secure — kind of like having your own private room in a big hotel. In a VPC, you can choose how your network is set up, including things like IP addresses, firewalls, and gateways. You control who can come in, what they can access, and how your data moves around.

- **Monitoring and management (e.g., network observability, traffic flow/shaping, capacity management, fault detection and handling).** Monitoring and management means keeping an eye on how a network is working and making sure everything runs smoothly. This includes network observability, which helps you see what's going on inside the network — like a dashboard showing traffic and system health. Traffic flow and shaping is about controlling how data moves so the network doesn't get too crowded or slow down. Capacity management means planning ahead to make sure there's enough space and power for everything the network needs to do. And fault detection and handling is about finding problems quickly and fixing them before they cause bigger issues.

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 infrastructure (e.g., redundant power, warranty, support)**. 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. Be familiar with these 3 examples:

  - **Redundant power.** In a home network, most devices (such as computers and routers) have a single power supply. If that power supply fails, the device loses power. In a work environment (small or large business or organization), it is a good practice to use multiple power supplies to provide redundant power. Redundant power is typically reserved for infrastructure components such as servers, routers and firewalls, but not for end-user devices such as laptop or desktop computers. Redundant power is often combined with other redundancies (redundant servers, redundant switches, etc.) to provide high availability.
  - **Warranty.** One aspect of managing hardware is managing warranties. Hardware typically comes with warranty coverage that covers hardware failures in the first year or few years. Organizations should opt for warranties whenever possible to minimize the costs of repairing and replacing hardware.
  - **Support.** Another aspect of managing hardware is having support. Support for hardware is important. Many IT shops are not skilled in repairing some types of hardware (for example, if a motherboard goes out in a router or if a screen goes out in a laptop computer). Hardware support typically offers on-site repair or replacement and many guarantee their service level (same-day, next-day, etc.). Because procuring replacement hardware is time consuming, organizations can't count on it in the way they can count on support from the vendor.

- **Transmission media (e.g., physical security of media, signal propagation quality).** 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. New for the 2024 exam update is a call out of the physical security of media and the signal propagation quality. When it comes to CAT5E and CAT6 cables, they are susceptible to electromagnetic interference (EMI) whereas fiber optic cables are immune to EMI. Coaxial cables offer some EMI protection but are not immune to EMI. Fiber optic offers high security but physical taps are possible. In all cases, you should provide protection against tapping or tampering of transmission media through the use of locked closets, server rooms, and cable routing in private places (ceiling as one example). The signal propagation quality determines how well the network transmission fairs from a source to a destination. Many factors come into play, such as attenuation (the longer the distance, the weaker the signal), latency, jitter, and crosstalk (one channel interfering with another channel).

- **Network Access Control (NAC) devices (e.g., physical, and virtual solutions).** New for the 2024 exam update is the mention of physical and virtual solutions. Much as you need to control physical access to equipment and wiring (locked closets, server rooms), you also 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 perimeternetwork. 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 (e.g., host-based).** 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 is 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. Most enterprises use a combination of edge security, network security, enterprise security, and host-based solutions to maximize protection. Beyond the traditional endpoint protection methods, there are others that provide additional endpoint 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 strategy 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.

Physical security is one of the most important aspects of securing network components. 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, video, and collaboration (e.g., conferencing, Zoom rooms).** 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 by 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. New for the 2024 exam update is the addition of video and collaboration in the title.

  - **Video.** Video refers to one-on-one video calls and video streams (for example, corporate meeting or training). From a security perspective, you need to ensure only authorized people have access to the video (to minimize the risk of data exfiltration as one example). Similar to protecting other communication mediums, such as instant messaging, you should configure end-to-end encryption and update related video software (computer, mobile) regularly. With the advent of AI, it is now possible for malicious actor to appear as somebody else on a video call. To combat AI impersonation
  - **Collaboration.** This topic includes all of the tools you would use to work with somebody else such as instant messaging, file sharing, and screen sharing. Enterprises must protect against data leakage, unauthorized access (such as sharing with somebody outside of the company), and shadow IT (users installing unapproved

collaboration software). The most common ways to combat against collaboration attacks is using strong access controls (principle of least privilege via permissions, MFA), implement a Data Loss Prevention (DLP) solution to stop sensitive information from leaving the organization, and configuring the collaboration tools to maximize security (for example, screen sharing requires authentication, meetings have passwords and are restricted to being accessible from your corporate network).

- **Remote access (e.g., network administrative functions).** 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 (IPsec on layer 3 or SSL/TLS on layer 7) 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). New for the 2024 exam update is the reference of network administrative functions which is a reference to the various configurations of network-related hardware and software related to remote access. Remote access, whether VPN or similar, should be strongly authenticated (for example, username/password + second authentication factor). Many organizations opt to require a VPN connection to use other remote access methods, such as SSH and RDP.

- **Data communications (e.g., backhaul networks, satellite).** Whether you are physically in an office or working remotely, communication between the devices being used should be encrypted to prevent 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, communication should be encrypted using TLS, SSL or IPSec. New for the 2024 exam update is the addition of backhaul networks and the reference to satellite as a data communication technology. Backhaul networks are used by very large enterprises with an extensive Wide Area Network (WAN), ISPs, and cellular providers to connect various sites to a headquarters (or main) site. Satellite connectivity is a newer method of wireless networking that reaches remote areas that are difficult for traditional network providers to reach. Satellite's strongest selling point is the wide availability (can be used on cruise ships, and rural towns "in the middle of nowhere"). However, satellite has high latency vs. traditional Ethernet or local Wi-Fi and is vulnerable to eavesdropping and/or interception.

- **Third-party connectivity (e.g., telecom providers, hardware support).** This refers to third-party organizations (vendors, such as IT auditing vendors) connecting to your network. Attackers are routinely looking at creative ways to gain access to organizations, and third-party connectivity is one option. Back in 2014, attackers gained access to a major retail chain (Target) by using credentials obtained from third-party connectivity (a heating/refrigeration/air- conditioning company that serviced Target).

As organizations evaluate third-party connectivity, they need to look carefully at the principle of least privilege and at methods of monitoring use and misuse.

The following topics were removed from the 2024 CISSP version of the exam (however, we are leaving the content in the guide to provide additional information):

- **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.

- **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

Domain 4 Answers to 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, such as 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 (such as read, write or delete) for the specific asset. 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 (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. Your access is based on your group memberships. For example, one security group might have read access to an asset, while a different security group might get read/write/execute access to that 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 (multifactor authentication and the different factors are explained in section 5.2). 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 multifactor 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.

- **Applications.** There are a variety of methods to control application access. In some cases, you can map application roles (read-only, power user, admin, etc.) to security groups (such as Active Directory security groups), for example, so that members of the "Power Users" group get power user rights in the application. Be familiar with role-based access control (RBAC) for application access. For example, you might have a role named "Database Administrator" that is made up of 4 groups — SQL DBAs, DB Report Writers, Oracle DBAs and Crystal Reports Admins; each of those groups has rights to various database- related applications.

- **Services.** In Identity and Access Management (IAM), services are tools or systems that help control who can access what. These services make sure that the right people get into the right places — and keep the wrong people out. For example, some IAM services handle authentication, like single sign-on (SSO), so users only need one password to access many tools. Others check permissions, making sure a person can only see or use what they're supposed to. Some services even watch for anomalous behavior and send alerts if something is unusual.

5.2 Manage identification and authentication strategy (e.g., people, devices, and services)

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

- **Groups and Roles.** This is a new topic for the 2024 exam update.

  - **Groups.** Groups are a collection of users. Groups simplify the management of users for actions such as granting access, sharing, and communicating. Without groups, administrators have to individually manage users, which adds administrative overhead and increases the risk of an administrator error. There are security groups, which are used to provide access to group members. There are distribution groups, which are mostly used to communicate to group members.
  - **Roles.** Roles are a collection of groups or a collection of permissions. Roles simplify the management of user permissions. Let's look at an example. A company has an infrastructure IT team that manages servers, databases, and 3 applications. There is a security group used to provide access to the servers. There is a group used to provide access to the databases. And there is 1 group for each application to provide admin access to each application. Without roles, each infrastructure team member has to be added to each group. Every time a new team member starts, they have to be added to each group. With roles, users can be added to a single role and get access to the servers, databases, and applications. Identity software often helps facilitate roles. Some smaller companies mimic the functionality of roles by nesting groups within a parent group (with the parent group become the "role"). Roles also ensure that every team member has the same access. In large organizations, it is common for people to have more than one role based on their job function. It is also common to have a yearly process to review role assignments to ensure they are still needed for each user.

- **Authentication, Authorization and Accounting (AAA) (e.g., multi-factor authentication (MFA), password-less authentication).** In the 2024 exam update, this topic was updated to reference accounting and password-less authentication. Single-factor authentication requires a single method to authenticate, such as a username/password combination or a smartcard. Multifactor authentication requires methods from two or more factors as part of the authentication process. There are 3 category factors: something you know (such as a password), something you have (such as a mobile phone) and something you are (such as an iris scan). For example, in MFA, you might first enter your username and password; if they are valid, you then provide a soft token (random number) from a security application (such as a mobile app) as the second method. This process adds security because an attacker who steals your credentials is unlikely to also have access to your phone. In general, requiring more methods enhances security, but different factors provide different levels of security. For example, answering a question isn't usually as secure as a smartcard or token from a security app, because the information to answer the question might be readily available on the internet or otherwise known by a malicious user. One downside to multifactor authentication is the complexity it introduces. If a user doesn't have their mobile phone or token device, they can't sign in. To minimize issues, you should

provide options for the second method (for example, the user can use biometrics on their computer for MFA). Password-less authentication refers to a user authenticating without typing a password. One common example is Windows Hello for Business, which enables users to create a unique PIN per device or use biometrics (fingerprint or facial recognition) to logon to a computer without a password. Note that Single Sign-On (SSO) provides the appearance of password-less authentication since a password isn't used. Instead, most SSO providers rely on the user having a successful authentication to their device.

- **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, proofing and establishment 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's Forefront Identity Manager which also uses the FIM 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 providingthe 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 standards for exchanging authentication and authorization information. SAML is the most common standard in use today, but 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.

- **Credential management systems (e.g., Password vault).** New for the 2024 exam update is the call out of a password vault (which is another way to say "credential management system"). A credential 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 accounts, 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 check out a database admin account in order to perform some administrative work using that account, and check the account back in when the job is finished; then 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. Many large organizations have two or more systems for managing credentials: a centralized system where shared accounts are stored for teams and a personal app for users to maintain their own passwords.

- **Single sign-on (SSO).** SSO 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 (discussed above). 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 re-authenticating. Typically, the SSO experience will last for a specified period, such as for 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 enhanced too. For example, a user might be able to use SSO to seamlessly gain access to a web-based portal. However, if the user attempts to make a configuration change, the portal might prompt for authentication or even require MFA. 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. Multifactor authentication can help mitigate this risk.

- **Just-in-time (JIT).** Historically, users and administrators were given permanent access to the apps and services they need as part of their job. But having standing access makes it easier for attackers to gain access or expand their access. Now, many organizations are moving to a model where users and/or administrators gain access to what they need only when they need it. This is known as "just-in-time" (JIT). Imagine I'm a database administrator and I need to make a change to the logging confirmation on a server. With JIT, I request access to do that and it is granted either manually or automatically based on a set of rules. Once I finish with my change, the access automatically goes away. This means I don't have elevated database admin rights while I'm reading email or performing unrelated admin work. JIT helps organizations adhere to the principle of least privilege. JIT is often combined with granular permission models that provide the minimum access required for the task, instead of defaulting to "full admin" for every task.

For the 2024 exam update, the following topic was removed. We are keeping them in the study guide to provide more information and context.

- **Identity management (IdM) implementation.** This topic covers IdM implementations in general. When thinking about implementing an IdM solution, the same implementation factors as for other core infrastructure services are relevant. For example, you need to figure out the best way to ensure the apps and services are highly available, site resilient and secure. For IdM related apps and services, performance is important. You need to minimize latency. Identity-related systems are generally spread out across an entire organization, so authentication and authorization should occur as close as possible to the user, app or service to maximize performance. Because IdM components are part of securing your environment, you must be vigilant with security patching and bug fixes. You should also be familiar with the components, apps and services that are involved in a complete IdM implementation. In an enterprise environment, it is common to have the following: a directory service (such as Active Directory), a password/key/secrets management service, a single sign-on solution (such as Ping Identity or Microsoft Azure Active Directory), identity services in the cloud such as Amazon AWS IAM, an identity lifecycle solution (that manages the creation, modification, and deletion of roles, users, and groups), a self- service password reset tool, an auditing and compliance tool (for scanning and reporting), a multifactor authentication solution, and a password manager (for users and/or administrators).

5.3 Federated identity with a third-party service

There are many third-party services offering identity services in the cloud and on-premises. In most cases, these identity services integrate with your existing identity and access management systems.

- **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, enabling 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. For example, Microsoft Azure offers Azure Active Directory (along with premium versions that have more features). Ping Identity and Okta provide cloud-based federated identity solutions. 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 to get started with.
  - 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 management. 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 managing just 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 multifactor authentication solution, it might not work seamlessly with your cloud-based identity service.

- **Hybrid.** Many organizations cannot get by using just an on-premises identity provider or just a cloud-based identity provider; instead, they require both. In such cases, the identity components are often integrated or work together to provide a complete solution. This mix of on-premises solutions and cloud solutions is referred to as a hybrid approach.

Third-party identity services are provided by a vendor, often as a complement to your identity service. For example, Ping Identity is a vendor that provides an identity platform that you can integrate with your existing on-premises directory (such as Active Directory) and with your public cloud services (such as Microsoft Azure or Amazon AWS). Cloud vendors are beginning to encroach upon the third-party identity services market and the solutions are often competitive with third-party identity services. The key facts about third-party identity services are:

- Often, you still need an on-premises directory service.
- Some third-party identity services cannot 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, multifactor authentication and meta-directory services (pulling in 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, which has rights to workstations, the anti-virus software and a software installation shared folder. If a new desktop technician starts at your company, they can be added to the role group and quickly have the same access as other desktop technicians. RBAC is a non-discretionary access control method, because there is no discretion — the role has what it has. RBAC is considered an industry-standard good practice and is in widespread use.

- **Rule-based access control.** Rule-based access control implements access control based on predefined rules. For example, you might have a rule that enables read access to marketing data for anyone who is in the marketing department, or a rule that enables you to print to a high-security printer only if you are a manager. 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 enables anybody in the New York office to access a file server there. 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 (see the next bullet) 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).** In most identity and access management systems, user objects have optional attributes. For example, user objects might have an attribute to store the employee number or the date of hire. Many organizations use attributes to store data about users, such as their department, cost center, manager and location. These attributes can tie into your authorization systems to help automate authorization. For example, if you deploy a wireless network for your office in Paris, you can configure the wireless authorization

such that only users that have “Paris” listed as their office location (in the attribute of their user objects) can use the wireless network. Attribute-based access control also help secure authorization, because you can layer it into existing authorization mechanisms. For example, if a user has to be a member of a group to gain access to the HR shared folder, you can add in attribute-based access control by configuring the shared folder to accept members of the group only if their department attribute is “HR”. In such a scenario, you need two items to gain access — the group membership and the attribute populated correctly.

- **Risk-based access control.** Risk-based access control is gaining in popularity. When access to a resource is requested, the risk is evaluated to determine whether to allow the access, deny the access or require additional authentication (such as MFA). Risk is evaluated on factors such as authentication metadata like location (known, unknown, on blacklist), IP address (known to be malicious or previously used), and device information (device type, OS and version, type of client app). Some risk-based access control solutions can use a plethora of metadata. Risk based access control is often combined with other access control methods, such as role-based access control.

- **Access policy enforcement (e.g., policy decision point, policy enforcement point).** This is a new topic from the 2024 exam update. Access policy enforcement is about making sure people can only get to the stuff they’re allowed to — no more, no less. When someone tries to access something, like a file or a system, there are two main parts that help make the decision. The Policy Decision Point (PDP) is the ”brain” — it looks at the rules and decides if access should be allowed. The Policy Enforcement Point (PEP) is the gate — it blocks or allows the person based on what the PDP says. Together, they help enforce access rules and keep things secure.

5.5 Manage the identity and access provisioning lifecycle

The identity lifecycle includes the creation of users, the provisioning of access, the management of users, and 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. Such systems are often 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 user’s email address and role might be added.
5. As needed, 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. After a specific period (such as 7 days), the IT department permanently deletes the user account and all associated access.

Beyond the steps captured above, there are additional tasks involved in managing identity and access. For example, you should perform periodic access reviews in which appropriate personnel attest that users have the appropriate rights and permissions. You should review the configuration of the identity service to ensure it adheres to known good practices. You should review the directory service for stale objects, such as user accounts for employees who left the company a long while back. The primary goal is to ensure that users just have what they need, nothing more. If a terminated user still has a valid user account, then you are in violation of your primary goal. The lifecycle focus for the exam is on the following three areas:

- **Account access review (e.g., user, system, service).** Account access review, sometimes referred to as periodic access review (PAR), is a formal review of access. Does the user account, system account or service account have only what is required to perform their job or task? Was the access established based on the company's access request policies? Is the granting of access documented and available for review? User accounts are accounts used by users to perform their daily job duties such as read email, research, and create new content. System accounts are accounts that are not tied to one-to-one to a human. 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 and/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 of access and where the access was provided to. While "system account" generally covers all non-human accounts, a subset of those accounts are referred to as "service accounts." A service account is generally used to run a service (like a web service or a database service). Sometimes, "system account" and "service account" are treated as synonymous, but a service account isn't always tied to a service and a system account could be used to run a service.

- **Provisioning and deprovisioning (e.g., onboarding, offboarding and transfers).** 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/or delete accounts and you also have dormant accounts that can be targeted. It is a good practice to automate provisioning and deprovisioning whenever feasible, because automation speeds the process and reduces human error (although the automation code could have human errors). Your company should establish guidelines for account provisioning and deprovisioning. For example, your company might have a policy that an account is disabled on an employee's last day or as the employee is being notified that they are being terminated. You also need to ensure proper provisioning and deprovisioning during transfers When somebody moves to a new position or a new department, their access should change to match their new job responsibilities, which can involve both adding new access and deleting any access they had before that is not needed for the new role. By using RBAC, you can automate much of the work around transfers.

- **Role definition and transition (e.g., people assigned to new roles).** For the 2024 exam update, this topic was updated to include a reference to transition. The roles used in role-based access control need to be created and properly defined. For example, the Accounting role should grant access to run accounting software. Many organizations use roles to define entire jobs. For example, a network administrator requires access to manage the wireless access points, the routers, the switches, and the firewalls. Instead of individually adding users to several groups, the users can be added to roles and the roles contain all the groups needed to gain access to the systems for their job. Sometimes, people change roles and require new access. This process is often referred to as the "mover" process in a joiner (new hire), mover (role change within a company), and leaver (somebody retiring or switching to a new company) system. During job transitions, it is important to remove access the user no longer needs and to provide access to roles that the new job requires.

Privilege escalation (e.g., use of sudo, auditing its use). This topic is was updated in 2024 with a reference auditing sudo use). IT administrators should use an account that doesn't have elevated access for tasks like email and document creation. Only when they need to perform an administrative task that requires more access should they escalate their privilege (sometimes via built-in methods in the operating system and sometimes with third-party privilege management solutions). Privileges should be escalated only as much as needed; for instance, admins should not use high-level accounts for day-to-day administrative tasks like stopping and starting a service.

  - **Managed service accounts.** A managed service account is a service account whose password is automatically managed by either the operating system or a third-party app. Administrators don't have to manage the service account and don't need to know the password. These types of accounts reduce risks because the passwords are unknown and password management is automatic.
  - **Auditing the use of sudo.** Sudo is a method used by administrators to run commands as the root account on Linux (although you can customize the configuration to use other non-root accounts). With sudo, users do not need to know the root password and they don't need to log on as root (both good things from a security perspective). It is a good practice to use sudo instead of logging on as root or a root equivalent account. While sudo is built into Linux distribution, other operating systems such as Windows have an escalation mechanism too (on Windows, "Run as administrator" or "Run as different user"). Organizations need to be able to determine who is using sudo, when they are using it, and what they are doing with it. This is often referred to as "accounting". Some organizations opt to record sudo sessions via auditing software or screen capture technology.
  - **Minimizing its use.** While sudo provides a way for administrators to temporarily escalate their privileges, it has some downsides. First, it isn't a centralized solution in which all servers are controlled by a single sudo configuration file, so there is some overhead with setting up and maintaining sudo in a large environment. Second, the configuration is controlled by local files on servers, which can be modified by administrators. In addition, Sudo lacks capabilities of other privilege escalation solutions, such as screen recording, multifactor authentication and risk-based authentication.

- **Service accounts management.** This is a new topic introduced in the 2024 exam update. Service accounts are special types of accounts that are used by computer programs or systems — not often by people — to do tasks like running apps, running services or daemons, or connecting to databases. Just like user

accounts, service accounts need to be managed carefully as part of the identity and access provisioning lifecycle. That means keeping track of when they're created, what they have access to, and when they're no longer needed. If service accounts have too much access or are left enabled when they're not being used, they can become a big security risk. Service account credentials (passwords, secret keys, certificate, or similar) should be stored in a centralized credential management system.

5.6 Implement authentication systems

The sub-topics in this section were removed in the 2024 exam update. The content below is still valid for the new exam. This section is very similar to 5.2 but covers things from a higher level, focusing on the authentication systems instead of the authentication process and details. This section does not specifically call out design, but you should understand the design considerations of authentication systems in addition to the implementation information. At this end of this section, I walk through some design and implementation considerations. First, we look at the specific authentication systems called out in the exam blueprint.

- **OpenID Connect (OIDC)/Open Authorization (OAuth).** OAuth is an open framework used for authentication and authorization protocols. It is designed for the internet, where identity federation is mostly used. The most common protocol built on OAuth is OpenID Connect (OIDC). OAuth 2.0 is often used for delegated access to applications — for example, a mobile game that automatically finds all of your new friends in a social media app is likely using OAuth 2.0. On the other hand, if you sign into a new mobile game using a social media account (instead of creating a user account just for the game), that process might use OIDC.

- **Security Assertion Markup Language (SAML).** SAML is a long-standing authentication solution that many companies use to give their users access to applications on the internet. It enables users to sign into their company computer and then gain access to all of their applications without having to re-enter their credentials.

- **Kerberos.** Kerberos is a network authentication protocol widely used in corporate and private networks and found in many LDAP and directory services solutions such as Microsoft Active Directory. It provides single sign-on and uses cryptography to strengthen the authentication process.

  Remote Authentication Dial-In User Service (RADIUS)/Terminal Access Controller Access Control System Plus (TACACS+). RADIUS is a network authentication, authorization and accounting protocol that is often used by VPNs and network components. RADIUS is mostly used on internal networks, as it wasn't designed for today'sinternet. TACACS+ is a remote authentication protocol which sometimes competes with RADIUS. TACACS+ is newer, encrypts more of the data and offers support for more protocols. Both protocols are used to authenticate a user (such as a network admin) to a destination (such as a firewall that the admin manages).

- **Design considerations.** Authentication systems participate in authenticating users, computers, service accounts, and more, so they are mission-critical. Often, they are at the most critical tier of an IT infrastructure, along with the network components. Accordingly, there are some important design considerations:

  - **High availability.** Your authentication systems should be highly available within a site, such as an office in Frankfurt. That might be as simple as having two servers, but it can be more complex, based on the size of the network and the security requirements. Everything that your authentication systems rely on should be highly available, too — storage and network being a couple of examples.
  - **Site resilience.** When your authentication systems are site resilient, it means that people can continue to authenticate even if your main office goes down. With critical IT services like authentication, you should always design for site resiliency.
  - **Performance.** In many organizations, people need to authenticate often, sometimes a few times an hour or even 50 or 100 times a day. Authentication should be seamless, occurring in the background quickly and easily. To achieve a seamless authentication system, you need to design for adequate performance. If it takes too long to authenticate, users become frustrated and productivity drops. To achieve adequate performance, you need to go through a sizing exercise. It is helpful to use vendor documentation for performance, especially if the vendor offers a reference architecture or sizing guides. To maintain optimal performance, you must consider monitoring your authentication environment: Establish a baseline and set up alerting when performance goes outside of normal.
  - **Security.** Because authentication systems are the foundation for gaining access to apps and data, you need to design them with high security. Limit administrative access to authentication systems to authentication administrators only. Break authentication administrators into levels; for example, Level 1 can perform basic admin tasks, Level 2 can perform basic and advanced admin tasks, and Level 3 can perform all admin tasks. Place authentication systems only in highly secure facilities such as data centers or other locations with biometrics, multiple authentication factors, locked cabinets, security cameras and security guards. Require MFA for all administrators. Require the use of administrative accounts and administrative workstations for performing admin work.

- **Implementation considerations.** The implementation goal is to meet or exceed the design requirements. Steps include figuring out how to meet the design (including choosing options), performing the implementation, and validating the outcome. Keep these implementation considerations in mind:

  - **High availability.** If your design calls for high availability, you need to choose a method. Should you use failover clustering? Hardware or software load balancing? DNS round robin? Partly, it will depend on the requirements of the high availability and the hardware and software available to you. In some scenarios, you might combine multiple options. For example, you might have 4 identical servers and use load balancing to provide high availability.
  - **Site resilience.** If your design calls for site resilience, how will you achieve that? Are you starting with a single data center or do you already have a secondary site to use? Are you able to scale your high availability across multiple sites? Does your load balancing solution offer multi-site load balancing (geographical load balancing)? Not only do you want to provide site resilience, but you want to do it in a way that doesn't add too much administrative overhead.

- **Performance.** Often, a design calls out performance at a high level or in a vague way. For example, the design might call for supporting 10,000 users working across 3 time zones. Sometimes, a design might already have sizing ready for you. But you should be familiar with achieving a desired performance level that a design calls out. You need to understand "scale out" (you have 20 servers now, and you add 20 more) versus "scale up" (you have 20 medium-powered servers now and you upgrade them to become high-powered). You also want to understand the basic performance characteristics of hardware and software such as CPU, memory and storage.
- **Security.** You need to ensure that the implementation meets or exceeds the security requirements. For example, a design might call out limiting administrative work to administrative workstations. But how do you accomplish that? Restrict the source IP addresses of SSH and RDP to only the administrative workstation subnets? Restrict the workstations that administrative accounts can logon to? Use just-in-time administration? All of that and more? You must decide.

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 multifactor 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 multifactor 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

Domain 5 Answers to 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 changed 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 audit strategy will depend on its size, industry, financial status and other factors. A small non-profit, a small private company and a small public company will have different requirements and goals for their audit strategies. 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 (e.g., within organization control).** New for the 2024 exam update is a reference to "within organization control". An internal audit strategy should be aligned to the organization's business and day-to-day operations. For example, a publicly traded company might have a more rigorous internal auditing strategy than a small 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 (e.g., outside organization control).** New for the 2024 exam update is a reference to "outside organizational control". 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. For external audits, the auditors often define their own assessment, test, and audit strategies and tactics. For example, they might randomly pick an application in your environment to audit. They might pick a few user transfers to review the transfer accounting processes.

- **Third-party (e.g., outside of enterprise control).** New for the 2024 exam update is a reference to "outside of enterprise control". 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. Third-party audits are often helpful for providing compliance for industry regulations, such as Payment Card Industry Data Security Standard (PCI DSS).

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 (e.g., red, blue, and/or purple team exercises).** New for the 2024 exam update is the call out of red team, blue team, and purple team exercises. 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. As with vulnerability assessments, this 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.

  - **Red team.** The red team is the team that attacks. They have offensive security skills. They are a hands-on team that tries to bypass security controls, evade defenses, and identify security gaps.
  - **Blue team.** The blue team is the team that defends against attacks. Their goal is to design and configure IT systems to thwart attacks, minimize lateral movement, and sometimes respond to attacks (generally a dedicated team).
  - **Purple team.** This team is often a sub team, made up of members of the red team and blue team. Some organizations also have a dedicated purple team. The goal of the purple team is to work with the offensive and defensive sides to optimize organization defenses.

- **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/benchmarks.** While user monitoring captures actual user actions in real time, synthetic — scripted or automated — transactions can be used to test system performance, usability, or security. New for the 2024 exam update is the addition of benchmarks. From the security perspective, benchmarks help you evaluate your environment against industry-specific guidance, frameworks, and vendor recommendations. You can establish a baseline, similar to how you would for performance testing, to look for configuration drift from a security perspective.

- **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 unsecure.

- **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.

- Coverage analysis. For the 2024 exam update, the word test was removed from the title. The content remains as is. 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 (e.g., user interface, network interface, application programming interface (API)).** New for the 2024 exam update is the reference to user interface, network interface, and application programming interface (API). This topic 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 APIs, plug-ins, error handling and more. You should be aware of the different testing types for each system.

- **Breach attack simulations.** As part of your testing, you need a way to simulate real-world attacks. Breach attack simulations provide that by simulating attacks across your entire environment. Often, they are automated and always running. Tools are always being updated with the latest techniques. Breach attack simulation tools reduce the amount of work required for red teams and blue teams. Breach attack simulation tools often provide remediation steps and a lot of documentation for your teams to work with.

- **Compliance checks.** Companies need to perform regular compliance checks to assess whether they are following their controls at that point in time. For example, does that web server respond only on port 443 and not 80? Does that database server have the security patch from last week? Many organizations automate compliance checks using third-party tools or tools developed in house. Results from compliance checks are often sent to teams via email. But a failed compliance check will often result in a ticket or a task so teams can investigate and remediate issues in a timely fashion.

6.3 Collect security process data (e.g., technical and administrative)

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 includes documenting the creation of a user account, as well as 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 when the access method is also prevented? There should also be separate processes for managing the accounts of vendors and other people who might need temporary access.

- **Management review and approval.** Management plays a key role in ensuring that account management 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 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 report

The teams that analyze the security of your apps and services should be aware of how to handle the results and subsequent processes. Any information that is of concern must be reported to the management teams immediately so that they are aware of possible risks or alerts (in some cases, this could lead to immediate remediation while final reports are built). The level of detail given to the management teams might vary depending on their roles and involvement (think "need to know").

- **Remediation.** Teams will remediate deficiencies found during security testing once they become aware of them. For example, if a remote code execution vulnerability is found, the appropriate team should be notified immediately so that the remediation process (even if it is a temporary workaround) can be prioritized. Often, remediation is categorized by risk and priority. For example, an information disclosure vulnerability might be labeled as low risk because it is available only on the corporate network, while a buffer overflow might be labeled as high risk with high priority since it is available to the internet.

- **Exception handling.** When an app or service has an error or exception (the most common are errors on websites), it should generate a message or perform some other tasks so that the user can continue or restart what they were doing. By default, some apps or services will disclose too much information, such as internal configuration details, internal IP addresses, hostnames, software version numbers and lines of code, which could be useful for attackers.

- **Ethical disclosure.** Sometimes referred to as "responsible disclosure," this type of disclosure keeps the full details of a vulnerability private, known only by those that need to know. For example, imagine that you find a vulnerability in a cloud database service. You report the information to the vendor but do not disclose the details to the public. That's ethical disclosure. If you release everything you know to everybody, that's "full disclosure." Full disclosure is often frowned upon because it can elevate risks to many organizations (especially when a fix or workaround is not available).

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 (e.g., within organization control).** 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 (e.g., outside organization control).** 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 (e.g., outside of enterprise control).** 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 a check and balance of the internal and external audits or to perform a more in-depth auditing procedure.

- **Location (e.g., on-premise, cloud, hybrid).** This is a new sub-topic from the 2024 exam update. When doing security audits, it's important to think about the location of the systems and data — whether they're on-premise, in the cloud, or in a hybrid setup (a mix of both). On-premise means everything is stored and managed in your own building, so you have full control but also full responsibility for security. Cloud means you're using someone else's servers, like from Amazon or Microsoft, and while they handle some of the security, you still need to make sure your data is protected. A hybrid setup mixes both, which can be helpful but also adds more to check. From an audit perspective, some vendors don't allow for customer audits. Some do but have specific processes in place for the details.

Note that 6.1 "Design and validate assessment, test, and audit strategies"also addresses location in a similar context.

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 email 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

Domain 6 Answers to 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 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 comply with 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 the 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. For all interviews, consider recording them, if that's legal.

- **Reporting and documentation.** 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 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 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.

- **Artifacts (e.g., data, computer, network, mobile device).** The 2024 exam update added data to the list of artifacts). In an investigation, artifacts are things that are left behind that sometimes leave a trail of what happened and when it happened. Artifacts are very important in an investigation and should be preserved throughout the investigation. You can find artifacts on all of the following:

  - **Data.** Data is critical in an investigation. It is important to leave data where it was during an incident and not alter the hardware, software, or configuration. Much of the investigation relies on the data being in the same state as during an incident. Copying off data to somewhere else sometimes loses critical information and could impact an investigation or how a court of law sees the evidence. It is important to have a chain of custody (documented evidence of anybody who handled the evidence) to provide integrity of the data. To make a copy without impacting the data evidence, you need to make a forensic copy with special tools.
  - **Computers.** On a computer, you can find artifacts such as URLs, file hashes, file names, remnants in the Windows registry and IP addresses.
  - **Network devices.** On a network device, you can find artifacts like irregular outbound network traffic, DNS requests, firewall rule hits, unusual traffic patterns, unusual numbers of requests and unexpected software updates.
  - **Mobile devices.** On a smartphone, you can find artifacts similar to those a computer, such as URLs and IP addresses. You can also find unexpected app installations, unexpected running processes and unusual files.

7.2 Conduct logging and monitoring activities

This section covers logging and monitoring.

- **Intrusion detection and prevention system (IDPS).** For the 2024 exam update, the title was slightly updated by adding the word 'system' and the acronym 'IDPS'. 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, often referred to as an IDPS.

  - An **intrusion detection system (IDS)** is 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 inline 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 it can quickly become hard to manage and work with. Many companies deploy a SIEM solution to centralize the log data and make it simpler to work with. For example, suppose you were looking for failed logon attempts on web servers. You could individually look through the logs on each web server. But if you have a SIEM solution, you can go to a portal and look across all web servers with a single query. A SIEM is a critical technology in large and security- conscious organizations.

- **Continuous monitoring and tuning.** Continuous monitoring is the process of having monitoring information continuously streamed in real time (or close to real time). Such information presents the current environment's risk and information related to the security of the computing environment. Some SIEM solutions offering continuous monitoring or 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 are related to 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.

- **Log management.** At its most basic, log management is the organization and lifecycle of logs. You don't want to keep logs forever and you don't want to fill up the available disk space on endpoints with logs. Often, log management is handled by way of a third-party tool or service. Sometimes, there is a dedicated solution for log management (such as a tool that maintains the logs locally) on endpoints. And often, an endpoint log management system is combined with a SIEM that provides a centralized view and reporting. You need to ensure that logs are kept around long enough to be ingested by your SIEM and long enough to enable administrators to troubleshoot issues.

- **Threat intelligence (e.g., threat feeds, threat hunting).** Threat intelligence is threat data that you can use to minimize threats. Two examples are:

  - **Threat feeds.** These feeds can include specific vulnerabilities being targeted in the wild, new malicious IP addresses, and new malicious URLs and websites. Feeds can also include things like your organization's name, IP addresses or email addresses appearing for sale or in a dump on the dark web.
  - **Threat hunting.** Threat hunting is the act of scouring your environment for threats. Often, organizations have dedicated threat hunters whose job is to look for unusual activity in your environment. If a threat hunter identifies something interesting, it can lead to a formal investigation and subsequent response. Threat hunting relies heavily on data and intelligence, which might lead to a threat hunt in a specific area or on a specific endpoint.

- **User and entity behavior analytics (UEBA).** UEBA is a newer area of IT security. It looks at user and system behavior to establish baselines of (what's normal) and detect anomalous behavior. UEBA can look at things like the amount of network traffic used for an endpoint, the times a user typically works and the location a user works from. It's similar to how a credit card company detects fraudulent use of a credit card: If you typically use your credit card only on Mondays and for only $50 at a time, your card being used on a Wednesday for $1,500 will certainly trigger an alert.

7.3 Perform configuration management (CM) (e.g., provisioning, baselining, automation)

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. You can configure it so that most changes to the system by a user result in automatic remediation by the configuration management 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. The three examples below are important to know for the exam:

- **Provisioning.** With configuration management, you can automate the provisioning process. As new computers arrive, your configuration management tool can automatically install the organization's operating system image, including the latest security updates. Automatic provisioning helps to ensure your computing environment is consistently configured while saving on administrative overhead.

- **Baselining.** Baselining is the act of capturing the normal configuration, which enables you to detect anomalies with a server or even with the network (systems are not configured to the baseline). Baselining is often used with configuration management to remediate configuration drift. But it can be extended to include performance data for a typical day of use in order to spot performance issues.

- **Automation.** While configuration management enables you to capture and maintain configurations, it has limits. To extend the capabilities of configuration management, you can add automation. Sometimes, this will require another product (such as an automation product that integrates with your configuration management software). With automation, the options are almost limitless. For example, if configuration management detects a service stopped, it can start it, turn on verbose logging and send an email to the administrator.

7.4 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. These are both recommended practices. 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 an example from outside of IT: If Bob trusts Jon and Larry trusts Jon, should Bob trust Larry? Probably not. In high-security environments, it isn't uncommon to see non-transitive trusts used, depending on the configuration and requirements.

- **Segregation of duties (SoD) and responsibilities.** New for the 2024 exam update is the change from 'separation of duties' to 'segregation of duties'. Segregation of duties refers to the process of separating certain tasks and operations so that a single person doesn't control all of 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 or segregation of duties is to make it more difficult to cause harm to the organization (via destructive actions or data loss, for example). With segregation 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.

- **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 an SLA 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.

7.5 Apply resource protection

This section covers media. When we think of media, we think of hardware that can store data — hard drives, USB drives, optical media, backup tapes, etc.

- **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.

- **Media protection techniques.** Media protection is the act of ensuring that media doesn't get damaged or destroyed (physically or virtually). When you think about virtual media protection, think of techniques such as "write once, read many" (WORM), whereby media can be written just once but not overwritten or tampered with thereafter. When you think about physical media protection, you should think about data centers (secure facilities), locked cabinets, security cameras and other physical security.

- **Data at rest/data in transit.** This is a new topic that was introduced with the 2024 exam update. Data at rest and data in transit are two types of data that need protection. Data at rest is information that's stored somewhere, like on a hard drive or in the cloud — it's not moving. Data in transit is data that's traveling, like when you send an email or copy files from a source to a destination. Both types need to be kept safe from hackers or leaks. One common way to protect them is by using encryption, which turns the data into secret code that only the right people can read. A second way is using information rights management, which is embedded in the document and can prevent unauthorized access to data, even if the person obtains the data.

7.6 Conduct incident management

Incident management is the management of incidents that are potentially damaging to an organization, such as a distributed denial of service (DDoS) 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 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.

7.7 Operate and maintain detection 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. Note that during the 2024 exam update, the title of this section was changed from 'detective' to 'detection' but the meaning remains the same.

- **Firewalls (e.g., next generation, web application, network).** While operating firewalls often involves adding and editing rules and reviewing logs, there are other tasks that are important, too. For example, you should review the firewall configuration change log to see which configuration settings have been changed recently. Know about these 3 examples of different types of firewalls:

  - **Next generation.** The latest firewalls often bring new features to the traditional role of the firewall. Features such as deep packet inspection, application-level inspection, third-party security data evaluation and IDS/IPS features (described below). Firewalls offering these features are often referred to as "next generation" firewalls.
  - **Web application.** A web application firewall (WAF) is designed to work with HTTP and HTTPS (web applications / websites). It protects not just users but also apps. Think of it like a reverse proxy that sits in front of an application. It can block bad things coming from the user before they get to the app.

- **Network.** A network firewall is a traditional firewall used to protect networks (such as corporate networks) from unauthorized access (such as from the internet). A network firewall is also used to segment an organization's network to minimize or eliminate some types of network traffic. Traditional firewalls look at source IP address, destination IP addresses, ports and protocols and allow or deny the traffic based on a pre-defined list of rules.

- **Intrusion detection systems (IDS) and intrusion prevention systems (IPS).** You need to routinely evaluate the effectiveness of your IDS and IPS systems. You also need to review and enhance the alerting functionality. If too many alerts are sent (especially false positives 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.

- **Third-party security services.** 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 forensic services. 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 / 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 uses 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.

- **Machine learning and artificial intelligence (AI) based tools.** Traditionally, many security solutions relied solely on rules, signatures or a defined configuration to know what to do (allow, deny, quarantine or something else). With machine learning and AI, new tools are able to rely on vast amounts of data (from your own network and other networks) before deciding what to do. Machine learning and AI can look for patterns in data, events and actions and perform deep analysis to spot malicious events.

7.8 Implement and support patch and vulnerability management

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

- **Patch management.** When a vendor provides an update to their software to fix security issues or other bugs, the update is often referred to as a patch. 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. Some experts claim 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 this process so that you know 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 have to figure it 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.

7.9 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.

7.10 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 (e.g., cloud storage, onsite, offsite).** For the 2024 exam update, the examples of cloud storage, onsite storage, and offsite storage were added to the tile). 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. Onsite storage is often a good place to store backups, as long as it isn't the only place. Offsite storage can be in the cloud or can be at another location (a remote office or co-location facility). 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 and cloud 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 and cloud providers offer 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 (e.g., cold vs. hot, resource capacity agreements).** The 2024 exam updated added the examples of cold vs. hot and resource capacity agreements. 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/hot standby site). A warm site has some equipment, but requires some work to fail over to it. A hot site is a fully operational mirror of your production site and provides very fast recovery (but is also the most expensive). 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 (HA), quality of service (QoS) and fault tolerance.** To prepare for the exam, it is important to be able to differentiate between these subtopics:

  - **System resilience.** Resilience is the ability to recover quickly. With site resilience, 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 (HA).** While site resilience is about recovery 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. Many organizations want both high availability and site resiliency. While clusters are often the answer for high availability, there are many other methods available too. For example, you can provide a highly available web application by using multiple web servers without a cluster.

  - **Quality of service (QoS).** Certain services might require a higher quality of service than other services. For example, an organization 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 make its own websites 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 in order to provide fault tolerance. Fault tolerance by itself isn't valuable. For example, a server might have fault-tolerant CPUs, but if the server's power supply fails, the server is done. Therefore, you must account for fault tolerance across your entire system and across your entire network.

7.11 Implement disaster recovery (DR) processes

Establishing clear disaster recovery processes helps minimize the effort and time required to recover from a disaster. Testing the plans is also important; it is discussed separately in the next section (7.12).

- **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.

- **Communications (e.g., methods).** For the 2024 exam update, the title was updated to refer to communication methods. In a real disaster, communicating with everybody will be difficult and, in some cases, impossible. Sometimes, companies use communication services or software to facilitate emergency company-wide communications or mass communications with personnel involved in the disaster recovery operation. 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 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 smartphones (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, who 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. Examples include 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 and 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 in section 7.12.

- **Lessons learned.** After a real disaster, you should plan to go over the lessons learned — the things that happened that you didn't plan for (or plan for correctly), didn't account for or didn't anticipate and had to take corrective actions for during the event. The goal is to revise the DR processes so that the teams do not deal with the same challenges during the next disaster. Of course, lessons learned can also be valuable for DR tests.

7.12 Test disaster recovery plans (DRPs)

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 step through the recovery operations to uncover 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 to perform a real-world recovery operation. 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 downsides are serious, full interruption tests are a good practice for most organizations.

- **Communications.** This is a new topic introduced in the 2024 exam update. Communication with business stakeholders, including the current status, and industry or government regulators is important. It's not enough to just have a disaster recovery plan; people need to understand it and know their roles. That's why it's important to test the plan and talk about it clearly with the team. This includes explaining what will happen during the test, who's doing what, and how success will be measured.

7.13 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 recover from a disaster. Business continuity focuses on a business operating with minimal or no downtime (with a hope that a disaster recovery process won't be needed). Think of business continuity as a strategy and disaster recovery as a tactic. In the bullets below, the steps required to plan business continuity are shown. 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 the communication plan.** Finalize who needs to be contacted, figure out primary and alternative contact methods, and ensure that everybody can work, possibly from backup locations.

- **Coordinate with external entities.** Communicate with external entities such as the police department, government agencies, partner companies and the community.

7.14 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, you should restrict who can enter your facilities. 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 (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.** For internal security, we are focused 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. 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 to keep sensitive information locked away.

7.15 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. 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. Keep 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 (e.g., insider threat, social media impacts, two-factor authentication (2FA) fatigue).** For the 2024 exam update, the topic was updated with examples of types of 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 people. 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. Introducing additional security measures, such as two-factor authentication, can lead to 2FA fatigue if users are prompted too frequently. Users should be made aware of the latest threats in their training, including the use of AI, MFA fatigue (2FA fatigue). In addition, your training and awareness campaigns should account for an insider threat (current co-worker) vs. just talking about anonymous hackers.

- **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 brandishing 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?

   a. The company has a zero-day vulnerability.
   b. The company was compromised by a zero-day exploit.
   c. The computer does not have a configuration management agent.
   d. 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 option 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

Domain 7 Answers to 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 that 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 in the software development lifecycle

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 (e.g., Agile, Waterfall, DevOps, DevSecOps, Scaled Agile Framework).** For the 2024 exam update, this topic was updated to mention the Scaled Agile Framework. 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

Methodology	Key Concepts
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
Scaled Agile Framework	▪ SAFe helps big teams work together by using agile methods across the whole company, not just in small groups — it's like getting many teams on the same page so they can move faster and smarter. ▪ It focuses on planning, teamwork, and quick updates, making it easier to deliver products or changes in small steps instead of waiting months for one big release. ▪ Security needs to be part of the process, so in SAFe, things like risk management and secure coding are added early and often — not just at the end. This helps catch problems sooner. The industry sometimes refers to this as "shifting left".
DevOps	▪ DevOps combines development and operations teams so they can work better together, build software faster, and fix problems quickly. ▪ Automation is a big part of DevOps, which means using tools to do tasks like testing, building, and deploying code without having to do it all by hand. ▪ Security should be built into every step, not just added at the end — this is often called DevSecOps, where security is part of the team from the start.

Methodology	Key Concepts
DevSecOps	▪ DevSecOps stands for Development, Security, and Operations, and it means building security into every part of the software process — not just at the end. ▪ Everyone shares responsibility for security, including developers, security teams, and IT — they all work together to find and fix issues early. ▪ Security tools are used automatically during coding, testing, and deployment to check for things like bugs, weaknesses, or risky settings, helping to keep software safe from the start.

- **Maturity models.** There are five maturity levels in 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 that a change is requested. Each change request must be evaluated for capability, risk and security concerns, impacts to the timeline, and more.

- **Integrated product team (IPT).** 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 software development ecosystems

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. Several of the bulleted items below are things you should understand before taking the exam.

- **Programming languages.** 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. **Generation 1:** Machine language. This is the binary representation that is understood and used by the computer processor.
  2. **Generation 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. **Generation 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. **Generation 4:** Very high-level language. Generation 4 languages aim to simplify the previous generation languages by reducing the amount of code that is required. This allows the language and code to focus on algorithms that are used for specific programming instruction. Python, C++, C# and Java are examples of Generation 4 programming languages.
  5. **Generation 5:** Natural language. Natural language is the ability for a system to learn and change on its own, as with artificial intelligence. Instead of code being developed with a specific purpose or goal, only the constraints and goal are defined; the application then solves the problem on its own based on this information. Prolog and Mercury are examples of Generation 5 programming languages.

- **Libraries.** Libraries are typically third-party pre-developed code that handle common functions or features. Often, they are free and/or open source. Many developers use third-party libraries to speed up the development process (since they don't have to write code to handle the function or feature). One security challenge with libraries is keeping track of updates and reported security issues. Often, libraries are used but never updated, which leaves organizations vulnerable. Some organizations use internal repositories and limit or block downloads from the internet. Configuration management can also help track down usage.

- **Tool sets.** Organizations provide developers with tools to enable them to create, debug and maintain code. Tools are included with software development kits, suites and tool sets. Sometimes, tools are standalone. As with other applications and services, you need to watch carefully for security patches, bug fixes and version updates. Developers are often focused on their code (and the security of their code) but the tools behind the code remain an important part of securing a software development environment.

- **Integrated development environment (IDE).** An IDE is itself an application (software) that integrates core developer tools in a single platform (often a GUI). For example, there might be a source code editor (text editor), a debugger and a build tool in an IDE. Some IDEs offer features such as automatic code completion, real-time syntax and bug identification, and automated compiling. Many IDEs offer plugins and extensions, which like, browser plugins and extensions, post risk, so some organizations limit or prevent the use of unapproved plugins and extensions.

- **Runtime.** Runtime begins when a program or app is launched and it continues until the program or app is closed. It is difficult to stop runtime attacks because of the complexity of software. Some vendors offer a runtime application self-protection (RASP) solutions that intercept calls made by a program to the system to ensure the calls are safe before allowing them to proceed.

- **Continuous integration and continuous delivery (CI/CD).** Sometimes referred to as the CI/CD pipeline, this development methodology enables developers to deliver code changes more often and with more reliability. It is categorized as an agile development methodology. The premise is to introduce small changes and introduce them frequently. Often, organizations use automation to build, package and test. This helps to enhance consistency and quality. The delivery part of this methodology automates the code change pushes to various environments, such as QA, test and production.

- **Software configuration management (CM).** 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 have been made. The repository can track versions of an application so that you can easily roll back to a previous version if necessary.

- **Code repositories.** The version control system that houses source code and intellectual property is the code repository. A development lifecycle might have 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 to access the source code.

- **Application security testing (e.g., static application security testing (SAST), dynamic application security testing (DAST), software composition analysis, Interactive Application Security Test (IAST)).** For the 2024 exam update, the title added examples of software composition analysis and Interactive Application Security Test (IAST). Testing the security of your applications is an important step to minimizing the risk of vulnerabilities and other security issues. The exam focuses on two types:

- **Static application security testing (SAST).** SAST is used to automate manual code reviews. Instead of having engineers manually look through lines of code, scanning software and/or services scan it at a rapid pace to identify vulnerabilities and mitigate them. SAST solutions excel at finding vulnerabilities in the code (especially well-known code issues) but are not suited for finding issues outside of the code. SAST solutions will not find all vulnerabilities and issues. Some organizations use multiple tools or combine SAST with other methods (such as DAST).
- **Dynamic application security testing (DAST).** DAST is used to find vulnerabilities in an application while it is running (which is much different than scanning static code like SAST). Therefore, DAST tools can be used only after an application is developed and able to run, while SAST tools can be used earlier in the development process. DAST tends to be less expensive and can find issues outside of the application itself. Many organizations combine DAST, SAST and additional methods.
- **Software composition analysis.** Software Composition Analysis (SCA) is a way to check what parts make up a software program, especially the pieces that come from other sources like open-source libraries. Think of it like checking the ingredients label on a food package — you want to know what's inside and if anything could be harmful. Some of these outside parts might have security issues or licenses that cause problems. SCA tools help find and keep track of these parts so developers can fix or update them before they cause trouble.
- **Interactive Application Security Testing (IAST).** IAST is a way to check if an app has security problems while it's actually running. It watches the app from the inside, kind of like a coach on the field giving feedback during the game. IAST tools are installed in the app and work while testers are using it, so they can spot real issues as they happen — like broken code, unsafe data use, or things that could lead to hacking. It gives more accurate results than just scanning from the outside.

This topic was removed from this specific section of the exam, however, is still covered elsewhere on the exam.

- **Security orchestration, automation and response (SOAR).** SOAR is a collection of software that helps security teams automate routine operational tasks such as scanning for vulnerabilities and looking through log files. It also includes integrating various security tools together to streamline operations and facilitate automation. Organizations choose to implement SOAR to reduce administrative overhead, as well as to enhance the consistency and timeliness of the work.

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 meets 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 acquires another organization, the other organization's 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. Finally, you need to understand the impacts of the software. During an acquisition, your organization might be taking on hundreds of new applications (and their associated risks).

- **Commercial-off-the-shelf (COTS).** Boxed software or software that you download and install is sometimes referred to as COTS. Some of the same software security principles that apply to securing software developed in house apply to COTS. The primary difference is that you typically won't have access to the code base and libraries, so you will be limited to dynamic testing and will have to rely on the vendor for security patches.

- **Open source.** Open source software, by definition, has everything available to the public, which makes it easier to perform the same types of security scans as you can with in-house applications. The primary difference is that you can't always fix things, especially at the pace you might with your own code.

- **Third-party.** Third-party software typically refers to software that y is used in your network but that you don't own, install or manage it. For example, if you have a vendor come in to perform a wireless network upgrade, they might use several applications and tools that you don't own or manage. But those tools are on your network and could have vulnerabilities. Some organizations choose to require vendors to use organization computing devices, such as VDIs. This helps limits exposure to unknown software.

- **Managed services (e.g., enterprise applications).** For the 2024 exam updater, the title removes a reference to SaaS and PaaS and instead calls out 'enterprise applications'. The SaaS and PaaS were moved to the next section. Software from managed services is software that you use but don't own, install or maintain. Often, you might configure it. From a security perspective, you are limited in what you can do (based on what type of security testing the vendor allows) and you are limited in your visibility (for example, you won't have access to the codebase). Many of the large service providers publish their security audit information for you to review. And many allow for specific types of security review and/or penetration testing.

- **Cloud services (e.g., Software as a Service (SaaS), Infrastructure as a Service (IaaS), Platform as a Service (PaaS)).** When using cloud services like SaaS, PaaS, and IaaS, it's important to think about how they affect security. Each type of cloud service gives you different levels of control. With SaaS (Software as a Service), like Google Docs or Microsoft 365, the provider handles most of the security. With PaaS (Platform as a Service), you manage your apps, but the provider manages the system underneath. With IaaS (Infrastructure as a Service), like Amazon EC2, you get more control, but also more responsibility to protect things like your servers and data.

8.5 Define and apply secure coding guidelines and standards

In many organizations, there is a defined security strategy, along with policies, procedures and guidelines. Often, the strategy is focused at the infrastructure level and deals with network security, identity and access management, and endpoint security. However, having a security strategy should also extend to coding if the organization develops code internally. Some of the strategy is focused at the security of the code, while some of it is dedicated to standards, such as the preferred programming language in defined 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, such as 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 all 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 (APIs).** 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. Security of APIs starts with authentication. Authentication techniques such as OAuth or API keys should be used. Additionally, authorization should be used. For example, with one API key, you might be able to read information, but you need a separate API key to alter or write information. Many companies use an API gateway or an API security gateway that centralizes API calls and performs checks on the calls (checks tokens, parameters, messages, etc.) to ensure they meet your organization's requirements. Other common methods to secure your APIs is to use throttling (protects against DoS and similar misuse), scan your APIs for weaknesses, and use encryption (such as with an API gateway).

- **Secure coding practices.** When coding, there are established practices you should follow to maximize the security of your code. Some of the most common practices and areas are:

  - **Input validation.** Validate input, especially from untrusted sources. Any validation failures should result in a rejection of the input.
  - **Don't ignore compiler warnings.** When compiling code, ensure that you are receiving the highest warning level available.
  - **Deny by default.** By default, everybody should be denied access. Access is granted as needed.
  - **Authentication and password management.** Require authentication for everything, unless something is meant to be available to the public. Hash passwords and salt the hashes.
  - **Access control.** Restrict access using the principle of least privilege; deny access if there are issues checking access control systems.
  - **Cryptographic practices.** Protect secrets and master keys. Establish policies and procedures and 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.

- **Software-defined security.** Whenever you see "software- defined" in front of something, it typically means that software and automation are in control. In this case, software and the associated automation controls the security of your applications and services. Often, this control happens based on defined policies. This area grew out of software-defined networking (a precursor to software-defined everything else). Software-defined security can manage everything from provisioning to firewall rules to ports used on a web server.

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

Domain 8 Answers to Review Questions

1. Answer: A

Explanation: Agile development emphasizes efficiency and iterations during the development process, and 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

Webinars

How to Protect Your Organisation Against Lateral Movement Attacks

Learn from Industry Experts: How Businesses Like Yours Can Protect Sensitive Data

Behind the Scenes: 4 Ways Your Organization Can Be Hacked

Top 5 Things to Do to Stop Attackers in Their Tracks

Pro Tips for Defending Your Organization from Data Breaches

Securing Your Network Devices in the Era of Cyber Threats

[Deep Dive] Force IT Risks to the Surface

Withstanding a Ransomware Attack: A Step-by-Step Guide

Best Practices

Privileged Access Management Best Practices

Data Security Best Practices

Data Security and Protection Policy Template

Data Classification Policy Example

Best Practices: How to Harden Privileged Account Security

Windows Server Hardening Checklist

Information Security Risk Assessment Checklist

How to Prevent Ransomware Infections: Best Practices

Best Practices: How to Minimize the Risk of Insider Threats

Best Practices: How to Implement Audit Policy

eBooks

Addressing Modern Cybersecurity Challenges through Enterprise-Wide Visibility

To SIEM or Not to SIEM: Is there a better way to secure your data?

10 Questions for Assessing Data Security in the Enterprise

Insider Threat Playbook: How to Deter Data Theft by Departing Employees

Defending Against Crypto-Ransomware

Reduce Your Risk of a Data Breach by Extending Visibility Beyond SIEM

Blogposts

Data Security Explained: Challenges and Solutions

What Is Privileged Access Management (PAM)?

Understanding Insider Threats: Definition and Examples

What Is Security Information and Event Management (SIEM)?

10 Security Tips for Malware Prevention

What to Know about a Data Breach: Definition, Types, Risk Factors and Prevention Measures

Top 5 Human Errors that Impact Data Security

Must-Have Data Security Controls

Cybersecurity Assessment: Definition and Types

Risk Analysis Example: How to Evaluate Risks

Five Reasons to Ditch Manual Data Classification Methods

How to Build an Effective Data Classification Policy for Better Information Security

A Perfect Storm in Cybersecurity

Choosing the Right Security Certifications: CISSP vs CISM, CISA and CRISC

Expert Advice: Is CISSP Worth It?

Blogposts

[(ISC)²: ISC2 Certification Guide]

[The Importance of Certification Training]

[How to Choose Between CompTIA, AWS, and ISC2 Certifications]

About Chauster

Chauster UpSkilling Solutions is a leader in technology education, delivering dynamic, real-world IT training designed to meet the demands of today's fast-paced digital workforce. We offer an extensive range of computer training programs that span mission-critical domains such as cloud computing, networking, cybersecurity, and DevOps—ensuring learners at every level gain the specialized knowledge needed to thrive in modern tech environments.

What sets Chauster apart is our device-based learning model, which goes beyond conventional web-based platforms and outdated classroom formats. By integrating hands-on, mobile-enabled, and accountability-driven training, we empower learners to engage with the material in meaningful, personalized ways. This approach ensures better retention, greater flexibility, and immediate application of skills in professional settings.

Our programs are carefully curated to support not just knowledge acquisition but career transformation. Whether you're a job seeker looking to break into IT, a seasoned professional aiming to upskill, or an employer building a high-performance team, Chauster delivers practical, industry-relevant solutions that move people forward.

At Chauster, we don't just teach IT—we build confidence, open doors, and drive innovation. We're proud to equip our students with the tools, certifications, and confidence they need to succeed in roles that shape the digital world.

Learn more about our approach and impact in our featured blog:

Chauster UpSkilling Solutions Practical Training for Real-World Careers