Domain 1

1. Answer: C

Trusted Recovery is required in high-security systems and allows a system to terminate its processes in a secure manner. If a system crashes, it must restart in a secure mode in which no further compromise of system policy can occur.

2. Answer A

Open design is often thought to be better than closed design, as the openness allows for review from others in the community. The idea is that if others have access to the code, they will help examine and review the code, and ultimately improve it. That was not the case unfortunately with OpenSSL. The point being that it is not necessarily that open source is more secure. If the code is not reviewed, it might as well be closed source. Also, ultimately the quality of the code dictates the security, much more so than whether it is open or closed.

3. Answer C

Dual Control is a security principle that requires multiple parties to be present for a task that might have severe security implications. In this instance, it is likely best to have at least two network administrators present before a private key can be recovered. A subset of dual control is called M of N control. M and N are variables, but this control requires M out of a total of N administrators to be present to recover a key.

4. Answer A

Project Initiation is traditionally the phase in which senior management pledges its support for the project. Often in this phase, management provides a project charter, which is a formal written document in which the project is officially authorized, a project manager is selected and named, and management makes a commitment to support.

5. Answer B

Before any work should be done on a Business Continuity Policy, there must be a BCP policy signed by management. Without one, the BCP Coordinator/Project Manager will not know management's objectives, scope, and level of commitment. The policy will also include management's degree of support and funding for the project. Without this information and commitment, the project is doomed from the start.

6. Answer D

Some organizations group risk analysis into the process of conducting the Business Impact Analysis, while others consider it a separate function. The purpose of the BIA is to identify business processes and prioritize them based on criticality. After this step, risk analysis would identify the threats (and their likelihood) that could compromise those business processes.

7. Answer B

Though senior management is responsible for testing the plan, they cannot be expected to be involved in testing the technology that will implement the plan. Functional managers or department heads will oversee the technical systems that will achieve the overall goals that senior management has laid out. For instance, senior management may set a goal of data recovery within an hour but is up to the head of the department to ensure that the company's backup/recovery strategy can meet those goals.

8. Answer D

Senior Management (or possibly the BCP coordinator, if specified in the plan) should fulfill the responsibility of declaring a disaster. The plan should explicitly define the characteristics of a disaster, and senior management should determine if the current environment meets that criterion. If so, then senior management should begin phase one of the plans, which is to notify employees.

9. Answer B

Arguably, the BCP committee's most important function is to conduct the Business Impact Analysis. This document is the point from which all other plans will begin. The BIA will specify the metrics and objects to be met as a result of the Disaster Recovery Plan, as well as others.

10. Answer B

The Salvage Team is responsible for reconstitution (also known as failback) to a state of permanence. Reconstitution will require restoration of LEAST critical services first, ultimately leading to the full restoration of operations at the permanent facility. Only after reconstitution is a disaster considered to be over.

11. Answer A

The Occupant Emergency Plan will detail how employees are to evacuate a facility and reach a safe environment. It will often include how to assist those with limitations, assign responsibility for activities such as ensuring all members have reached safety as well as include evacuation and backup routes.

12. Answer C

The COOP is responsible for enabling the long-term (relatively speaking) operations after a disaster. Rescue plans address the protection of human life and property in the immediacy of the disaster. The recovery phase deals with restoring critical operations as quickly as possible. The COOP begins after operations have been restored and is designed to provide guidance on running the organization until full operations can be resumed.

13. Answer D

The Disaster Recovery Plan is usually focused on restoring IT services based on their criticality. The DRP's counterpart that addresses business processes is called the Business Recovery Plan.

14. Answer D

The BCP should be distributed based on a need-to-know basis. The entire plan may contain sensitive information and plans about how to respond to security breaches and how to protect against them. This is not information that should be distributed indiscriminately. Individuals are granted access to the portion of the plan that is relevant to them. Most users are only given information about how to safely evacuate the building and any necessary steps following the evacuation.

15. Answer D

Most industry experts indicate that an annual review of the BCP is necessary to ensure the information contained within is current. Also, in the event of a major change, like acquisition or merger with another organization, a review is necessary.

16. Answer B

Redundancy is an important principle that provides high availability. Because of the inherent importance of Disaster Recovery and Contingency plans, copies should be kept at multiple locations and should be stored digitally and as a hard copy.

17. Answer D

A test in which an offsite facility is activated, and a portion of operations are performed at this offsite facility is called a parallel test. It is riskier than paper-based tests because if the alternate facility isn't properly operational, a portion of operations can be lost. It is, however, less risky than a full- interruption test in which all operations are ceased at the normal facility, and resumed at the alternate facility.

18. Answer C

The BIA will determine metrics such as MTD (Maximum Tolerable Downtime) which defines how quickly a service or data should be restored. RPO (Recovery Point Objective) will dictate how current data must be. These pieces of information will determine what controls will be put in place. For instance, if an organization needs to be able to provide data that is current within one hour, but only conducts daily backups for redundancy, it will be impossible to guarantee the RPO. Nightly backups have a possible loss of a day's worth of data (Systems could fail at 4:59 and we would only have last night's backup to use for recovery.)

19. Answer B

The DRP has three phases: Notification, Recovery, and Reconstitution. The recovery phase of the DRP should address the function and recovery of critical operations, often at other locations. These locations can include an offsite facility (hot, warm or cold site) that the organization uses to restore operations. It also, however, can describe an environment in which employees perform their operations from home (or elsewhere), usually for very limited periods of time, and not for long-term disasters.

Domain 2

1. Answer C

To significantly mitigate risks on the network, we have to implement security that limits connectivity to our network from external devices. Additionally, we are concerned with monitoring software being installed on our hosts, so we want to limit the ability of such software to be installed. Further, we want to ensure that other basic security requirements are satisfied, such as using strong passwords, lockout policies on systems, physical security, etc. Remember: Proactive devices PREVENT an attack, as opposed to responding to it. Network scans often detect these devices, but they rarely prevent.

2. Answer D

Separation of Duties is frequently used to limit the amount of information to which any one individual has access. For instance, a user cannot likely leak the password for a file server because that information is exclusively available for those for whom their jobs require that information. Separation of duties frequently goes hand-in-hand with need-to-know and the principle of least privilege

3. Answer A

Though B, C, and D may be part of what is detailed with the various levels of classification, the primary purpose of classification is to ensure that the appropriate controls are implemented to provide adequate and consistent security for the resource.

4. Answer B

One of the greatest benefits of configuration management is that it provides stability for systems on the network, as well as the network itself. Without a means of evaluating, controlling and documenting proposed changes, changes could be made at will. Often changes that seem like a good idea at first may have a long-term effect on systems and may have unanticipated results. Also, users frequently don't understand the functional and security ramification of application installation or modification of settings.

5. Answer C

Organizations that practice good configuration managements should have a well-documented policy on the change control process. Part of the policy should include the emergency change control process. Even if a lead technician or manager authorize a change, the change should still be presented to the Change Control Board through the emergency change control process.

6. Answer D

An organization's patch management strategy should include how to handle security-related patches, often with an expedited process. Never take it upon yourself to implement a patch, regardless of the reason. Patches may occasionally have an adverse reaction to systems, which is why there should be a well-documented policy.

7. Answer B

The best way to protect data is to encrypt it. Though a cable lock would indeed help prevent a laptop from being stolen, without encryption the data can still be compromised. Monitoring and the review of audit logs will probably not reveal access to sensitive information, and even if they did, the logs would only indicate that data had been accessed, and would not prevent that access.

8. Answer C

The TPM (Trusted Platform Module) chip is hardware contained on the motherboard originally designed for the limited purpose of hard drive encryption. Vendors today are frequently using this chip for other purposes, such as using it a location to store activation information in an attempt to prevent privacy.

9. Answer B

SSH is a secure protocol for remote administration. Additionally, it can be used to transfer files through the use of S/FTP. S/FTP is the SSH protocol with an FTP shell so that users experienced in FTP can use the commands with which they are familiar.

Domain 3 - Section 1

1. Answer B

The TCB (Trusted Computer Base) describes the elements of a system which enforce the security policy and are used to determine the security capabilities of a system. This term was coined by the Orange Book (Also known as the Trusted Computer System evaluation criteria.) Some components included in the TCB are the system BIOS, the CPU, Memory, the OS kernel.

2. Answer C

As a subject attempts to access an object, two of the main elements that control access are the Reference Monitor and the Security Kernel. The Reference Monitor is the conceptual rule set that defines access while the Security Kernel includes the hardware, software, or firmware that enforces the rules set.

3. Answer A

There is always a trade-off for security. Sometimes the cost comes in actual dollars spent. Often, other times, security negatively affects performance, backward compatibility and ease of use. An organization must look at the overall objectives of the business considering their primary needs. Whereas systems which house sensitive military information must be designed with much more security than a small home/office environment that has information of little to no value to an attacker. The amount of security that should be implemented should meet the needs of the business, without exceeding the amount of cost the organization is willing to pay.

4. Answer D

Secure by design is one of the most important concepts in system/software development. Often in the past, we have asked two questions: "Does it work?" and "Is it secure?" In following the "secure by design" philosophy, products are not considered functional unless they function securely. Security is addressed at each phase of the SDLC (System Development Lifecycle) including the initial phases which include the practices of risk assessment, functional design and implementation. By including security in each of the phases, we design a product to be secure, as opposed to considering security as an afterthought.

5. Answer C

One of the main benefits of thin clients is that the responsibility is taken off the client for the installation, upgrades, and management of resources. A central computer hosts the software and services and the clients access these services. The client can contain very minimum hardware/software, as the services are actually running on the server, whether it be a local server or a server accessed through a cloud service provider's network.

6. Answer C

A benefit of loose coupling is that the Components in a loosely coupled environment or system can be exchanged with alternative implementations which provide the same services, and are much less constrained by the same language, platform, operating system, or build environment.

7. Answer A

The *_Security Property of the Bell-LaPadula security module is designed to prevent users that have access to higher levels of data access from writing to an area of lower access. For instance, it would prevent a document classified as "top secret" from being written to a folder classified as "secret."

8. Answer A

Startup of a system is difficult to secure, as many protective mechanisms have yet to be loaded. Some of the more successful malware has been designed to load early in the process—perhaps when the kernel or virtual device drivers load to evade detection.

9. Answer B

The Clark-Wilson security models enforce separation of duties. Rather than allow an untrusted entity to have full access, we limit the untrusted entity to a limited access of an interface. The interface would then control and end enforce a well-formed request. The Clark-Wilson model is implemented in many ways in the Information Security world. We use a firewall as an interface between the public internet and our trusted internal network. We use application programming interfaces to allow an application to access the trusted resources it needs.

CYBRARY

Kelly Handerhan's CISSP Preparation Course

Domain 3 - Section 2

1. Answer A

The data owner has the responsibility of determining the classification of data based on pre-defined criteria. The data custodians primary responsibility is to implement the security controls based on the classification and to provide the day-to-day oversight, including ensuring that backups are current.

2. Answer A

For an information system, the potential impact values assigned to the respective security objectives (confidentiality, integrity, availability) shall be the highest values (i.e., high water mark) from among those security categories that have been determined for each type of information resident on the information system.

3. Answer C

Trust is typically defined in terms of the security features, functions, mechanisms, services, procedures, and architectures implemented within a system. Security assurance is the measure of confidence that the security functionality is implemented correctly, operating as intended and producing the desired outcome based on the reliability of the processes used to develop the system.

4. Answer A

A TOC/TOU attacks when an attacker (or a system process) creates a variance between when a resource is verified and when it is used. In this instance, the network operating system has authenticated the user and allowed him access to the domain. The OS continues to use the information learned in the initial check for the user's authentication. The user continues to "Use" the system, as no updated information about the account suspension is passed along. There are numerous instances when this attack can be used, causing multiple issues including privilege escalation.

5. Answer D

The best means of mitigating the threat of resource exhaustion is implementing a means of detecting and limiting access to the resource. Input validation can help ensure that an attacker doesn't input a data value greater than expected. Throttling might include tracking the rate of requests received from users and blocking requests that exceed a defined rate threshold.

6. Answer C

Data resides in storage much longer than it does in transit and must be stored in a secure manner. Encryption of data helpful to enforce confidentiality and protect application data, keys, passwords, etc. However, even when encryption is used, it may not be used properly. Common mistakes include:

• Failure to encrypt sensitive data

• Weak protection for the storage of credentials (keys, certificates, and passwords)

• Improper storage of confidential information in memory/swap files,

• Poor statistical randomness

• Weak cryptographic algorithms.

Domain 3 - Section 3

1. Answer D

Pattern analysis is often the easiest way to crack a pure substitution cipher. For instance, knowing things such as the most commonly used letter of the English alphabet is "e" can lead us to make a reasonable assumption that whatever character most commonly appears is likely substituted for "e." Also, it is estimated that as many as 60% of emails start with the letter "h." The more assumptions we are able to make correctly the quicker we can compromise a substation cipher.

2. Answer D

Sessions keys are used for a single session and are then discarded, as is the one-time pad. Additionally, each session key must be statistically unpredictable and unrelated to the previous key, as the one-time pad requires, as well. Any technology that takes advantage of a short-term password or key can ultimately be traced back to the one-time pad.

3. Answer B

DES was originally the standard used to protect sensitive but unclassified information for the US Government. Once DES was compromised we needed a quick means to increase the security. 3DES literally tripled the length of the key from 56 bits to 168 bits. Often a quick means to strengthen a compromised algorithm is to increase the key length or the length of the initialization vector.

4. Answer B

Non-repudiation is the combination of authenticity and integrity and is implemented through the use of digital signatures.

5. Answer: C

Integrity provides assurance against modification of data, whether malicious or accidental. Though non-repudiation (which includes integrity) would also provide detection that messages have been corrupted, it would also provide the additional security services of authenticity and non-repudiation, which would cause additional overhead.

6. Answer D

Non-repudiation combines integrity (which guarantees the message has not been modified) and authenticity which verifies the origin of the message. Only non-repudiation would meet the above requirements.

7. Answer D

Seeds or salts are added to provide additional randomness to passwords as part of the second layer of defense against password cracking.

8. Answer A

XOR (Exclusive Or) is a process frequently used by stream ciphers to provide bit-by-bit encryption. Typically this type of encryption is very fast and efficient but does not usually provide the same security that block ciphers provide.

9. Answer A

Another term for the key is crypto-variable which indicates that the randomness and variability of the crypto process comes from the key.

10. Answer C

Rijndael was selected by the government to satisfy the Advanced Encryption Standard specified by the government in 2002 and is the default algorithm that many applications use to provide security.

11. Answer A

One major challenge in a purely symmetric system is how to share the secret key. Encrypting the key with a passphrase is out of place here, since we still have the fundamental problem of sharing the passphrase. Answers b and d refer to asymmetric cryptography.

12. Answer B

Due to complexity and security provided, the most commonly used type of symmetric cipher is a block cipher. DES, 3DES, AES, Twofish, Blowfish and others are examples of block ciphers. Generally, blocks ciphers provide greater security than stream ciphers. However, performance suffers.

13. Answer B

Authenticity is provided through the use of the sender's public key. Both symmetric and asymmetric provide privacy. Integrity is provided by hashing algorithms, which rely on one-way math (not a key) and non-repudiation requires a hash.

14. Answer D

In symmetric cryptography, a secret key needs to be shared between two parties to encrypt private messages. However, in asymmetric algorithms, the recipient's public key is used to provide privacy. The

public key contains no sensitive information and does not need to be kept secret.

15. Answer C

There is no such thing as a public key compromise as there is nothing sensitive attached to a public key. The secrecy of asymmetric algorithms comes from the relationship between the public and private key and the fact that it should be impossible (or at least highly unlikely) to determine the private key from the public key.

16. Answer B

Symmetric keys can provide the same strength of encryption with much shorter keys. RSA Security 1024-bit RSA keys are equivalent in strength to 80-bit symmetric keys, 2048-bit RSA keys to 112-bit symmetric keys and 3072-bit RSA keys to 128-bit symmetric keys RSA claims that 1024-bit keys are likely to become crackable sometime between 2006 and 2010 and that 2048-bit keys are sufficient until 2030. An RSA key length of 3072 bits should be used if security is required beyond 2030. NIST key management guidelines further suggest that 15360-bit RSA keys are equivalent in strength to 256-bit symmetric keys.

17. Answer C

Symmetric ciphers provide good, fast privacy, however exchanging the shared key requires some other means than the symmetric algorithms can provide. Frequently, the key exchange is handled by an asymmetric algorithm while the data exchange is provided by the symmetric algorithm.

18. Answer B

Though while Bob could also read documents destined for Alice, being able to sign documents as Alice would affect the accountability of the system.

19. Answer D

In asymmetric cryptography, privacy comes from using the receiver's public key to encrypt the information. In this event, only the receiver's private key can decrypt (which only the legitimate receiver should have.)

20. Answer B

When initiating a secure connection with a web server using https, the server responds by sending the client its public key on a certificate, ideally signed by a trusted Certificate Authority. The server's public key will then be used to encrypt a session key from the client.

21. Answer C

A digital signature provides non-repudiation (a combination of integrity and authenticity) for a message. With a digital signature, the message is hashed with a hashing algorithm like SHA-1 or SHA-256. The hash is then encrypted with the sender's private key using an algorithm like RSA.

22. Answer D

Diffie-Hellman is described as providing a means for two parties to agree upon a key without having to send that key across the network. It has traditionally been used as a means for the two parties to agree upon a session key, which will then provide symmetric encryption for the data.

23. Answer A

ECC (Elliptical Curve Cryptography) is a very fast and efficient protocol used to protect communications on devices with limited processing power. Its secrecy is based on the algebraic structure of elliptic curves over finite fields.

24. Answer D

RSA has replaced DSA as the current algorithm used as the standard for digital signatures.

25. Answer C

Hashes are based on one-way math—math that is very easy to perform one way, but exceedingly difficult to reverse. Passwords are frequently stored as hashes for this reason. If a password is forgotten, a network administrator can't view the password, though they can reset it.

26. Answer A

A collision is caused when two different contents produce the same hash. In this instance, the hash has been broken and is no longer reliable as it doesn't detect a change in content. However, as everything encrypted can be decrypted, with another effort all hashes can have a collision. The strength of the hashing algorithm is in its resistance to collisions.

27. Answer C

A birthday attack is based on the idea that it is easier to find two hashes that just happen to match rather than trying to produce a specific hash. It is called a birthday attack based on the fact that it is easier to find two people whose birthdays just happen to match, rather than someone with a specific birthday.

28. Answer A

Because there is no indication of the origin of the message or file, there is no guarantee against spoofing if only a hash is used. Authenticity must be added in order to get an assurance against spoofing.

29. Answer C

A MAC includes a message plus a symmetric key that only the sender and receiver should know. Because two users share this symmetric key, we can't get true non-repudiation. Even though this doesn't supply the same assurance that a digital signature does, it requires less of an infrastructure.

30. Answer B

A sender uses his or her private key to encrypt the hash, producing a digital signature. The receiver verifies the digital signature by using the sender's public key to decrypt the hash. If the hash can be decrypted using the sender's public key, it had to have been encrypted by the sender's private key (Which only the sender has.)

31. Answer B

A private key should never be on a certificate or any other mechanism that is made public. As a matter of fact, even the Certificate Authority will not know the server's private key. As part of a server's request to a CA for a certificate, the server generates a public/private key pair. The public key is registered with the CA, and that key is added to the certificate.

32. Answer D

A message indicating a certificate has not been signed by a trusted authority indicates that the Certificate Authority's public key is unavailable to verify the authenticity of a web server's certificate. The way CAs certificates are made available to web browsers is that they are loaded into the certificate repository within the browser (often by the vendor who provides the browser). A trusted CA is one whose certificate is accessible on the client's system.

33. Answer C

OCSP is a protocol that streamlines the process of verifying the revocation status of a certificate. An OCSP server or responder is responsible for checking with the CAs CRL (Certificate Revocation List) periodically and provide a reasonable current assessment of whether the certificate has been revoked.

34. Answer B

Encapsulation "wraps" the data into some sort of packaging—usually a header and a trailer. Encryption is a transformation process that involves taking plaintext and transforming it into ciphertext through the use of a key and an algorithm. IPSec provides security for the portions of the packet that are encapsulated.

35. Answer C

IPSec, in tunnel mode, provides encryption for the entire IP packet. IPSec adds its own header and trailer to the packet. The IP entire packet is the IPSec Payload. Though this can take longer, it provides better security.

36. Answer C

In creating a secure tunnel from one site to another, IPSec is normally configured to operate in tunnel mode. Tunnel mode provides greater security by encrypting the header, payload, and trailer of the IP packet.

37. Answer A

Diffie-Hellman is an algorithm whose sole purpose is to allow key agreement without pre-shared secrets and is used by Oakley, a sub-protocol of IKE.

38. Answer B

ESP is the only sub-protocol that provides encryption. AH provides non-repudiation, but no privacy services.

39. Answer A

NAT (Network Address Translation) has the primary function of hiding internal IP addresses from hosts located outside the network. A NAT device does this by removing the original source address and replacing that address with its own external interface's address. Though this service is very helpful in enhancing network security, the header modification is detected by AH. For this reason, NAT and AH are natively incompatible, though solutions like NAT-Traversal are used to make the two work together.

40. Answer A

A medium security organization is best suited to an area with high visibility and natural surveillance. Security through obscurity is a myth and often leaves an organization more vulnerable.

41. Answer D

Anyone trying to access the building without proper credentials should be escorted to security. If they are simply denied access, they will wait for someone else to come along that will let them in. Additionally, even if that individual is a recognized employee, they should still be escorted to security. It is possible that employee has been terminated and his credentials have been revoked. Disgruntled employees have been the source of numerous attacks resulting in the loss of life, property and data.

42. Answer D

Generally, access points should be placed in the center of the building, allowing the walls and other physical aspects of the facility to absorb the signal and help contain access to Wi-Fi to the building. Additionally, signal strength can be manipulated to reduce the chance of outside access.

43. Answer C

CCTV cameras could provide surveillance to disprove employee claims of improper physical access. Though doors data center doors should certainly be locked and badged access to a building is helpful, these solutions don't protect against employee actions once in the building. Further, though the policy is important, it is an administrative control that simply deters fraud. It will not detect the fraud.

44. Answer B

A kick plate is designed to protect the bottom of the door against cosmetic damage but doesn't really enhance its physical security. A strike plate is the part of the locking mechanism that re-enforces the door at the doorknob area. Hinges can be protected by encompassing them or by reinforcing them, so they are resistant to tampering.

45. Answer B

Positive air flows are designed such that air flows out of a room instead of into it. This limits the ability of contaminants to flow from room to room.

Domain 4

1. Answer A

Layer 1 of the OSI Reference Model is referred to as the "Physical Layer" and provides physical connectivity to the network. Cable, connectors, hubs and any device that is only concerned with creating a means for the physical signal to traverse the network are Layer 1 devices. Though there is an element of a NIC (Network Interface Card) that does provide physical connectivity, it is considered by most to be a Layer 2 device.

2. Answer C

All copper cable is susceptible to eavesdropping to some degree. Even shielding of twisted pair cabling only makes an improvement to its resistance to tapping and eavesdropping. However, if the goal is to find a type of cable that is truly immune to interference and much more difficult on which to eavesdrop, fiber optic cable is the best choice. Though fiber has traditionally been more expensive and more difficult to work with, it is becoming more commonly used, and prices are dropping.

3. Answer C

Lower layer devices are usually faster than upper layer devices, as these devices are not concerned with complicated inspection and decision making. In order to make decisions at Layer 7 for instance, the lower Layer headers would have to be stripped away, to provide deep packet inspection and direction. Layer 1 devices just provide a medium for the signal to travel, without taking the time to analyze or inspect.

4. Answer D

Ethernet Media Access uses CSMA/CD. This indicates that hosts will "sense" the cable to determine if data is being transmitted. However, multiple hosts could have sensed that the media was available at the same time. In this case, if multiple hosts transmit on the cable it causes a collision which should be detected immediately. A hub would not help with this problem. In order to limit collisions, a switch is necessary.

5. Answer C

In order to resolve a known IP address to an unknown MAC address, a host uses an ARP (Address Resolution Protocol) broadcast. ARP uses a broadcast to query the MAC address for a specific IP address. That MAC address is then added to the ARP cache, so as to eliminate the need for another broadcast should that information be needed again.

6. Answer C

CSMA/CA which is specified for devices following IEEE standard of 802.11, or Wi-Fi systems. In this method, a client sends a signal to indicate its desire to transmit. As a result of this signal, no other host transmits. In CSMA/CA environments, collisions are not simply reduced but eliminated.

7. Answer B

Switches maintain a CAM table that maps MAC addresses on the network to physical ports on the switch. This function allows the switch to direct data out of the appropriate physical port where the host is located, as opposed to indiscriminately broadcasting the data out all ports as a hub does. In a MAC flooding attack, the attacker sends the switch many Ethernet frames, with each one containing a different source MAC addresses. The intention is to consume the limited memory set aside for the CAM table. Ultimately, this process overwrites the legitimate entries that the switch has learned. Once the switch no longer has legitimate entries in its CAM table, it broadcasts data until it re-learns the MACs of the legitimate hosts.

8. Answer B

A switch serves two main functions on a network. First, it directs traffic out the appropriate physical port for the destination device. This prevents the need for the switch to send all traffic out all ports, as a hub did. Secondly, each physical port on a switch is its own collision domain. By lessening the number of hosts in each collision domain, there are fewer systems competing for time on the cable.

9. Answer A

With switches being used, traffic is directed out the appropriate physical port that is mapped to the recipient's MAC address. Since most likely there is no traffic addressed to the sniffer the only traffic being directed out that port would be ARP broadcasts used to learn the MAC address of the recipient.

10. Answer C

A router is usually considered a Layer three device because of its capability to handle the best path determination and to use IP addressing. However, routers must have some form of physical interface which is Layer 1. Also, once traffic is sent to the proper interface on the router, it uses an ARP broadcast (Layer 2) to locate the local client.

11. Answer B

The primary purpose of a VLAN is to create separate broadcast domains on a network. This function has traditionally been the responsibility of routers. However, routers are expensive and more difficult to logically configure, so this capability has been incorporated into switches.

12. Answer C

Broadcasts domains are subnetted and identified by their network addresses. IP addressing is a Layer 3 function. Though a VLAN can provide this segmentation on either type switch, when a standard switch is employed, the switch (Layer2) can't "understand" the difference between the network IP (Layer3) addresses. In this case, the VLANs would not be able to communicate. With a Layer3 switch that understands IP segmentation, the VLANs would be able to communicate.

13. Answer A

Media Streaming would best benefit from using UDP as its transport Layer protocol. Because media streaming is so very bandwidth intensive, speed and throughput are essential. Though UDP can also be used for file downloads through the TFTP (Trivial File Transfer Protocol,) usually TCP is used for small files.

14. Answer D

Since UDP is connectionless, it has no needs for fields that assist with guaranteeing communication or handshaking. However, UDP still requires the use of port numbers in order to identify the protocol or service being transmitted.

15. Answer B

The main difference between the protocols FTP and TFTP is that they use different layer 4 protocols. FTP uses TCP that provides connection-oriented delivery. TFTP uses UDP for faster connectionless delivery of data.

16. Answer B

The Presentation Layer sends data to the Application layer. This Layer provides a translation into standard formats, encryption, and compression. Though there are no specific protocols that work at The Presentation Layer (6,) most application Layer protocols are considered to function across the top three Layers.

17. Answer D

A session hijack occurs at the Session Layer (5.) In session hijacking, an attacker uses session-based information, such as Session ID, Username, and any other cached information, to step in and take over an existing session.

18. Answer D

The best way to mitigate sidejacking is a well-designed and secure website. The server should use https://

for all pages served instead of just the ones for login information. On the client side, the best way to protect against this attack would be to secure your network to ensure that there are no unauthorized devices and packets are not being sniffed.

19. Answer A

An application proxy is the best choice in this question. In order to make decisions based on content, a screening device would need full access to all layers of the OSI stack. Application layer devices are the only ones who have this degree of access.

20. Answer B

Though Application Proxies do provide a high degree of security through deep packet inspection, the can cause a significant performance decrease. The first line of defense is often a screening router that has very basic ACLs (Access Control Lists) to evaluate traffic very quickly.

21. Answer C

Though blocking all downloads would keep modified files from being downloaded, it would interfere with normal operations. The best means of ensuring that files downloaded are from the true server, as presented, and to ensure these files have not been modified is to ensure only files digitally signed are able to be downloaded. Digital signatures provide both authenticity and integrity.

22. Answer C

The earlier standards for Wi-Fi (802.11 a, b, g) did not support WPA II and were only capable of using WEP and later WPA (which provided much less security than their successor). WPA II was required to be supported by any standards after 802.11i.

23. Answer A

RADIUS (Remote Authentication Dial-in User Services) allows authentication through a central authentication server. This technique is frequently implemented in corporations that do not wish to manually configure authentication rules on each of their Wi-Fi access points (or VPN servers, RAS, or other network access devices.) RADIUS is only available in Enterprise mode.

24. Answer C

The most significant change brought by WPA II was the use of the AES algorithm. AES is a block cipher, which is a sizeable improvement over the stream cipher RC4; both WEP and WPA used RC4. Block ciphers are generally much stronger than stream ciphers, though they are slower. RC4 also had a short encryption key (either 40 or 104 bit) whereas AES can provide 128, 192, or 256-bit encryption.

25. Answer B

One of the big benefits of a cloud infrastructure is the elasticity it offers. Elasticity is the degree to which systems are able to adapt to changes in workload by provisioning and de-provisioning needed resources automatically so that each time the available resources match the current demand as closely as is possible.

26. Answer A

IaaS stands for Infrastructure as a Service and provides cloud-based access to routers, switches, servers, storage and other elements necessary to support a network infrastructure.

27. Answer D

In a community cloud deployments, storage is usually provided to clients of the same or similar industries that require the same security implementations, usually due to compliance issues. In this case, there is likely a cloud service provider that houses medical information from other healthcare providers or others required to maintain HIPAA compliance. This solution will most likely be cheaper and easier to manage than hosting their own private cloud.

28. Answer B

In a SYN flood attack, the malicious host sends a large number of SYN packets to the recipient, who in turn opens up space in memory to process the data that should be coming as the result of the handshake. Eventually, the system's available memory is exceeded, causing a DoS.

29. Answer D

Blocking ICMP at the firewall is almost always mandated. ICMP is a frequently exploited protocol. Even though it is useful inside a network for troubleshooting, there is no need to allow ICMP packets from outside the networks. However, numerous upper Layer services like DHCP, DNS, and TFTP (as well as others) require UDP to work properly. Therefore it is more difficult to protect against Fraggle attacks. Nevertheless, there are other strategies to mitigate against Fraggles. For one, directed broadcasts should be blocked. Directed broadcasts are those that originate from outside the firewall.

30. Answer B

An ARP poisoning attack is implemented when an attacker overwrites legitimate entries in the cache and replaces them with the addresses of rogue devices. Malicious modification of cache is usually referred to as poisoning or pollution attacks.

31. Answer B

DNSSEC (Domain Name System Security Extensions) is a set of extensions that provide security to the DNS service through enabling DNS responses to be validated. DNSSEC provides origin authenticity and integrity. With DNSSEC, DNS is much less susceptible to spoofing.

32. Answer D

If a rootkit is detected, the best way to ensure that it is removed is to wipe the system, reinstall the operating system from original media, then restore data from backup. It can be difficult to tell when a rootkit was installed, so restoring the operating system from backup could potentially reinstall the rootkit as well.

33. Answer B

This degradation is most likely the result of a worm infestation on the network. Because things were fine on Friday, the indication is that the issue is not a result of a virus, because a virus requires user interaction. A worm, however, consumes a tremendous amount of network resources and is able to spread throughout the network on its own.

34. Answer A

A packet-filtering firewall provides layer 3 and 4 inspection of headers for determining if traffic should be blocked or allowed. Some of the information that can be found at these layers is source and destination IP address (Layer 3), source and destination port (Layer 4) and protocol (Layer 4.)

35. Answer C

The primary purpose of AH is to detect spoofing, which means, it is designed to protect against modification of the source addresses. Because NAT modifies that source address, the two are natively incompatible.

36. Answer C

An application proxy is the best choice in this case. Application proxies have time awareness, Active Directory integration (which is likely needed to limit specific users,) as well as deep packet inspection which allows access to the content of data. Though Application Proxies provide much more in-depth inspection, they are usually slower and more expensive than lower layer firewalls.

37. Answer B

Packet Switching technology like MPLS (Multi-Protocol Labeled Switching,) VOIP and ADSL divides data into packets. Each packet finds its own best pathway to the destination. Packet switching is a much faster technology than circuit switching.

38. Answer D

Any type of traffic on an IP network is susceptible to sniffing. Natively, VOIP uses insecure protocols like RTP (Real Time Protocol) that does not provide encrypted communications. Though more secure protocols can be used, natively VOIP offers no inherent security. Tools such as Wireshark can very easily sniff VOIP traffic and reveal the details of the communication.

39. Answer C

Multiprotocol Label Switching (MPLS) is provider-based network designed for networks which need high-performance communications. MPLS networks direct data from node to node in the network based on short labels rather than long network (IP) addresses. This process is quicker than using complex routing tables. The headers added to the packet before traversing the MPLS network includes a field for QoS, so that VOIP traffic is prioritized.

Domain 5

1. Answer A

Hot/cold aisles are used in the server room and other areas where there isn't always much room for air to circulate properly. A major concern would be that as one system expel hot air, another system would use that hot air to cool those systems. In order to prevent this problem, systems are set up to expel hot air back to back (hot aisle) and to pull in only cool air from the cold aisle.

2. Answer C

Planting bushes directly underneath windows makes it more difficult for an attacker to gain entry. Fences, lighting and surveillance cameras will help enhance security but are not environmental. Security through obscurity is the false idea that being less visible improves security (in fact, that makes an organization less secure as there is no visibility and crime is more likely to go undetected.)

3. Answer B

Burglar alarms are reactive devices that are activated by some sort of trigger. This trigger indicates the breach has happened or is happening. Lighting usually considered a deterrent, but motion-detection lighting would be considered detective. However, since this fact was not mentioned in the question, the best answer is B.

4. Answer B

Group policy can be used to enforce rules in relation to passwords. Password complexity requires users to have passwords which meet certain criteria, such as length, uniqueness, etc. Also, the length of time for which a password is valid, and password history can all be controlled with group policy.

5. Answer B

Cognitive passwords are knowledge-based authentication consisting of words or phrases which a user should intrinsically know. Mother's maiden name, name of someone's first pet, high school mascot, etc. are examples of cognitive passwords. Keep in mind that in today's world of information sharing many of these pieces of information may be readily available on the internet.

6. Answer C

A rainbow table is a precomputed table designed to be used for reversing cryptographic hash functions. Since frequently hashes are stored as passwords, the most frequent use of rainbow tables is to crack passwords.

7. Answer C

Cookies are often placed on user systems when the user first opens an account with a financial server or other server wanting to provide seamless two-factor authentication. When a user tries to log in from a new system, they get a warning message telling them that they are logging in from an untrusted system. At this point, the user is prompted to provide additional authentication information.

8. Answer B

One-time password generators allow a one-time password to be used without dramatically increasing the overhead on the user.

9. Answer C

Multi-factor authentication is not simply providing multiple means of authenticating; it requires providing at least two different types. A smart card is only single-factor authentication—a card is something you have. In almost every imaginable instance, the smart card is coupled with a password or PIN. Then and only then does it provide multi-factor authentication. Answer C uses a password (Type I) and a thumbprint (Type II.)

10. Answer C

Though biometrics offer the best authenticity for single factor authentication, multi-factor authentication is always best. Adding a password (Type I) or a Smart card (Type II) would offer multifactor authentication when used in conjunction with biometrics.

11. Answer C

The type of technology that will be chosen is based upon the other three options. For instance, an organization will have a cost in mind; they will have a reasonable understanding of the accuracy needed and the degree to which their users will be required to submit to verification. The answers to these questions will determine what technology type to choose.

12. Answer A

FAR (False Acceptance Rate) indicates the number of times that someone is able to gain entry without having the appropriate credentials. This number is inversely related to FRR. When FARs go down, FRRs go up. However, you're not wanting to accomplish a high FRR, though that might be a result of changing the settings.

13. Answer C

In Kerberos, a user enters his or her password onto a system. The password is stored locally. The

username is sent to the authentication server. The authentication server generates a TGT (Ticket Granting Ticket) and encrypts the TGT with the user's password. If the user had entered the correct password, then the TGT can be decrypted. The face that the user has a decrypted TGT proves that the user authenticated properly.

14. Answer A

When a client requests a session with a principle in a Kerborized environment, the TGT issues a ticket. This ticket contains two copies of the exact same session key. One copy of the key is encrypted with the user's password. The second session key is encrypted with the service's password. With this technique, only the correct password will decrypt the session key on the client side and only the correct key of the service. Kerberos is a purely symmetric environment, so the key exchange is cumbersome.

15. Answer D

In Windows-based systems, an authentication token contains a list of the groups in which the user is a member. This list of group membership is compared up against the access control list for the resource and the determination is made whether to allow access.

16. Answer B

The above answer uses context, not content-based decisions. The member is not being blocked to the content of the payroll information—she has access to it all day. Context-based access control evaluates accessed on HOW the information is being accessed.

17. Answer C

The Clark-Wilson security model states the need to protect trusted resources from untrusted entities. In order to do so, an interface is used to enforce well-formed transactions. By constraining the interface, we constrain the activities that the junior admin can perform.

18. Answer C

Almost all firewalls use some form of rule-based access control to filter traffic. The rules on the firewall are usually referred to as ACLs (Access Control Lists.) In the question, the most basic firewall of the four listed is the packet filtering firewall. This is a layer three device which inspects information in the packet header at the network layer, which would include source and destination IP address, port number, and protocol.

19. Answer B

The IEEE 802.1x standard for EAPoL. 802.1X authentication involves three elements: the supplicant, the

authenticator, and the authentication server. The might be a dial-up client, a VPN client, a Wi-Fi device or some other device requesting access. The authenticator is a network access device, such as a wireless access point, a VPN server, etc. The authentication server is typically a server running RADIUS or other similar software.

20. Answer C

The greatest benefit of a decentralized environment is granularity. Each individual network access device could have its own individual policies and access control criteria and could be more closely aligned with the individual roles of each server.

21. Answer D

With CHAP, when a peer tries to authenticate, the authenticator sends a challenge to the peer. The peer performs an algorithm on the challenge and responds with the result. If the result is what the authenticator expected, the peer is authenticated.

22. Answer A

Heavy metal absorbs stray signal and is frequently used to prevent leakage. A faraday cage is made of heavy metal and can describe an actual cage, room, building or any other casing that can absorb the signal.

23. Answer A

Data encryption, though important for privacy protection, is not a protection against data emanations. Often the study of the emanations analyzes the frequency, power consumption and other details which encryption would not mitigate.

24. Answer B

Though Cloud-based solutions provide centralized management and ease administration of users and accounts, CSPs (Cloud Service Providers) are not regulated and not required to provide the degree of security your company may need. Obtaining a well-written contract and auditing that contract are two ways to ensure your company's security requirements are met.

25. Answer D

Identity as a Service typically indicates that the directory database is cloud-based and managed by a cloud service provider. Though the organization can host its own directory service, it is less likely to use IdaaS if storing the database on the internal network.

26. Answer D

SAML (Security Assertion Markup Language) is an XML-based, open data format to facilitate the exchange of authentication and authorization information between parties, often across organizational boundaries.

27. Answer A

User account provisioning creates, modifies, and disables/deletes user accounts as well as their profiles across the IT infrastructure and business applications as needed. Many provisioning tools can use approaches such as cloning, roles, and rules to automate onboarding, offboarding or other administration workforce processes (new account creation, transfers, promotions and/or termination.) Provisioning tools can also automatically aggregate and correlate identity data from entities such as HR,CRM, mail systems or other "identity stores." Fulfillment can be initiated via self-service, from a management request or changes to HR systems.

28. Answer C

In the provisioning lifecycle, before an account is created, or credentials assigned, there must be a policy in place to determine how an individual provides proof of their identity. Perhaps reference checks, certification verification or other procedures must be followed before a user is granted access to company systems.

29. Answer C

In request-based provisioning, users or their managers search for and request access to applications, privileges, or resources with a system. These requests are then validated by workflow-driven approvals. Finally, they will audit for reporting and compliance purposes.

Domain 6

1. Answer B

A vulnerability assessment will have the least impact on your network, while still verifying that common security vulnerabilities have been mitigated. These tests are generally considered passive, as they are looking for weaknesses but not attempting to exploit them.

2. Answer A

The Rules of Engagement document provides important information detailing any limitations to a pen test. Certain systems, tools, times, etc. may be off limits, and this information needs to be clearly understood. Pen tests introduce risk to the environment, and ideally, these risks should be reduced as much as possible.

3. Answer C

Full knowledge penetration begins with providing the tested the same amount of information an administrator would be expected to have. This type of test emulates a scenario when it is the network administrator or some other privileged user who is committing the attack.

4. Answer D

The first step of any type of network assessment is to meet with management and determine the goals. How we approach testing will depend on what our ultimate purpose is.

5. Answer D

Most of the information listed above is easily accessible to the general public. Names of managers, office locations, and phone numbers are obtained from the internet or simply from querying the organization. This information is often used to form the basis for a social engineering attack. Internal IP addressing schemes, however, are almost never published publicly.

6. Answer B

The purpose of footprinting is to gather information about the configuration of the network. An attacker will use this technique to learn about the services on the network and the hosts which provide them. An attacker may also learn about the various connectivity devices and where they are placed, as well as other critical information. Once the network has been footprinted, and the attacker has located a desirable system, that system is often fingerprinted. The goal of fingerprinting is to determine the operating system running on the host, in the hopes of finding known vulnerabilities.

7. Answer C

The scenario above describes entrapment, as the attacker is tricked into accessing a system that he might not have accessed otherwise. A honeypot should entice an attacker away from other resources without persuading them to commit a crime or violate policy.

8. Answer D

A pseudo-flaw is an intentional fault written into the code of an application or operating system in order to distract or trap an intruder.

9. Answer A

In order to mitigate the risk of your honeypot being compromised and causing damage (either to your network or someone else's) the most logical place for a honeypot is in the DMZ. The honeypot can attract attackers and can be placed alongside other legitimate DMZ servers, providing early warning of threats.

10. Answer B

Profile matching systems look for activity on the network that is unexpected and label that activity malicious. Behavior and anomaly based systems fall into this category and frequently report false positives. They greatest problem with false positives is that they can desensitize administrators to alerts and lead them to be complacent.

11. Answer C

Since zero-day attacks are those for which no signature exists, signature-based systems cannot detect these attacks. It can take weeks or even months before a signature is developed for an attack. Until that signature is developed, the IDS cannot detect the attack as malicious activity.

12. Answer C

An anomaly-based IDS monitors network traffic and compares it against a baseline. The baseline is created and will then be used to identify what is "normal" behavior for that network. Considerations can include the amount of bandwidth, which protocols are used, ports frequently utilized, etc.

13. Answer B

The Rules of Engagement document should include the details necessary for the penetration tester to determine necessary action in the event that a critical security error is found. The tester should never act on his own to correct problems as this would violate the separation of duties and change control policies.

14. Answer A

In DDoS (Distributed Denial of Service) attacks, unsuspecting network hosts are commandeered to launch an attack on another network. These hosts are often referred to as zombies or bots. These systems are usually configured to send packets with spoofed source addresses.

15. Answer B

The Executive Summary of your penetration testing report should present the meaningful information summarized in such a way that the senior managers can understand. Many executives are not technical experts and need the information broken down and simplified.

Domain 7 - Part 1: Investigations and Daily Processes

1. Answer D

The primary job of a first responder is to preserve the evidence. Digital evidence is extremely volatile, and one must be certain that the integrity of the evidence is preserved before the investigations begin. Documenting the Chain of Custody should begin as soon as evidence is identified.

2. Answer B

Typically, CPU registers store instructions or addresses for a very short period of time. These registers are extremely volatile elements of the system.

3. Answer C

One of the most important requirements in forensics investigations is that evidence should not be modified as a result of its collection. The first responder should immediately preserve the evidence to the best of their ability, and whenever possible, an examiner should work with a copy and not the original system or device.

4. Answer C

A signed contract is considered to be "Best Evidence." The "Best Evidence Requirement" is a legal principle that considers the original version of a document as the superior form of evidence. The rule specifies that a copy or fax would not be admissible if an original of the document exists and is obtainable.

5. Answer D

Expert witnesses, such as forensic experts, cryptography experts, etc. are considered to present secondary evidence.

6. Answer C

One of the exceptions to the fourth amendment (which protects citizens from illegal search and seizure by law enforcement) is in cases of exigent circumstances. Exigent circumstances describe a situation in which evidence is in immediate harm of being destroyed.

7. Answer A

Copies of documents are ruled as second-hand, or hearsay evidence. In order to be admissible in court, steps need to be taken to prove their authenticity and integrity. Hashing, Digital Signatures, private keys and other controls can assist in providing the logs' legitimacy.

8. Answer C

In relation to a policy of this nature, email auditing should take place and become a part of normal business operations. For instance, if this policy was only used to investigate a particular employee, it may appear as if that employee is the only one to whom the policy applies. Best practice dictates that we create policy, implement policy, audit policy and enforce the policy to all to whom the policy applies.

9. Answer D

In order to reduce the risk of an attacker modifying audit logs, all choices above are valid. Write-once media obviously should not be able to be overwritten or modified. Hashing detects any modification. And finally, the regular review of audit logs will help an administrator familiarize himself with standard activity so that (hopefully) an anomaly will stand out.

10. Answer B

Provisioning provides users access to data and technical resources. The term is used in reference to organizational resource management. Provisioning combines the duties of the human resources and Information Technology departments in an enterprise, where users are given access to data or granted authorization to systems, software, and databases based on their unique user identity, and secondly, users are granted access to hardware resources such as computers, mobile phones, and tablets. The process requires that the rights and privileges are monitored and tracked to strengthen the security of an enterprise's resources.

11. Answer D

Self-service account provisioning allows users to participate in certain aspects of the provisioning process, helping to reduce the administrative overhead. Frequently, users are able to request an account and choose, manage and reset their own passwords.

12. Answer D

Automated account provisioning requires each account to be added through a centralized interface, usually in an HR application or database. Every person has an account which is linked to each one of their corresponding accounts. Any changes to the primary account (credential changes, role changes, workflow

changes, termination, etc.) are automatically updated to all accounts.

13. Answer B

DAI (Dynamic ARP inspection) is a security feature which rejects invalid and/or malicious ARP packets. This feature prevents a type of MITM attacks in which an attacker intercepts traffic for other systems by poisoning the ARP cache of its neighbors.

14. Answer B

A VLAN provides logical segmentation of networks. Though VLANs are created on switches, not all switches support VLANs (this is why answer A is incorrect.) A router would also create this segmentation, but on a port-by-port basis, a router is much more expensive

15. Answer A

Any filtering mechanism that uses whitelisting will block all traffic, except for what is specifically allowed on a so-called "whitelist." This filtering method works well with firewalls but is likely to be entirely too restrictive for situations like spam filtering for mail servers. It is hard to imagine having a mail server that blocks all traffic except for that from a particular network or domain. So in that instance, we would use blacklisting. Blacklisting would allow all traffic, except for that which is on the so-called "blacklist."

16. Answer D

The primary function of incident response is to minimize the impact of the attack on the organization as a whole. Often one of the first things we consider is to isolate the affected system or subnet from the rest of the environment, so the attack doesn't spread and affect other systems.

17. Answer C

Most monitoring software includes the ability to configure alerts in the event that certain thresholds are exceeded. This is the timeliest means of detecting these issues. Reviewing logs and querying metrics may work, but would only be done periodically. An alert will message the admin immediately.

18. Answer C

Traffic on the internal network should have an internal network address. If outgoing traffic has an external address, it is often an indicator that the systems have been compromised with malicious software that allows them to be remotely controlled and can access the internet through public addresses. Traffic coming into the internal network with an internal address might indicate a spoofing attack.

19. Answer A

Because in the above scenario, only a small representation of audit logs are presented, it could easily be ruled as incomplete. In order to have a better likelihood of admissibility, it would be better to collect data from the entire week or even month. When a small amount of data is presented, it may appear that the only information presented is that supports the goals of the investigator, and may not represent the complete picture.

20. Answer D

In regards to forensics, one of the most important rules is that the investigations process should prevent alteration of the evidence. First responders are responsible for ensuring that the identified evidence is preserved in such a way as to prevent modification.

21. Answer C

Three forensic hashes are necessary to provide the proof that the hard drive has not been modified as a result of the investigation. When it is determined that the drive needs to be analyzed, the drive should be placed in a write-protected system and hashed immediately, documenting the hash. Next, a bit-by- bit copy of the drive should be created, and that copy hashed (and documented.) Finally, after analyzing the copy in a write-protected system, the drive should be hashed again. All three hashes should be exactly the same.

22. Answer B

Mutual authentication requires both parties to provide authentication. Though most environments require users to authenticate, we often fail to require authentication of our network systems. Certificates, keys, and other mechanisms could provide a way for access points and other systems, such as DNS to prove their identity.

23. Answer A

Network Access Control is a network service designed to inspect systems and allow or deny access to network services based on client health. Good health might indicate a system has anti-virus software, anti-spyware, a firewall, as well as being up to date on patches and upgrades. Other criteria can be specified as well.

24. Answer C

The access list above is a typical ACL which might be found on any router. Traffic is denied from any source host to any destination host on port 23, which is telnet.

25. Answer A

A system that cannot be patched to the current level poses a threat to a network environment. However, since the payroll system is only supported on the current OS patch, the best way to protect the rest of the network is to isolate the unpatched server.

26. Answer D

Implementing a patch management server can streamline the patch management process. Patches and updates can be downloaded, tested and made available to users. Group policy can require the users to connect to the patch management server and download only those updates which were approved. Of note, even though security patches should be given priority, they should never be distributed without testing.

27. Answer B

Slipstreaming is a technique in which software updates are integrated into the original operating system media. With slipstreaming, the operating system and the updates are installed as part of the same installation, providing a more integrated process and fewer reboots.

28. Answer D

In order to promote the stability of systems, a change control process should be in place and should be strictly followed. When a change is proposed, the first step is to refer the change to the company's Change Control Board. The CCB will evaluate the change for risk and determine if the change should be made. At that point, the proposed change will be implemented and tested in a lab environment before being implemented.

29. Answer B

Though it is essential to follow the formal change control process whenever possible, at times, a change will have to be made to limit the impact an incident has on current business functions. At that time, the change should be implemented, as per your emergency change control process, which will likely include documenting the change and then referring the change for review by the CCB.

30. Answer A

When a modification or new installation works properly in a lab environment, but not in production, it is usually due to a discrepancy between how the lab configuration and the production environment.

CYBRARY

Kelly Handerhan's CISSP Preparation Course

Domain 7 - Part 2: Redundancy and Business Continuity

1. Answer B

"Mean Time Between Failures" is a metric that indicates the amount of time a hardware device should function before it fails. Once the MTBF is known, then an administrator or technician can be prepared for the failure of the device.

2. Answer D

RAID 5 is often defined as "Disk Striping with Interleave Parity" provides the same performance improvement as RAID 0 (Disk Striping.) However, RAID 5 adds parity information interleaved through the RAID array. The parity can be used to rebuild data from a failed drive.

3. Answer B

One-half of disk space is always set aside for redundancy in a RAID 1 array. Each drive is an exact replica of the other, so the array must be comprised of equal disk size.

4. Answer C

A cluster can be simply defined as multiple physical servers that function as a single node for the purpose of fault tolerance and often load-balancing. Of note, not all clusters provide load balancing though many today do.

5. Answer A

An Active-Passive cluster is fairly easy to implement and doesn't require a large investment or a monthly payment. Often in active-passive clusters, the primary server is the device that handles the entire workload; the passive cluster can be a low-end system that only comes online in the event that the primary fails.

6. Answer B

Redundant servers are usually unique devices on the network that are independently accessible. With clustering, nodes are incorporated into the cluster and are no longer accessible individually except through an administrative access.

7. Answer D

Unscheduled backups should be performed as a "copy." The copy function neither looks for nor cares about the archive bit. If a full backup was performed at 4:00 in the afternoon, the archive bit would have been cleared. The nightly backup, then, would've only contained changes to files that occurred since 4:00

8. Answer C

In order to have the assurance that the backup process is working, backups should be fully restored. Only then do you have the assurance that the backup is accessible and complete.

9. Answer D

When using incremental backups, the full backup must be restored and then each of the corresponding incremental backups. In this case, Sunday's full backup as well as the backup from Monday, Tuesday, and Wednesday must be restored.

10. Answer C

Electronic vaulting allows an organization with high availability needs to transmit transactions in batches to another facility or location numerous times a day. This allows for data to be more current in the event that a restoration is necessary.

11. Answer D

RPOs (Recovery Point Objectives) relate to data that must be recovered and the required age of the data. With an RPO less than 24 hours, nightly backups would not be frequent enough. Remote journaling, vaulting or shadowing should be considered.

12. Answer B

Database shadowing provides the quickest restoration and least amount of data loss in the event of a disaster or corruption. Transactions are written simultaneously to two separate databases, sometimes using different storage media for high availability of data.

13. Answer B

The recovery plan provides instructions on returning the most critical services to operation as quickly as possible. Criticality is determined in the BIA (Business Impact Analysis) and indicates the loss suffered without the process or service. Most critical processes cost the organization the most money while they are down. Reconstitution is the process by which operations are returned to the original or permanent facility and begins with the restoration of least critical, working to most critical.

14. Answer A

The RPO is the company's tolerance for data loss. If the company merely runs backup once a day, then the possibility is that a full day's worth of data could be lost. The organization may have determined that the need for current data is not worth the cost of more frequent backups. Remember, RAID is not a redundancy for data. If a malicious file infects one drive in an array, they are likely all infected.

15. Answer C

When leasing a cold site from a vendor, it is important to be aware of the fact that vendors frequently lease the same space to multiple organizations. This assumes that companies will just need these sites for a disaster affecting only their company. However, in the event of a regional disaster, the facility is available to the first of those leasing the site to show up. Cold sites are the least expensive of the other options.

16. Answer C

A simulation test goes through the motions to verify that the plan is accurate and complete. A structured walkthrough is sometimes referred to as a tabletop test because despite the name "walkthrough" it is actually a discussion based process involving the members of the disaster recovery team. The parallel test is one in which a portion of business operations are conducted at the offsite facility, while other processes take place at the original facility.

17. Answer B

Test verify the plan for accuracy and completeness. Employee response is evaluated in drills and exercises. Usually, by the time drills are conducted, the plan has already been tested and found to be complete.

18. Answer D

A full interruption test is the riskiest test because after fail-over, all business operations begin at the offsite facility. If for any reason the site were not ready, then the organization will likely lose some or all of its new transactions.

19. Answer C

The purpose of the BIA is to identify business processes and prioritize them based on criticality. Often risk analysis is lumped in with the BIA but should really be a separate function which examines threats and vulnerabilities that could lead to the compromise of those functions.

20. Answer A

A BCP policy is essential because it will include the commitment of senior management to support and fund the BCP process. This process is complex, lengthy, and has no direct ties to profitability. For this reason, not all managers buy into this project.

21. Answer B

A Business Continuity Planning team should include members throughout the various business processes so that each department's interests are represented. It is also helpful if those carrying out the plan are the same people who create the plan.

22. Answer B

When someone attempts to enter a building without providing the correct credentials, he or she should be escorted to security immediately. If you don't let him in or ask him to leave, he will simply wait for someone that he can follow to come along later.

23. Answer: D

It is best if the humidity is around 50%. Anything below this could lead to problems with static electricity. More than this can lead to condensation, which among other issues, can cause components to rust.

24. Answer A

An eight-foot fence is required to deter an intruder. Often barbed wire or concertina wire is used atop fences to add extra deterrence. Remember, there is no height fence that will prevent a determined intruder. There is always a taller ladder, or a means to go around, over, or under any type of fence. To truly protect your perimeter, use layered defense.

25. Answer D

A pre-action system holds water in a reservoir which is released into the pipe when the alarm is triggered. A plastic valve holds the water back until it melts, providing mitigation in the event of a false alarm.

26. Answer B

The Occupant Emergency Plan deals with the most important aspect of disaster recovery: Safety of personnel. It will include information such as safe evacuation of employees, how to determine that all employees have been evacuated, and any special procedures or processes that are necessary.

27. Answer B

Class C fire extinguishers should be located within fifty feet of electrical distribution systems. Class C extinguishers are designed specifically for electrical fires, though many extinguishers today are rated for multiple types of fire. Always check and be sure the correct type of extinguisher is provided and clearly marked.

Domain 8

1. Answer A

Input validation prevents improper entries from being passed along to the backend data. Examples of validation might include verifying the length of the input, examining for data control languages and data type. Input sanitization will attempt to "clean up" data before entry, strip improper characters or change single quotes to double quotes.

2. Answer C

In tests that involve fuzzing, large amounts of random data, referred to as fuzz, are entered into the software in order to ensure that validation techniques are effective.

3. Answer A

White box testing is a type of testing in which the tester has full access to the software's code and examines the code for structure and logic.

4. Answer C

Script kiddies are individuals with little true knowledge of hacking, and instead, are known for copying and pasting script from other, more knowledgeable attackers. When script kiddies run code, often they don't truly understand the potential for the loss they could be inflicting upon a system or network.

5. Answer D

A Highly structured attack is one that is instigated by attackers with more technical skill and competency than most attackers. Often these attacks can persist for long periods of time, and because the attacker is usually quite motivated, they will often continue until they have accomplished their objective.

6. Answer A

Ethical hacking or white-hat hacking are other ways to describe penetration testing. Though the term "hacking" has long held a negative connotation, in reality, it is neither positive nor negative. As long as the penetration test is authorized by the organization, then it is ethical to conduct these tests.

7. Answer B

Often the IP address and subnet mask would be known before beginning a scan (usually necessary to connect to the system.) Network services running would indicate open ports. Operating systems and applications have known vulnerabilities that may help an attacker gain access to the system.

8. Answer C

The company policy and mission statements are not likely to give an attacker much useful information. However, job postings for a Unix administrator would indicate that Unix systems are in place. The WhoIs database will provide information about publically registered domain names and may include information (technical contacts, name servers, addresses) that could be used in a technical or social engineering attack. Knowing branch office locations and phone numbers may also be helpful in a social engineering attack.

9. Answer B

Due diligence describes the research necessary to make good business decisions. By authorizing a vulnerability scan, the company is determining where their weaknesses lie. Once they take steps to correct the vulnerabilities, they are demonstrating due care.

10. Answer B

The above scenario describes code injection. If forms do not have a means of input validation, then there is the risk of an attacker inserting code into the available fields. If the code is passed along to the back end, it can be processed causing data loss and modification. The best defense is, as stated, input validation.

11. Answer D

It is recommended that SSL/TLS be used to connect to web servers for a secure connection. One of the reasons for this recommendation is that HTTP is a stateless protocol. Stateless protocols don't hold information based on the previous sessions, and either have to resend information or have the information cached. For example, authentication information must be transmitted for each request and often session information, such as the session id, is stored in cookies.

12. Answer B

An XSS (Cross-site scripting attack) is the most common attack on web applications. This attack relies upon exploiting a trusted website lack of input validation. Many client-side browsers check for pages that may be vulnerable, but it is best mitigated by good web application design with input validation.

13. Answer A

An XSRF (Cross Site Request Forgery) attack occurs by exploiting the trust a web server has in a currently logged in client. Through the use of pre-established session IDs and cookies, the malicious intruder is able to masquerade as the legitimate client and authorize transactions without leaving a trace. Often phishing emails with links to financial institutions or other desirable sites are used. Users should not sessions concurrently running that consist of secure and insecure connections.

14. Answer B

Indirect object access can occur when an application allows access to a resource solely based on user input. Providing additional authentication and access control, as well as using obfuscating the reference and ensuring it is not predictable will help mitigate this attack.

15. Answer D

Missing Function level access control is an attack very similar to exploiting direct object access, except the former allows additional privileges, where the latter allows unintended objects. Lack of predictability and greater access control will mitigate both of these issues.

16. Answer C

Java uses a security measure in its development environment to limit the behavior and (some) functions which are applied when the applets are sent as part of a web page. The term "sandbox" is a term that references the area of containment. For instance, the applets are sandboxed in the browser.

17. Answer C

A front-end application will allow the users an interface which will ultimately modify the backend database. However, the application will help ensure consistency and better-formed transactions through the use of data typing, drop-down arrows, field length limits and other restrictive means.

18. Answer A

Tokenization will remove the credit card information from the company's internal network while replacing it with a pointer, or "token." Merchants then use only the token to access, modify or maintain the individual customers' credit card information. The actual credit card information is stored at a secure offsite location.

19. Answer B

DNS is a distributed, hierarchical database, with different servers responsible for different portions of the namespace. For instance, there are root servers, top-level servers (.com, .net, .edu, etc.) as well as 2nd level and beyond.

20. Answer: D

The hierarchical database organizes data in an inverted tree, with the top-level as the root of the tree and the sub-levels branching out. The root is the ultimate parent object and objects directly below the root are its children. This continues throughout the hierarchy. This model mandates that each child object may have only one parent object.

21. Answer D

Relational databases store information in tables. Each table contains records and attributes describing the individual entities contained. Keys are used to build relationships between the tables, allowing information to be aggregated across tables.

22. Answer A

The primary key is a field necessary to identify each record as unique. Key fields are used to provide links between these tables to aggregate information.

23. Answer A

The cardinality of a database describes the number of rows in a relation. For instance, a common cardinality might be a one-to-many relationship. This would indicate that the primary key would appear once in its primary table and many times in a secondary table. For instance, customer 123, would only appear once in the Customers table but could appear many times in the Orders table.

24. Answer C

The schema of a database contains the complete description of the structure and contents of a database. One can think of the schema as the "blueprint" describing the logical elements of the database.