Mc Graw Hill

CCSP®
Certified Cloud Security Professional

Sample Chapter

CHAPTER 4:
Cloud Platform and Infrastructure Security

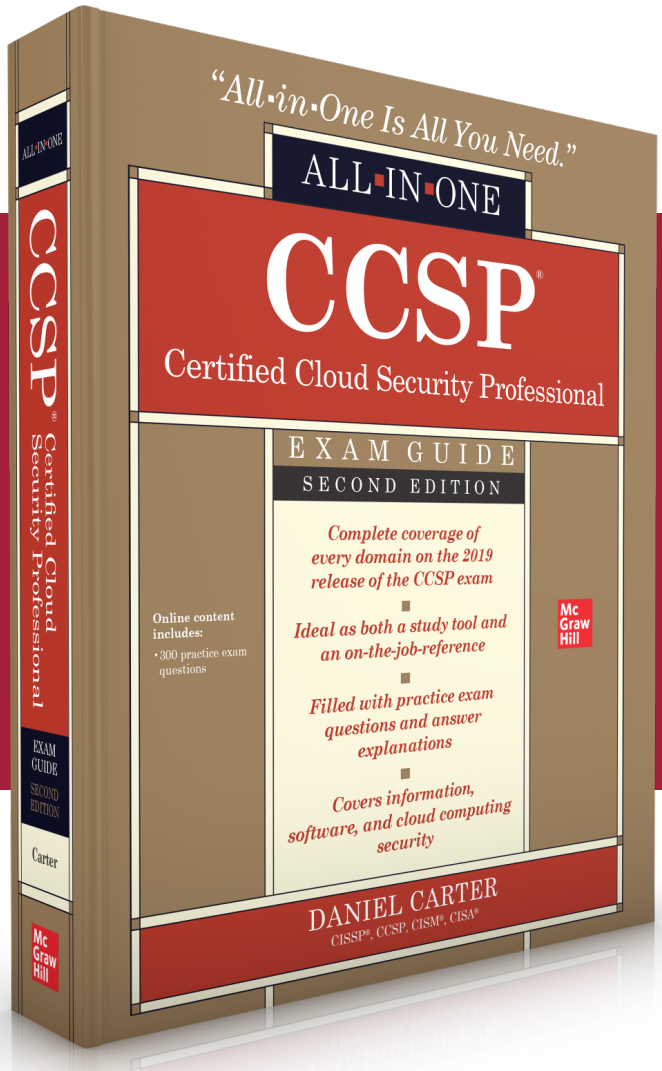

LEARN MORE

BUY NOW

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

CHAPTER 4

Cloud Platform and Infrastructure Security

This chapter covers the following topics in Domain 3:

- Physical aspects of a cloud environment
- Key components that make up a cloud environment
- Risks associated with cloud computing
- Designing and planning for cloud-based security controls
- Auditing in a cloud environment
- Disaster recovery and business continuity in a cloud environment

Cloud platforms bring unique benefits to an organization and have many attractive capabilities, including performance and scalability, the removal of focus on hardware, the placement of focus on business requirements, and measured service—all for a possible lower total cost and investment than the organization running its own data center. However, these platforms also bring unique challenges and risks because of the very same factors, and cost savings are not always what they may appear to be at the onset, or not on the scale often presumed. This chapter goes over those risks and challenges, how to address and mitigate them, as well as disaster recovery and business continuity requirements and benefits in a cloud environment.

Comprehend Cloud Infrastructure Components

The cloud infrastructure is made up of many of the same components that a traditional data center has, just applied from the perspective of a cloud environment. The cloud infrastructure also adds some unique components, as shown by Figure 4-1.

Physical Environment

While the model of cloud computing has been a revolutionary technological approach that organizations take for hosting systems and applications, the underlying architecture and requirements in a cloud environment are no different from the traditional data center model; the cloud environment simply abstracts that level of concern and detail from the

LEARN MORE BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

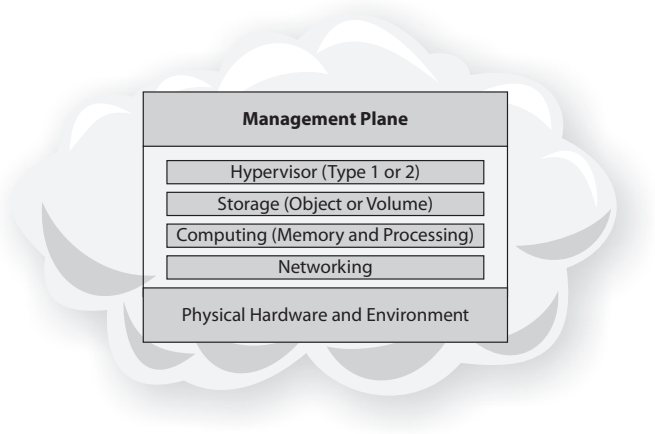

Figure 4-1 Cloud infrastructure components

cloud customer. However, especially with the large public cloud systems, the scale and coordination required in a cloud environment can be far more intricate and complex.

A traditional corporate data center, especially for a large company, will have thousands of computers and incredible cooling and utility needs. With a major cloud environment, you are typically looking at tens of thousands or hundreds of thousands of servers, spread across multiple (sometimes dozens of) physical locations.

Having such large-scale data centers requires enormous power and cooling resources. With the expectations in a cloud environment of high availability and resiliency, all systems must absolutely be redundant and allow maintenance to be performed in a way that does not cause any downtime or create single points of failure during any maintenance periods. With most cloud environments hosting a significant number of customers, any downtime will have enormous impact and be very visible to the customers and to the cloud provider. On the positive side, with so many customers pooling resources and the cloud provider focusing on an infrastructure specifically built for its cloud offerings and not for hosting numerous different types of systems with different needs, economies of scale can be leveraged in a way that an organization hosting its own data center would not be able to do.

Internally, a data center needs redundant power and cooling, and the actual physical grounds have additional redundancy concerns. A cloud provider needs multiple and independent power feeds, on top of typically having generator power and battery backups to serve in the interim or in the event that the power feeds become unavailable. Key needs for cloud redundancy are shown in Figure 4-2.

To minimize risk from environmental and natural disaster concerns, a cloud provider should seek out the best possible locations for its data centers, without being bound

LEARN MORE
BUY NOW

Because learning changes everything.®
©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

Figure 4-2 Areas of focus for cloud redundancy

External Redundancy	Internal Redundancy
• Power feeds/lines • Power substations • Generators • Generator fuel tanks • Network circuits • Building access points • Cooling/chilling infrastructure	• Power distribution units • Power feeds to racks • Cooling chillers and units • Networking • Storage units • Physical access points

geographically to a headquarters location or office locations like a typical organization would have to contend with. Because cloud access by definition occurs via networks (and not physical access), as long as a cloud provider has significant and sufficient network bandwidth, the location of a cloud environment can be anywhere in the country or the world to take advantage of cheaper facilities, land, and utilities. Another enormous physical security benefit of a cloud environment comes in the economies of scale with large data centers and the number of customers that leverage them. Sophisticated and redundant levels of security can be very expensive for a data center, but when the costs can be spread among all customers, each customer benefits from far greater and more technologically advanced security than they would be able to afford on their own.

Network and Communications

Networking is essential to a cloud environment because it provides the only way for customers and users to access their systems, applications, and software tools. In the sense of cloud offerings, the network is fully the responsibility of the cloud provider and something that the cloud customer and users will just expect to always work and never have issues.

Networking Hardware

When it comes to building out a network, multiple layers come into play that have their own issues and challenges, even if the customer and users do not really see these aspects. At the basic level are the physical network components such as the actual wiring and cabling. Especially in large data centers, the volume of wiring is extremely high, and often teams are dedicated just to organizing the physical wiring.

Once the physical wiring is in place, it has to be hooked into devices and machines. This forms the next layer of the network in a data center. A large network of switches, routers, and network security devices make up this next level. These are typically constructed in a tiered system that physically segments networks for isolation and security in layers. Segmenting a network physically offers additional security by separating different tiers of servers or restricting traffic within certain sectors. If a successful attack manages to penetrate a data center to the network layer, this physical separation can minimize the extent of vulnerabilities and access.

LEARN MORE BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

Beyond the physical segmenting of a network, software/virtual separation is obtained through such mechanisms as virtual local area networks (VLANs). VLANs allow dedicated IP address spacing for servers that are in the same class or belong to the same application or customer, giving enhanced security and isolation from other systems at the network level. VLANs are not dependent on physical network devices and, as such, can span across data centers, regardless of where hardware is physically located; servers do not need to be in the same racks or even connected to the same switches or routers.

Software-Defined Networking

An important aspect of cloud computing is the use of software-defined networking (SDN). With SDN, the decisions concerning where traffic is filtered or sent and the actual forwarding of traffic are completely separate from each other. With cloud computing, this separation is important because it allows the administrators of the cloud network to quickly and dynamically adjust network flows and resources based on the current needs and demands of the cloud customers. With the separation from the actual network components, a cloud provider can build management tools that allow staffers using web portals or cloud administrative interfaces to make changes to the network without having to log in to the actual network components or needing the command knowledge of a network administrator to make changes. With the level of access provided and the types of resources available to control, a high level of security needs to be attached to any SDN implementation, with access tightly controlled and monitored regularly.

Compute

As with a traditional data center model, cloud computing is built around processing capabilities. Simply put, computing and processing capabilities are defined as the CPU and memory (RAM) of the system and environment. In a traditional server setup using physical servers, it is easy to define and manage both resources because each server represents a finite and unique unit, both in configuration and in the ability to run metrics and observe trends. Within a cloud environment, considering resource pooling and multitenancy, the computing capabilities become far more complex in both planning and management. With large virtual environments, it becomes imperative for the cloud provider to build out enormous resources that can be shared among all the systems, applications, and customers, and done in such a way that each has the resources it requires at any given point in time to meet high availability, performance, and scalability demands.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

Limits

As opposed to reservations, limits are put in place to enforce maximum utilization of memory or processing by a cloud customer. These limits can be done at either a virtual machine level or a comprehensive level for a customer. They are meant to ensure that enormous cloud resources cannot be allocated or consumed by a single host or customer to the detriment of other hosts and customers. Along with cloud computing features such as autoscaling and on-demand self-service, limits can be either "hard" or "fixed," but can also be flexible and allowed to change dynamically. Typically, when limits are allowed to change dynamically based on current conditions and consumption, it is done by "borrowing" additional resources rather than making an actual change in the limits themselves.

Shares

The concept of shares within a cloud environment is used to mitigate and control customer requests for resource allocations in case the environment does not have the current capability to provide these resources. Shares work by prioritizing hosts within a cloud environment through a weighting system that is defined by the cloud provider. When periods of high utilization and allocation are reached, the system automatically uses the scoring of each host based on its share value to determine which hosts get access to the limited resources still available. The higher the value a particular host has, the more resources it will be allowed to utilize.

Storage

Mass storage in a cloud environment from the hardware perspective is not much different than in a traditional data center or server model. Storage typically consists of RAID (redundant array of inexpensive disks) implementations or SANs (storage area networks), which are then connected to the virtualized server infrastructure.

Volume Storage

As covered in Chapter 3, volume storage is where storage is allocated to a virtual machine and configured as a typical hard drive and file system on that server. Although the storage is from a centralized storage system and/or is connected to the network, it will appear to the server as a dedicated resource, much the same as other computing and infrastructure services appear to a virtualized operating system. With a volume storage system, the main infrastructure storage is sliced into pieces called logical units (LUNs), assigned to a particular virtual machine by the hypervisor, and then mounted via a particular method based on the operating system of the host. From the storage allocation perspective, this is only a reserved slice of storage that is given to the virtual machine. All configurations, formatting, usage, and file-system-level security are handled by the particular operating system of the host VM and by the administrators of the host.

Object Storage

As you learned from Chapter 3, object storage is where data is stored on a system separate from the application and access occurs via APIs, network requests, or a web interface. Oftentimes, object storage is implemented as an additional step of redundancy and as a

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

performance measure. By removing the storage from the actual host instances, a cloud provider can focus dedicated resources on managing an object storage system in a way that is specific to optimizing storage performance and security. Object storage also has its own redundancy and scaling systems that are separate from the host and can also be optimized for a particular function or mission.

Rather than a traditional file system with a directory and tree structure, object storage utilizes a flat system and assigns files and objects a key value that is then used to access them. Different implementations of object storage may call this value different names, but in the end it is the same concept—a unique value, oftentimes completely opaque, is used to access data versus using the traditional filename nomenclature. Many cloud providers use object storage for central pieces of infrastructure such as their library of virtual host images.

EXAM TIP Make sure you understand the differences between object and volume storage types that are used with IaaS. There will likely be questions relating to how and when both are used. In particular, make sure you remember that object storage within a cloud will be used for the storage of virtual machine images.

Virtualization

As previously discussed, virtualization forms the backbone of a cloud environment and all its hosting models. Virtualization is what allows a cloud environment to offer most of its top benefits to its customers, especially resource pooling, on-demand self-service, and scalability. The use of virtualization breaks free from the old paradigms and limitations of single servers, where the host is tied to the server. Instead, virtualization allows very large pools of resources to be leveraged across many hosts and applications. Virtualization also allows the abstraction from the hardware via the use of a hypervisor.

Hypervisors

As discussed in Chapter 2 as part of the "Cloud Concepts, Architecture, and Design" domain, there are two types of hypervisors within virtualization: Type 1 and Type 2. An overview of both hypervisor types is shown in Figure 4-3.

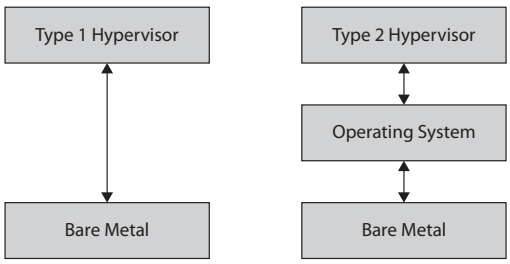

Figure 4-3 Type 1 and Type 2 hypervisors

LEARN MORE
BUY NOW
Because learning changes everything.®
©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

Type 1 Hypervisors As covered in Chapter 2, a Type 1 hypervisor is a native implementation that runs tied directly into the underlying hardware. In other words, it runs natively and directly on top of the hardware with direct access to its components and resources. The Type 1 hypervisor is specifically written and tuned to run on top of bare metal and provide the hosting environment; because of this, it has very tightly written code and is lean overall because it does not have to fulfill any additional requirements. As such, it also allows for much tighter security and controls because there are no additional applications or utilities running within the hypervisor other than those to fulfill its intended mission. Therefore, there are far fewer potential attack vectors and vulnerabilities than a traditional operating system that is designed to be highly flexible would contain.

Type 2 Hypervisors From Chapter 2, you know that a Type 2 hypervisor differs from a Type 1 hypervisor in that it runs under a host operating system rather than directly tied into the underlying hardware of the virtual host servers. With this type of implementation, additional security and architecture concerns come into play as the interaction between the operating system and the hypervisor becomes a critical link. The hypervisor no longer has direct interaction and control over the underlying hardware, which means that some performance will be lost due to the operating system in the middle needing its own resources, patching requirements, and operational oversight. It also means that any security concerns within the underlying operating system can impact the hypervisor as well.

TIP Due to the nature of Type 2 hypervisors and their reliance on the underlying operating system, the Cloud Security Professional needs to be extra vigilant in securing both the hypervisor and the host because of the added complexity. If a cloud provider has robust hypervisor security but is lacking in host security, the entire platform becomes vulnerable and exposed. However, most large and/or public cloud implementations do not use Type 2 hypervisors, so exposure to this type may be well limited. However, care should be taken to know what type is being used.

Management Plane

The concept of a management plane within the context of cloud computing is a bit different from the traditional definition of a network management plane, though the overall concepts are very similar. Within a cloud environment, the management plane is focused on the management of the environment and all the hosts within it. By utilizing the management plane, the cloud provider can manage all the hosts within the environment from a centralized location, without the need to go to each individual server to perform certain tasks. The management plane is typically run from dedicated servers, and it has its own physical connections to the underlying hardware so as to separate out its functions and dependencies from any other aspects of the environment, including the hypervisor.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

The management plane can be used to do the bulk of the tasks that enable cloud computing to be the unique technology that it is. From the management plane, virtual servers can be provisioned with the appropriate resources allocated to them, such as network configurations, processing, memory, and storage. Apart from provisioning and allocating resources, the management plane can also start and stop virtual hosts and services.

The functions of the management plane are typically exposed as a series of remote calls and function executions or exposed as a set of APIs. Those APIs are typically leveraged either through a client or more commonly via a web portal. The web portal is typically proprietary within each cloud implementation, with the appropriate underlying scripts and functions for that environment and the level of exposure that the cloud provider wants to make accessible with the management plane.

Given the access and privileges with which the management plane operates, concerns over security are of the highest level. A compromise of the management plane would give an attacker complete control of the entire environment and make the entire cloud vulnerable, which is well above the risk and threat a compromised hypervisor would provide because the management plane controls multiple hypervisors. Only the most highly vetted and limited set of administrative staff should have access to the management plane, and all access and functions should be tightly audited and reviewed on a regular basis.

Analyze Risks Associated with Cloud Infrastructure

Cloud-based systems have the same level of risk as other hosting models, but with the addition of risks specific to cloud hosting. A cloud-based system should be approached and managed as any other outsourced platform, with the same types of concerns, risks, and audit/governance requirements as an external hosting environment.

Risk Assessment and Analysis

A cloud hosting environment has the same areas of risk as all systems and applications, with cloud-specific risks on top of those risks or as key aspects expanding upon them.

From an organizational and regulatory perspective, there are risks related to lock-in, governance, data security and privacy, and any legal and regulatory controls and reporting required for a system or application. One of the biggest benefits of the cloud hosting model is portability and the ability to move between providers at will. If an organization chooses a particular cloud provider that has a lot of propriety requirements, it may get locked into that provider and incur substantial costs later if it decides to switch.

With any external hosting arrangement, a company loses substantial control over its systems and its governance. Even with strong contractual requirements and SLAs in place, the company's level of control and access will be less than what it would be in its own proprietary data centers. Depending on the regulatory requirements for the type of application and data to be hosted, the choice of cloud provider may be limited, or even nonexistent. Cloud providers have to serve a large number of customers with their business model, which makes complying with several types of certifications and requirements difficult. A primary concern for any company is where its data will be stored and whether sufficient protections are in place to ensure its confidentiality and integrity.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

A cloud environment presents many challenges for appropriate governance, made even more complicated by possible eDiscovery requirements, depending on the nature and type of the data, that many cloud providers might be unable or unwilling to meet. All these aspects need to be carefully evaluated and weighed prior to making a cloud hosting decision.

NOTE eDiscovery is the process through which electronic data is requested or required by a legal entity for use in a criminal or civil legal proceeding. It is the responsibility of the application or system owner to thoroughly analyze and search for all data that is within scope and pertinent to the official request, and then provide it to the legal entity with a certification that it is complete and secured.

Apart from the risk factors that play into any hosting environment, a cloud environment has additional factors that are unique in nature. A major risk in a cloud environment is ensuring that data can be completely removed from the system when necessary or required. As covered in Chapter 2, in a traditional data center, physical media can be destroyed to ensure data destruction, which is not possible in a cloud environment, so concepts such as cryptographic erasure and overwriting are prominently used. Along the same lines of data protection is the security of system images within a cloud environment. Because the images themselves are just files on a file system without any physical separation of servers, combined with the possibility of malware being injecting into an image even when it is not running, their security becomes important in a cloud environment, where the cloud provider bears sole responsibility for assurance.

With the self-service aspects of cloud computing, an enormous amount of software is involved in running a cloud environment—everything from image creation and deployment, to auditing and reporting tools, to user management and key management. Every piece of software throughout the environment—from full application suites down to small utility scripts—carries with it an inherent risk of compromise and vulnerabilities. This software is also removed from the customer's visibility and inspection, and it's solely the responsibility of the cloud provider to ensure its security. Because this is outside the run-time environment of the cloud system in general, in many instances this software will go overlooked in audits and monitoring, and the Cloud Security Professional should apply due diligence when selecting a cloud provider to ensure that they are aware of the risks associated with these tools and have strong auditing and monitoring policies and systems in place. The underlying software running the actual virtual hosts is also certainly susceptible to risk of compromise and vulnerabilities, which will we cover in the following section on virtualization risks. However, the key points include the hypervisor being compromised as well as the potential for guests to break out and access other systems and data in the cloud if security controls are not properly implemented and monitored.

Virtualization Risks

With the added layers and complexities of virtualization also come additional risks not found in a traditional server model.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

Countermeasure Strategies

While the unique challenges and additional risks of cloud computing are well known, so too are the many mitigation strategies that have become commonplace and best practices.

The same high degree of automation that a cloud platform provides between provisioning and reporting can be applied to the implementation of security controls. As systems auto-scale and expand, by using base images for virtual hosts that already have been hardened and scanned, you can ensure that new hosts, as they are brought online, are already secured in the exact same manner as the baseline. The same methodology can be applied within a cloud framework for patching and updating. Rather than performing upgrades and large-scale automated patching, which then require scanning and auditing to ensure everything was applied correctly and comprehensively, a cloud provider can instead opt to reimage the hosts using a new baseline image that has been patched and tested. Instead of thousands of hosts being patched and tested, efforts can be focused on one host and then deployed across the cloud to all others.

Cloud environments are designed to be high availability by nature with redundancy, auto-scaling, and rapid elasticity. This architectural design makes the patching, maintenance, and isolation of hosts in the event of a possible security breach much easier because they can be removed from production pools. It also allows for updating, scanning, and making configuration changes without impacting the customer and users of a system or application, thus reducing the risk to availability.

Design and Plan Security Controls

In order to ensure a sound security policy and overall governance, the Cloud Security Professional should concentrate on several different areas, as detailed in this section.

Physical and Environmental Protection

Physical and environmental protection relates to any and all physical devices and infrastructure components. While the access and technologies used with a cloud infrastructure offer a unique set of services to customers, underneath it is all is a classic data center model, albeit in most cases on a much larger scale. However, because a cloud by definition is a system that is accessible over broad networking, such as the public Internet, physical protections must also extend to those systems that are used to access the cloud.

The physical assets in the actual data center include servers, cooling units, power distribution units, networking gear, physical racks, and miles of cabling, as well as the actual physical facilities and the auxiliary systems located on the premises, such as battery backups, power conduits, generators, fuel tanks, and the surrounding periphery. Outside the data center property there are still more physical devices and infrastructure that are important to the Cloud Security Professional. These include the power and network conduits that the data center relies on, as well as the endpoints of access for the users and customers, such as mobile devices, workstations, laptops, tablets, and any other client systems.

From a physical perspective, the approach to security is no different than in any other system, with defense in layers being the guiding principle. Around the outside of a data center building, you will have typical security measures such as fences, cameras, lights,

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

vehicle access control and barriers, as well as personnel stationed at various locations. Access to the interior of the buildings should be tightly controlled by personnel, doors, key card access, identity proofing, and other various aspects that are similar to how you would gain access to an IT system, including multifactor authentication, in the form of badges, codes, biometrics, and so on. Once inside, access to the actual floor space where the systems reside should be further controlled and restricted by the use of cages and sectioning, as well as employing different levels of access control and monitoring. In some cases, there may be contractual or regulatory requirements to have different types of systems segmented off from others in physical cages, depending on the type of system, the jurisdictions that have authority over it, and the type of data that it stores and consumes.

Any utilities and facilities that the data center depends on, especially in regard to power and cooling, should be redundant both inside and outside the data center. From the outside there should be multiple independent power supplies coming into the data center. Inside the data center, the power distribution units that send power to the racks and servers should also be redundant. This way, whether there is a power failure internally or externally, there are independent redundant power supplies. The same applies to the network, with independent network feeds and redundancy both internally and externally.

Extensive and rigorous background checks should be employed for any personnel allowed access to the center in any capacity. Along the same lines of system security, personnel should be granted physical access based on the least privilege principle. Even with personnel, it is crucial to have proper monitoring and reviews in place, as well as continual training to remind them of policies, procedures, and proper safeguards.

System and Communication Protection

Although the cloud infrastructure is presented to the customer via virtualization, underneath is the same real hardware and systems that would be present in a traditional data center. Depending on the type of cloud hosting model employed, there are varying degrees of customer exposure and responsibilities.

While the cloud provider is responsible for the underlying hardware and network regardless of cloud service model, the remaining services and security responsibilities either lie with the customer or are split between the customer and cloud provider. It is incumbent on the Cloud Security Professional to clearly know the demarcation lines between the customer's and cloud provider's responsibilities, which should be clearly articulated and documented in the contracts and SLAs.

As with any application, the protection of data is a primary concern, and different methods are needed for different states of data:

- **Data at rest** The main protection point for data at rest is through the use of encryption technologies.
- **Data in transit** With data in transit, the main methods of protection are network isolation and the use of encrypted transport mechanisms such as TLS.
- **Data in use** Data in use is protected through secure API calls and web services via the use of encryption, digital signatures, or dedicated network pathways.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

Virtualization Systems Protection

As previously discussed, virtualization forms the backbone of any cloud infrastructure and allows for many of the features that make a cloud environment a unique and popular technology and platform. Given the visible and important role that virtualization plays, the components and systems that make up the virtualized infrastructure are the most obvious and attractive targets for attackers.

The management plane, with full control over the environment and exposed APIs for allowing administrative tasks, is the most visible and important aspect to fully protect. The management plane is made up of a series of APIs, exposed function calls and services, as well as web portals or other client access to allow its use. Any and all of these points are potential vulnerabilities. The Cloud Security Professional needs to develop a holistic picture of threats and vulnerabilities at each level and for each component of the management plane. If the Cloud Security Professional does not break down each component to realize its particular vulnerabilities and threats and then form them together into a comprehensive view, the entire system could be impacted due to the possibility of a missed weakness.

As with any system, role-based access controls are extremely important with the management plane and virtualization infrastructure. All administrative access must be tightly controlled as well as regularly and comprehensively audited. Very detailed logging should take place not only at the level of each piece of the virtualization infrastructure, but also at the level of the web portal or wherever the client accesses the management plane. All logs should be captured in a way where they are removed, indexed, and evaluated away from the actual system itself. This allows log preservation if the actual component is compromised, with sufficient administrative access to modify or delete the logs that are local in nature.

Apart from maintaining cloud features such as auto-scaling and resiliency, a major function of the virtualization environment and infrastructure is to promote and maintain multitenancy. Allowing for multitenancy is not a major challenge from a creation and implementation standpoint because, at its root, multitenancy is a management task versus a technological task. From a virtualization standpoint, virtual machines are just hosts; there is no dependency on what kind of customer or contractual requirements they contain. Where the importance of controls comes into play with multitenancy is the requirement to keep tenants separate and protected from each other. This includes keeping system interaction and security isolated between tenants, as well as resources and utilization, to ensure that all tenants have what their particular systems and applications need to meet contractual and SLA requirements.

While logging and auditing are certainly important to security, these are reactionary mechanisms and do not prevent an actual vulnerability from being exploited initially. Many of the same approaches and strategies that are used at the system level in a traditional data center also are adapted for a virtualization infrastructure in a cloud environment. A prime example is the establishment of trust zones, in the same manner a network separation between servers and systems would be implemented as a defensive layering strategy. Trust zones can be segmented in a variety of ways, depending on the particular needs of the environment or customer contracts. Before implementing trust zones, the cloud provider should undertake a rigorous threat and vulnerability assessment to

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

determine where weaknesses are in its infrastructure and where trust zones could be beneficial. You do not want to take an approach based on practices from other cloud providers without knowing and understanding your own cloud environment first. Going down that path may lead to a false sense of security, to additional complexity where it's not needed, and even to risks and vulnerabilities that were not present before due to the addition of unnecessary components and mechanisms.

Trust zones can be established and segmented using many different strategies. The most common way involves separating out the different tiers of the system architecture (for example, you could separate the web/presentation, application, and data zones of the infrastructure). This allows each zone to have its own protections and monitoring using tools and strategies that are pertinent and specific to its actual needs and functions. It allows the application and data zones to be isolated from external network access and traffic, thus further adding security controls and enhancements.

With this isolation and segmentation, those responsible for managing the systems will need access beyond just what the applications need to operate. Although communications channels will be open internally for the applications to function, that does not facilitate what an administrator needs to gain access, and under no circumstances from a security perspective would you want to open external connectivity for administrative access. The most common ways to allow administrator access in such a configuration is through the use of virtual private networks (VPNs) or jump servers. With a VPN connection, administrators can use most native methods and protocols to access their hosts because they will be secured from their device up to the inside of the cloud environment, and thus not exposed to the public Internet. The VPN connection will have its own authentication and authorization mechanisms, and it will use encrypted communications tunneling, thus adding additional layers of security. With jump servers, the concept is to have a server in the environment that is open and exposed to the public Internet to which administrators can connect and have access internally to the appropriate resources. This method allows security to be focused on the jump server rather than all servers, and it allows the appropriate access controls, monitoring, and logging to be put in place on a much smaller and more specialized scale. The security of the jump server becomes imperative, but it does remove the security concerns directly on the hosts themselves, and it allows the administrators to connect without the need for VPN software and profiles. Another concept that is often used is that of a bastion host. A bastion host is a server that is fully exposed to the public Internet; it is extremely hardened to prevent attacks and is usually focused for a specific application or usage only. This singular focus allows for much more stringent security hardening and monitoring. However, these implementations can, and often are, used together for additional security.

Identification, Authentication, and Authorization in a Cloud Infrastructure

Like applications, cloud systems require identification, authentication, and authorization. However, this need is also extended to include nonperson entities such as devices, clients, and other applications. Federation is another important aspect of cloud

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

computing, especially with public clouds that have a large number of customers and users. It allows the use of "home" or "native" systems to provide identification and authentication, without needing a user base to establish credentials with the cloud provider(s).

Federation

Federation involves a standard base of policies and technologies being employed by different organizations so that they can join their identity systems and allow applications to accept their credentials while still maintaining their autonomy. By establishing policies and guidelines that each member organization must adhere to, trust is established with identities provided by each member organization, and each application that accepts them has the understanding that the sufficient level of background checks and proofing has already been followed and approved under each identity provider. When a system or user who is part of a federation needs to access an application that accepts credentials from that federation, that system or person must obtain local tokens through their own authentication process, which are then passed to the application for acceptance and access. The members that participate run their own "identity providers," and the systems that accept them are known as the "relying party." The typical relationship flow between the user, identity provider, and relying party is shown in Figure 4-4.

Identification

As covered briefly in Chapter 2, identification is the process of singling out an entity—either a person or system/application—in a manner that makes them unique from any other identity. In pretty much all organizations there is already an established identity system, usually based on some form of LDAP. Many organizations, such as academic institutions, small businesses, and nonprofits, tend use open source or other similar identity systems, versus the corporate world where proprietary and commercially supported systems such as Active Directory tend to be the dominant players. With the emergence of large public cloud systems, many organizations and cloud providers have moved toward the OpenID and OAuth standards. Regardless of which particular flavor of identity provider is used by an organization and consumed by the relying party, the vast majority employ the SAML (Security Assertion Markup Language) and WS-Federation standards for passing identity information.

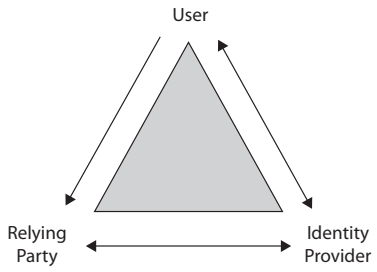

Figure 4-4 The relationship flow within a federated system

LEARN MORE
BUY NOW
Because learning changes everything.®
©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

NOTE Because many large companies and prominent cloud providers have moved toward OpenID, it would be well worth your time doing some additional exploration on the subject so that you have a good understanding of how it works. More information can be found at http://openid.net/.

Authentication

Whereas identification establishes a unique presence of an entity or person, authentication is the process by which one can be certain that the identification presented is true. By policy, this is done to an extent that a system can properly trust the access request. As previously discussed, this can be done through a variety of methods—from the basic user ID/password combination for lower security levels, to strong multifactor authentication, which should be used in all instances possible, but always for administrative and privileged access. The process of authentication is handled by the identity provider.

EXAM TIP Make sure you remember the use of multifactor authentication in all instances where authentication is performed, but it's especially crucial with administrative and privileged account users. You are very likely to see multiple questions about what multifactor authentication is, where it is used, and why it is essential to promoting security best practices.

Authorization

Once an identity has been established and authentication satisfied to the extent required, authorization grants the actual roles and entitlements appropriate for the user or system process to gain access to data and applications. During the authentication process, the identity provider sends certain predetermined attributes about the user to the relying party. The relying party then uses this information (name, location, job title, and so on) to determine the appropriate level and type of access to grant, or whether to grant access at all. The relying party, even in a federated system, makes this determination because it is tied to the actual application, and it makes the decisions based on the policies, requirements, or regulations of the data being accessed.

In making a determination as to whether to grant access and what level of access, the application can take from the identity provider any number of attributes about the entity. The determination can be based on a single attribute, all the way up to complicated conditionals and combinations of attributes.

As an example, take the case of a library system with online journals at a university. In most cases, the sole requirement for access based on the licensing agreements is that the individual be affiliated with the university, and the exact type of affiliation is not important. In this instance, the mere acknowledgment from the identity provider that the person is in fact a valid affiliate of the university satisfies the licensing requirements of the relying party, so access is granted.

On the opposite end of the spectrum, some systems and applications use very complex decision-making algorithms to determine what level of access to grant. For example, an application that a corporation runs may allow immediate access for those on the corporate

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

network but require more extensive validation for those not on the corporate network. As a first conditional step, the relying party can check where the user is connecting from. If the user is on the corporate network, access may be granted immediately. If the user is not on the corporate network, the relying party may require the user to authenticate through the identity provider. Upon successful authentication, the identity provider will provide the relying party with more information about the user to determine access.

No matter which route an entity attempts for access, the principles of least privilege and risk assessment must be adhered to through policy and auditing. Every access to a system is based on a risk assessment for data security, which is established by company policy and the manner in which access is attempted. If the party presents credentials and attributes that align with what has been determined as an acceptable risk to the data, then access may proceed.

Audit Mechanisms

Auditing in the traditional sense involves the ensuring of compliance with policy, guidelines, and regulations. Audits are performed from a security perspective to measure effectiveness at meeting security controls from a variety of sources that together form the entirety of the security requirements for a system or application. In many instances this will include internal security policy requirements, contractual requirements, regulatory requirements from the local, state, or federal government, as well as any industry or organizational requirements such as PCI DSS. Auditing is performed by analyzing security control requirements, meshing those with configurations and internal policies, and then collecting logs, screenshots, data sets, or anything else to prove and show compliance to the satisfaction of management, customers, or independent/government auditors.

In today's IT climate, auditing is also employed for an ever-growing list of reasons beyond the traditional requirements. Audits are used for ongoing security compliance and testing from an internal perspective, and are also increasingly used to provide evidence of value and performance to customers and clients. They can be used far beyond security compliance to show system performance, uptime, user access and load, and virtually any other metric that can be collected and quantified.

Within a cloud environment, auditing presents additional challenges and complexities over a traditional data center and hosting model. Cloud environments are large and distributed hosting platforms, in many instances spanning multiple data centers and even multiple countries and jurisdictions. Many systems are also hosted in a hybrid cloud setup where they span multiple cloud providers and/or hosting models, or are even a combination of cloud and traditional data center hosting configurations. Depending on the cloud model used (IaaS, PaaS, or SaaS), the audit scope will be defined based on the level of access a customer has and the level of information that will be supplied contractually by the cloud provider. The contract and SLAs between the cloud provider and cloud customer should clearly spell out the requirements for auditing and the delineation of responsibilities between both sides, as well as the frequency for any audit testing and reporting. Because a cloud is a multitenant environment, in almost all cases any penetration testing or audit testing must be coordinated between the cloud provider and the cloud customer to ensure that the tests do not conflict with or harm the needs of any other cloud customers hosted in the same environment.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

Because the underlying cloud systems and architecture are controlled by and are the responsibility of the cloud provider, any audit requirements will have to be clearly spelled out in the contract and with support and assistance provided by the cloud provider. With large public clouds especially, hosting hundreds or even thousands of different customers, there is no way for each customer to do their own exhaustive audits in any appreciable sense. The cloud customer will need to rely on audits commissioned on behalf of the cloud provider and have contractual language in place requiring them to be done. Typically, a cloud provider will have the cloud environment audited by a large reputable firm that meets the requirements of the cloud provider's individual tenants and customers, and will provide reports of sufficient detail to meet those audit expectations. A prominent method a cloud provider can use is to have its cloud environment certified under rigorous standards, such as those discussed in Chapter 2. If well-known and established standards are used, the cloud customer can accept the certification as evidence of security control implementations, policies, and ongoing audit requirements being met by the cloud provider.

An advantage a cloud provider has is to leverage its self-service capabilities to provide a set of auditing tools and reports to its customers. Through its self-service portal, the cloud provider can enable customers to access a suite of prebuilt reports, log-collection capabilities, and scripts that provide metrics for a large collection of control testing and data points. This enables the customer to collect audit reports and evidence from these tools whenever needed, without having to engage the cloud provider or auditors to do so.

Another big advantage for audit compliance that a cloud environment enjoys is the use of virtualization and server images. The cloud provider or cloud customer has the ability to construct a base image with all the security controls already implemented and tested, and then use that image to deploy any other hosts within the environment. This drastically cuts down on the time and effort a traditional data center would require for each server to be built and the need to have the baselines applied and tested before the server is deemed ready for use. In a cloud environment using the same identical image, the customer knows that each host instance from the onset is fully compliant and configured correctly to meet their security baseline requirements. The image can also be built to contain monitoring tools or data collection tools as well, so that as hosts are brought online, especially through auto-scaling, these capabilities are already established and running from the start. Having these hooks established on each server plays a big part in having self-service auditing capabilities and continual monitoring capabilities as well, without having to take the time to ensure they are installed, configured, and running properly in each instance.

Log Collection

Log collection within a cloud environment presents additional challenges beyond those in a traditional data center. While the same methods can typically be used for aggregating logs in both scenarios, the overriding challenge in a cloud environment pertains to what logs the customer has access to, with the level of access being highly variable, depending on the type of cloud deployment and contract with the cloud provider.

Within an IaaS environment, a cloud customer will have the most access to logs from an operating system and virtual device level, along with all logs from the platform and

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

application levels. This would require an implemented strategy to collect and aggregate the logs. However, this would not give access to logs from the hypervisor or broader network levels unless provided by the cloud provider through the contract terms.

Within a PaaS environment, the same level of comprehensive logs would be available to the customer at the platform and application levels, but not necessarily at the operating system and network device levels, unless made available by the cloud provider. The same issues are present at the application level within a SaaS environment.

Packet Capture

Packet capture in a cloud environment involves many of the same challenges that log collection presents. The level of packet capture available will be dependent on the level of cloud deployment and control over the necessary network appliances, as well as at what point capture is desired.

If packet capture is needed on the virtual machine, and within an IaaS environment, the cloud customer likely has the level of access they need. Within a PaaS or SaaS environment, though, they would be completely dependent on what services are available from the cloud provider. If packet capture is needed within the network or between network devices, access will be highly variable and dependent on the contract terms. In all instances, if packet capture is needed higher up the network level, or with border routers, the involvement of the cloud provider will be necessary.

Disaster Recovery and Business Continuity Management Planning

A cloud environment represents a natural opportunity for a robust business continuity and disaster recovery (BCDR) program for any organization due to its constructs around resiliency, portability, interoperability, and elasticity. However, the cloud environment also presents its own unique challenges, as we will discuss next.

Understanding the Cloud Environment

A cloud environment can be used for BCDR in a few different types of scenarios, such as hosting for either the primary or the BCDR site as a traditional data center or a cloud environment, or both environments being hosted in cloud environments.

The first scenario is where an organization has its primary computing and hosting capabilities in a traditional data center and uses the cloud to host its BCDR environment. This type of plan would typically revolve around already existing BCDR plans the organization has in place, where the cloud environment takes the place of the failover site should the need arise, versus having a BCDR site at another data center. This scenario leverages the on-demand and metered service aspects of a cloud platform, which makes it a very cost-effective method. With a traditional BCDR plan and a secondary data center, hardware must be procured and available for use, usually in a dedicated fashion, making costs significantly higher and requiring far more substantial prep time. Of course, as we discussed previously, extra care is required in going from a data center model to a cloud model to ensure that all security controls and requirements are being met, and there is no

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

reliance on local security controls and configurations that cannot be easily duplicated or duplicated at all in a cloud environment.

NOTE Apart from the cost benefits of not having hardware standing by and ready at a BCDR site is the benefit of being able to test and configure without having staff onsite. Traditionally, BCDR tests involve staff traveling to the location to configure equipment, but in a cloud environment everything is done via network access. However, do not overlook the fact that in a real emergency, unless staff is geographically dispersed already, some travel may be required if network access is not available at the primary location.

A second scenario is where a system or application is already hosted in a cloud environment, and a separate additional cloud provider is used for the BCDR solution. This would be used in the case of a catastrophic failure at the primary cloud provider, causing the migration of all servers to the secondary cloud provider. This requires the Cloud Security Professional to fully analyze the secondary cloud environment to ensure it has the same or similar security capabilities to satisfy the risk tolerance of the company. Although the system and applications may be portable enough that they do not suffer from vendor lock-in and can easily move between different cloud environments, the secondary cloud environment may offer completely different security settings and controls from the primary environment. There is also the need to ensure that images from one cloud provider can be used by the other cloud provider, or there is additional complexity in preparing and maintaining two sets of images should a sudden disaster occur. As with any BCDR approach, there is the need to have data replicated between the two cloud providers so that the necessary pieces are ready in the event of a sudden disaster. Many times this can be implemented by using the secondary site to back up the primary site.

A third scenario is where an application is hosted in a cloud provider and another hosting model is used within the same cloud provider for BCDR. This is more prevalent with large public clouds that are often divided geographically because it provides resiliency from an outage at one data center of the cloud provider. This setup certainly streamlines configuration and minimizes configuration difficulties for the customer, because both locations within the same cloud provider will have identical configurations, offerings, options, and requirements. This differs in regard to having a BCDR configuration between different cloud providers or data centers in that vendor lock-in is not a prevailing concern.

Understanding Business Requirements

Three big concepts are crucial to determining the business requirements for BCDR, whether implemented with a traditional data center model or a cloud hosting model:

- **Recovery point objective (RPO)** The RPO is defined as the amount of data a company would need to maintain and recover in order to function at a level acceptable to management. This may or may not be a restoration to full operating capacity, depending on what management deems as crucial and essential.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

- **Recovery time objective (RTO)** The RTO is a measurement of the amount of time it would take to recover operations in the event of a disaster to the point where management's objectives for BCDR are met.
- **Recovery service level (RSL)** The RSL measures the percentage of the total, typical production service level that needs to be restored to meet BCDR objectives in the case of a failure.

EXAM TIP Be sure to know the difference between these three concepts and to recognize them by their acronyms.

These three measures are all crucial in making a decision as to what needs to be covered under the BCDR strategy, as well as the approach to take when considering possible BCDR solutions. The prevailing strategy for any company will constitute a cost–benefit analysis between the impact of downtime on business operations or reputation versus the costs of implementing a BCDR solution and to what extent it is done.

Management first needs to determine the appropriate values for RPO and RTO. This step serves as the framework and guidelines for the IT staff and security staff to begin forming a BCDR implementation strategy. These calculations and determinations are completely removed from the possible BCDR solutions and are made strictly from the business requirement and risk tolerance perspectives.

Once management has analyzed and assigned requirements for the RPO and RTO, the organization can set about determining which BCDR solutions are appropriate to meet its needs, weighed against cost and feasibility. While this entire analysis is agnostic of the actual solution, there are some key aspects of cloud computing, spanning multiple areas of concern, that need to be addressed.

A primary concern when it comes to BCDR solutions in the cloud goes back to two main regulatory concerns with cloud hosting in general—where the data is housed and the local laws and jurisdictions that apply to it. This can be a particular concern for those opting for a model with a traditional data center and then using a cloud provider for their BCDR solution, where they are moving into an entirely different paradigm than their production configurations and expectations. However, it also plays prominently in the other two scenarios, because even within the same cloud provider, the other data centers will be in different geographic locations and possibly susceptible to different jurisdictions and regulations, and the same holds true for using a different cloud provider.

Understanding Risks

With any BCDR plan, there are two sets of risks—those that require the execution of the plan in the first place, and those that are realized as a result of the plan itself.

Many risks could require the execution of the BCDR plan, regardless of the particular solution the company has opted to take. These risks include the following:

- Natural disasters (earthquakes, hurricanes, tornadoes, floods, and so on)
- Terrorists attacks, acts of war, or purposeful damage

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

- Equipment failures
- Utility disruptions and failures
- Data center or service provider failures or neglect

Apart from the risks than can lead to the initiation of a BCDR plan, there are also risks associated with the plan that need to be considered and understood:

- **Change in location** Although cloud services are normally accessed over broad networking, a change in geographic hosting location in a BCDR situation can cause network latency or other performance issues. This can impact both the user and customer perspectives when it comes to the system or application, as well as the business owner's ability to update and maintain data, especially if large data updates or transfers are required. Latency can also implement timing issues between servers and clients, especially with many security and encryption systems that rely heavily on time syncing to ensure validity or timeout processes.
- **Maintaining redundancy** With any BCDR plan, a second location will need to be maintained to some extent, depending on the model used and the status and design of the failover site. Both sites will need additional staffing and oversight to ensure they are compatible and maintained to the same level in the event an unforeseen emergency happens without notice.
- **Failover mechanism** In order for a seamless transition to occur between the primary and failover sites, there must be a mechanism in place to facilitate the transfer of services and connectivity to the failover site. This can be done through networking changes, DNS changes, global load balancers, and a variety of other approaches. The mechanism used can involve caching and timeouts that impact the transition period between sites.
- **Bringing services online** Whenever a BCDR situation is declared, a primary issue or concern is the speed at which services can be brought online and made ready at the failover site. With a cloud solution, this typically will be quicker than in a traditional failover site because a cloud provider can take advantage of rapid elasticity and self-service models. If the images and data are properly maintained at the failover site, services online can brought quickly.

CAUTION A common practice is to leave images offline at the BCDR site when not in use. This can cause major problems in the event of a BCDR situation if the system images are not patched and up to date with configurations and baselines for the production systems. If the images are to remain largely offline, the Cloud Security Professional will need to ensure that appropriate processes and verifications are in place with images at the BCDR site.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

- **Functionality with external services** Many modern web applications rely on extensive web service calls out to other applications and services. With a BCDR situation, a crucial concern is ensuring that all hooks and APIs are accessible from the failover site in the same manner and at the same speed as they are with the primary location. If there are keys and licensing requirements to access services from the application, those also must be replicated and made ready at the failover site. If the service has checks in place for the origination of IP addresses or some other tie into actual hosts and locations, the company will need to ensure that the same service can be accessed from the failover site with whatever information necessary already available and configured. Although the failover cloud site may have on-demand and self-service capabilities, it is likely that any external tie-ins will not have the same capabilities and will cause complications while the staff is trying to get their own services up and running.

Disaster Recovery/Business Continuity Strategy

Once management has determined the requirements for the BCDR plan, weighed the appropriate risks and issues, and made the necessary decisions, it is time to create the BCDR plan, the implementation steps and processes for it, and the ongoing testing and validation plan. The continual process of the BCDR strategy is shown in Figure 4-5.

In order to create a BCDR plan for a system, we must take into account all previous material we have covered in this section. The steps for the actual formulation of the plan from the perspective of the Cloud Security Professional are discussed next.

Define Scope

The first step in plan formulation is to ensure that security concerns are an intrinsic part of the plan from the start, rather than trying to retrofit them into the plan after it is developed. This allows security to clearly define the roles and areas of concern during the planning and design phase and to ensure that proper risk assessments are conducted and accepted by management along the way.

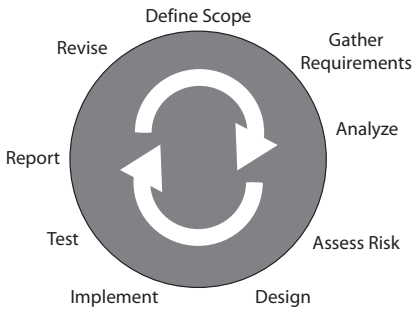

Figure 4-5 The BCDR continual process

LEARN MORE
BUY NOW

Because learning changes everything.®
©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

Gather Requirements

Requirements gathering takes into account the RPO and RTO objectives we previously discussed. This determines what needs to be included in the plan as well as gives a sense of what types and models of solutions are necessary to meet those objectives. From the perspective of the Cloud Security Professional, the determination of critical systems and the time necessary to establish operations in the event of a BCDR situation requires the analysis and application of threats and vulnerabilities that pose a risk to the data and systems. Regulations, laws, policies, customer expectations, and public relations will all play a role in this determination for possible solutions.

Analyze

This step is where the requirements and scope are combined to form the objectives and roadmap for the actual plan design. This step involves a thorough analysis of the current production hosting location for the system or application, a determination of components that need to be replicated for a BCDR situation, and the areas of risk associated with that. With moving to a new environment, new risks are inevitable due to differences in configurations and support models, as well as new hosting staff who do not have a history or familiarity with the application or system. A primary concern with moving to a secondary host provider for BCDR reasons is whether it can handle the load and expectations for the system or application like the primary production hosting arrangement does.

Assess Risk

In any IT system, regardless of hosting circumstances or providers, risk assessments are an ongoing and continual process to ensure security compliance and regulatory requirements. Here's a list of the main risks that are assessed:

- **Load capacity at the BCDR site** Can the site handle the needed load to run the application or system, and is that capacity readily and easily available? Can the BCDR site handle the level of network bandwidth required for the production services and the user community accessing them?
- **Migration of services** Can the BCDR site handle the bandwidth and requirements to get the production services mirrored and configured in a timely manner? If there are huge data sets, will incremental copies be performed regularly versus having to do a single large transfer at the time when services are down? How many services will be on standby and can be rapidly deployed and expanded when needed?
- **Legal and contractual issues** Are appropriate SLAs and contracts in place with the BCDR provider, and do they meet all the requirements the application or system owner must comply with for legal or regulatory reasons? Is all appropriate licensing in place—either system and software licensing or licensing to access external web services or data sources?

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

Design

The design phase is where the actual technical evaluation of BCDR solutions is considered and matched to the company's requirements and policies. In many aspects, the requirements are the same as what procuring a primary production hosting arrangement would entail. Both the technical and support requirements need to be firmly established and articulated in contracts, SLAs, and policies. This includes identifying the owners of each aspect of the system and process, the technical contacts, and the required turnaround times for service expectations. The plan must include the specific requirements that define when the BCDR situation would be declared and the plan put into motion so that the business owner and the hosting environments understand their requirements and what preparations are necessary. Other absolutely crucial parts of the plan are how the testing of the plan will be conducted and how the restoration of services back to the steady production state will be handled once the appropriate time arrives.

Implement the Plan

Once the plan and design have been fully vetted and filled in, the actual implementation of the plan will likely require changes from both technical and policy standpoints.

The actual production applications and hosting may need to be augmented or modified to provide additional hooks or capabilities to enable the BCDR plan to work. This can include system configuration modifications to bring the two hosting environments more in synergy, or providing data replication and configuration replication services to the BCDR host for continual updating and maintaining consistency. From this point forward, the BCDR planning and strategy should be integrated as a key component of ongoing management of IT services and all change and configuration management activities.

NOTE The idea of having to make modifications to existing production platforms in order to implement a BCDR solution is one that is often dismissed outright by management. Often, very good and sound solutions may get overlooked because management is set on finding a solution that can simply be dropped into place without incurring additional costs or modifications to existing systems. As a Cloud Security Professional, be sure to perform a cost–benefit analysis as to the extent of modifications required and the benefits that even minor modifications may bring to the overall organization.

Test the Plan

Once a BCDR plan is established, designed, and implemented, it can only really be considered sound and valid once testing has been performed to ensure its accuracy and feasibility. Also, testing should not be considered a one-time activity. As systems change, new versions of code and configurations are rolled out, new data is added, and any other

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

typical activities are done that denote change over time within a system or application, testing is required as an ongoing activity to ensure the plans are still valid and accurate. Testing essentially turns the plans and designs from theory into reality should a real event occur. It not only confirms planning and expectations, but also serves to provide the staff with experience and familiarity with the overall steps before they are executed in a real emergency situation.

With any test, the primary goal is to ensure that the RPO and RTO goals are achievable with the plans that have been designed and implemented. The testing plan should clearly define the scope of the tests, the applications involved, and the steps taken to replicate or simulate an actual BCDR situation. The tests should involve well-defined and plausible situations that simulate real disaster situations of varying degrees and with multiple variables at play. Tests should be conducted at least once a year, and more frequently if required by any certifications, laws, regulations, or policies. Tests should also be conducted whenever major system or configuration changes are made to ensure their continued validity. It is also crucial to ensure that BCDR validation tests do not impact current production capabilities or services. These tests should be done in a way to minimize any disruption to business operations and staffing.

When developing the test plan, staff should closely follow the objectives and decisions made as part of the RPO and RTO analysis and the steps designed to satisfy them. The test should evaluate each component of the plan to ensure that it is necessary, feasible, and accurate for a real-world situation. Starting at the beginning of the plan and going through each step in sequence, the tests can show any deficiencies or incorrect assumptions, thus providing feedback for augmenting the plan with the discovered issues. Once corrected, the test plans should be run again to test the changes within the scope of how an actual event would unfold.

It is imperative that testing not just be done in a one-way manner. Although testing to ensure that failover to a BCDR site is well documented, planned, and validated is important, it is equally important to fully test the recovery path and restoration of original production services as well. Without the restoration testing, a company could find itself running an alternate hosting arrangement and unsure of how to get back to normal operations, thus further impacting its business operations or data protection and security, as well as incurring potentially substantial additional costs and overhead of both staff and money.

Report and Revise

Once the testing has been completed, a full and comprehensive report detailing all activities, shortcomings, changes made during the course of testing, and the results of the effectiveness of the overall BCDR strategy and plan should be presented to management for review. Management will evaluate the effectiveness of the plan, coupled with the goals and metrics deemed suitable and the costs associated with obtaining such goals. Once management has had a full briefing and the time to evaluate the testing reports, the iterative process can begin anew with changes and modifications to the BCDR plan.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

Exercise

As a senior security officer within your company, you have been tapped to be part of a technical review committee to evaluate new disaster recovery solutions. You currently host your systems in a traditional data center that your company runs, but the committee has been told that system availability and reputation are the top priorities right now, and all options are on the table. Many committee members seem eager to jump to a cloud solution based on what they have heard about the financial benefits alone.

1. From a security perspective, what are the first steps you should advise the committee to consider?
2. Because all options are on the table, how do you formulate a plan to consider a cloud for a BCDR solution, and what impacts might this have on the current hosting model being employed?
3. How do you plan to drive and articulate concerns that must be considered with any cloud solution, weighed against the enthusiasm of others to push toward it based solely on costs?

Chapter Review

In this chapter, we covered the major pieces of the cloud infrastructure and the risks associated with them specifically, and with a cloud environment in general. We explored the various types of security designs and controls that are implemented to counter these risks, as well as the mechanisms and approaches available within a cloud environment to audit them for effectiveness and compliance. Lastly, we covered the very important topic of disaster recovery and business continuity planning, how it is similar and different with a cloud hosting configuration, and how a cloud environment can be used as a core component to a company's BCDR strategy.

Questions

1. Which of the following is the correct order for some of the steps of a BCDR strategy?
   - **A.** Define, Analyze, Design, Assess Risk, Test, Implement
   - **B.** Define, Assess Risk, Analyze, Design, Implement, Test
   - **C.** Define, Design, Analyze, Assess Risk, Test, Implement
   - **D.** Define, Analyze, Assess Risk, Design, Implement, Test
2. What is the entity called that takes the response from the identity provider to allow access to an application?
   - **A.** Relaying party
   - **B.** Relying party
   - **C.** Relaying system
   - **D.** Relying system

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

3. Which of the following storage methods provides a key value to call a file from rather than a directory tree structure?
   - **A.** Volume
   - **B.** Structured
   - **C.** Object
   - **D.** Unstructured
4. Which of the following concepts provides evidence that an entity is in fact who they claim to be?
   - **A.** Authentication
   - **B.** Authorization
   - **C.** Federation
   - **D.** Identification
5. Which of the following would be the *least* beneficial reason to consider a cloud platform as a BCDR solution?
   - **A.** Metered service costs
   - **B.** Hardware ownership
   - **C.** Broad network access
   - **D.** Virtual host replication
6. Which concept involves the prioritization of virtual hosts getting system resources when a cloud is experiencing high utilization and might not be able to serve all hosts?
   - **A.** Reservations
   - **B.** Limits
   - **C.** Shares
   - **D.** Quotas
7. Which of the following is the *most* important factor in defining the controls used to test an audit?
   - **A.** Regulations
   - **B.** Policies
   - **C.** Laws
   - **D.** All of the above
8. What do reservations define within a cloud environment?
   - **A.** Maximum level of resources available for allocation
   - **B.** Guaranteed minimum level of resources available for allocation
   - **C.** Maximum resources available for allocation to all customers
   - **D.** A reserved pool of resources held for emergency load spikes

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

9. What is the *main* objective of software-defined networking (SDN)?

- **A.** Make networking dependent on the operating system of the host and leverage its utilities.
- **B.** Separate the filtering of network traffic and administration from the actual transport of network traffic.
- **C.** Allow different operating systems to seamlessly communicate with each other.
- **D.** Use software to create virtual networks instead of relying on physical network cabling.

10. What is a major security risk with Type 2 hypervisors that does not exist with Type 1 hypervisors?

- **A.** Slower release of security patches
- **B.** Proprietary platform controlled by a single vendor
- **C.** Reliance on a small number of coding platforms
- **D.** Runs on top of another operating system

11. What is the main method for doing patching in a cloud environment?

- **A.** Scripts
- **B.** Host management software
- **C.** Reimaging
- **D.** Customers applying patches on affected hosts

12. Apart from annual testing, when would it be *most* crucial for a BCDR plan to undergo additional testing?

- **A.** During a change in senior management
- **B.** During major configuration changes to an application
- **C.** When new staff is hired
- **D.** During a change in encryption keys

13. What type of storage is the *most* likely to be used for virtual images?

- **A.** Structured
- **B.** Unstructured
- **C.** Volume
- **D.** Object

14. Which of the following issues would be the *greatest* concern from a regulatory standpoint of using a cloud provider for a BCDR solution?

- **A.** Location of stored data
- **B.** Scalability
- **C.** Self-service
- **D.** Interoperability

LEARN MORE BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

15. Which of the following relates to the acceptable duration of recovery to a BCDR location?
    - **A.** RPO
    - **B.** RSL
    - **C.** RDO
    - **D.** RTO

Questions and Answers

1. Which of the following is the correct order for some of the steps of a BCDR strategy?
    - **A.** Define, Analyze, Design, Assess Risk, Test, Implement
    - **B.** Define, Assess Risk, Analyze, Design, Implement, Test
    - **C.** Define, Design, Analyze, Assess Risk, Test, Implement
    - **D.** Define, Analyze, Assess Risk, Design, Implement, Test

    D. Define, Analyze, Assess Risk, Design, Implement, Test are in the correct order; the other options are all incorrect.

2. What is the entity called that takes the response from the identity provider to allow access to an application?
    - **A.** Relaying party
    - **B.** Relying party
    - **C.** Relaying system
    - **D.** Relying system

    B. The relying party takes the authentication tokens from the identity provider and then grants access and authorization based on its own business rules. The other terms and entities listed are not applicable or correct in this instance.

3. Which of the following storage methods provides a key value to call a file from rather than a directory tree structure?
    - **A.** Volume
    - **B.** Structured
    - **C.** Object
    - **D.** Unstructured

    C. Object storage is a flat storage system that resides on external services and references storage items based on a key value rather than a traditional file system and organizational structure. Volume storage is a traditional-type file system that contains directory structures and hierarchical organization as well as uses paths

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

and filenames for access within an Infrastructure as a Service deployment. Structured storage is used with Platform as a Service and is typically a system like a database, which has a predefined structure and organization methodology. Unstructured storage is also used with Platform as a Service and relates to data that does not fit within a predefined structure, such as web files or media objects.

4. Which of the following concepts provides evidence that an entity is in fact who they claim to be?

A. Authentication

B. Authorization

C. Federation

D. Identification

A. Authentication provides proof of the identification of an entity to an acceptable degree of certainty based on policy or regulation. Authorization, done after successful authentication, is the process of granting a user access to data or functions within an application and is based on the role or approved needs of the user. Federation is an authentication system that uses external identity providers that will accept authentication tokens for users, without requiring the user to create an account with the actual application. Identification is part of the authentication process.

5. Which of the following would be the *least* beneficial reason to consider a cloud platform as a BCDR solution?

A. Metered service costs

B. Hardware ownership

C. Broad network access

D. Virtual host replication

B. Hardware ownership would be the least beneficial reason because a cloud customer does not own the hardware; the cloud provider does. Metered service costs are a major benefit of using a cloud provider for BCDR, as the cloud customer would only pay for services when they are needed, unlike traditional BCDR, which typically involves idle hardware sitting in a secondary data center that will likely never be used. Virtual host replication is also a major benefit for a cloud platform and BCDR, as it enables production systems to be regularly mirrored to a secondary location and instantly used, unlike traditional backups, which would have to be recovered on top of another configured system before they can be used. Broad network access would also be highly beneficial for BCDR, as network availability would not be a concern and the ability to access the environment from anywhere on the Internet in case of a disaster would be a major factor.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

6. Which concept involves the prioritization of virtual hosts getting system resources when a cloud is experiencing high utilization and might not be able to serve all hosts?

A. Reservations

B. Limits

C. Shares

D. Quotas

C. The concept of shares is a prioritization and weighting system within a cloud environment that sets that order of specific applications or customers to receive additional resources when requested. Those with the higher prioritization number will receive resources first, and those with lower numbers will receive resources later, or not at all. Reservations involve the setting aside of resources so that cloud customers will be guaranteed a minimum level of resources to start and use their services, even if they cannot obtain additional ones. Limits are the upper bounds set on any level (host, application, customer) that constrain the amount of resources that can be allocated and consumed, in order to protect the overall environment from any entity consuming so many resources that it impacts other customers. Quotas are used by some to mean the same thing as limits, but the latter is the preferred terminology.

7. Which of the following is the *most* important factor in defining the controls used to test an audit?

A. Regulations

B. Policies

C. Laws

D. All of the above

D. All of the above are crucial to a security audit. Regulations, policies, and laws are going to be absolutes and require specific testing and validation, and none can be bypassed during a security controls audit.

8. What do reservations define within a cloud environment?

A. Maximum level of resources available for allocation

B. Guaranteed minimum level of resources available for allocation

C. Maximum resources available for allocation to all customers

D. A reserved pool of resources held for emergency load spikes

B. Reservations define a guaranteed minimum level of resources available to allocate to a host to power on and perform tasks. The maximum level of resources available for allocation would refer to limits, and the maximum resources available for allocation to all customers for the entire cloud environment would be the concern of the cloud provider and play into its overall resource pooling model.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

9. What is the *main* objective of software-defined networking (SDN)?

A. Make networking dependent on the operating system of the host and leverage its utilities.

B. Separate the filtering of network traffic and administration from the actual transport of network traffic.

C. Allow different operating systems to seamlessly communicate with each other.

D. Use software to create virtual networks instead of relying on physical network cabling.

B. The main objective of SDN is to separate the filtering of network traffic and administration from the actual transport of network traffic. This allows management to be performed from portals and API calls, rather than by networking specialists. Toolsets and provisioning systems can access and modify network capabilities that are specific to customer needs, without impacting the underlying actual routing and network transport of packets.

10. What is a major security risk with Type 2 hypervisors that does not exist with Type 1 hypervisors?

A. Slower release of security patches

B. Proprietary platform controlled by a single vendor

C. Reliance on a small number of coding platforms

D. Runs on top of another operating system

D. Running on top of another operating system versus being tied directly to the hardware is a major security risk with Type 2 hypervisors. This makes the hypervisor potentially subjected to any security exploits or issues the underlying operating system may have, as opposed to Type 1 hypervisors, which are tied directly to the hardware and do not rely on security patching and configurations of an external software package.

11. What is the main method for doing patching in a cloud environment?

A. Scripts

B. Host management software

C. Reimaging

D. Customers applying patches on affected hosts

C. Patching in a cloud environment is typically performed by reimaging hosts from the new, fully patched baseline image, rather than deploying patches and doing validations across all the various virtual machines. This allows for consistent and uniform management of patches against a tested and validated image, rather than having to validate on a host-by-host basis to ensure patches are properly received and applied.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

12. Apart from annual testing, when would it be *most* crucial for a BCDR plan to undergo additional testing?

A. During a change in senior management

B. During major configuration changes to an application

C. When new staff is hired

D. During a change in encryption keys

B. Major configuration changes with an application should entail new BCDR testing. Any major configuration change or update represents a significant shift in an environment, and, as such, proper testing is needed to ensure that all BCDR implementations and procedures are both still valid and still work as intended. The changes mentioned in the other answer choices are either minor or personnel changes that would not require new comprehensive testing.

13. What type of storage is the *most* likely to be used for virtual images?

A. Structured

B. Unstructured

C. Volume

D. Object

D. Object storage is the most likely type of storage used for virtual images. Object storage resides externally from specific systems and references each storage object through a key value, which is ideal for system images. System images also do not need any organization structure to them, such as what volume storage would offer. Structured and unstructured would not be appropriate choices as they are geared toward Platform as a Service and are not appropriate for storing system images.

14. Which of the following issues would be the *greatest* concern from a regulatory standpoint of using a cloud provider for a BCDR solution?

A. Location of stored data

B. Scalability

C. Self-service

D. Interoperability

A. Location of stored data would be the most important concern from a regulatory standpoint due to different jurisdictions and requirements. The other choices are all technological or cloud concepts that would not have any bearing on specific regulatory requirements.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

15. Which of the following relates to the acceptable duration of recovery to a BCDR location?

A. RPO

B. RSL

C. RDO

D. RTO

D. RTO, or recovery time objective, relates to the acceptable time for restoration of services. The other choices offered are acronyms that are not applicable here.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

C. Determine the qualifications of the person(s) who will perform the audit.

D. Determine scope, applicability, and purpose for the audit.

☑ **B**. According to ISO/IEC 27005 and other risk management frameworks, it is first necessary to establish the context of an audit. This means making a determination of the scope of the audit—which parts of the organization are to be included. Also, it is necessary to determine the purpose of the risk assessment; for example, determining control coverage, control effectiveness, or business process effectiveness. Finally, the criteria for the audit need to be determined.

☒ **A**, **C**, and **D** are incorrect. **A** and **C** are incorrect because any confirmation of qualifications would be determined prior to this point. **D** is incorrect because an audit that was not applicable should not be performed.

15. A risk manager recently completed a risk assessment in an organization. Executive management asked the risk manager to remove one of the findings from the final report. This removal is an example of what?

A. Gerrymandering

B. Internal politics

C. Risk avoidance

D. Risk acceptance

☑ **D**. Although this is a questionable approach, removal of a risk finding in a report is, implicitly, risk acceptance. It could, however, be even worse than that, and in some industries, this could be considered negligent and a failure of due care. A risk manager should normally object to such an action and may consider documenting the matter or even filing a formal protest.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because the term "gerrymandering" is related to the formation of electoral districts in government. **B** is incorrect because, although the situation may be an example of internal politics, this is not the best answer. **C** is incorrect because risk avoidance is defined as a discontinuation of the activity related to the risk.

16. Which of the following is *not* a risk management methodology?

A. FRAP

B. ISO/IEC 27005

C. NIST Special Publication 800-39

D. FAIR

☑ **D**. FAIR (Factor Analysis of Information Risk) is not a risk management framework, but a risk *assessment* methodology. Though closely related, a risk management framework is concerned with the outcomes of risk assessments, but not the performance of the risk assessments themselves.

☒ **A**, **B**, and **C** are incorrect because FRAP, ISO/IEC 27005, and NIST SP 800-39 are examples of risk management frameworks.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

17. What is the primary objective of the Factor Analysis of Information Risk (FAIR) methodology?

    A. Determine the probability of a threat event.

    B. Determine the impact of a threat event.

    C. Determine the cost of a threat event.

    D. Determine the type of a threat event.

    ☑ **A**. The primary objective of FAIR is to determine the probability of an event using "what if" analysis, which cannot be easily done using maturity models or checklists.

    ☒ **B**, **C**, and **D** are incorrect because FAIR is not used to determine the impact, cost, or type of a threat or threat event.

18. Why might the first control objective of CIS be "Inventory of Authorized and Unauthorized Devices"?

    A. Most organizations are required to have effective asset inventory processes.

    B. The CIS controls framework is hardware asset–centric.

    C. Several IT and security processes depend upon an effective hardware inventory.

    D. The CIS controls framework is an antiquated controls framework.

    ☑ **C**. It is postulated that CIS places hardware asset inventory as its first control because hardware inventory is central to critical processes such as vulnerability management, security event monitoring, and malware prevention and response.

    ☒ **A**, **B**, and **D** are incorrect. **A** is incorrect because this answer is a distractor. **B** and **D** are incorrect because these statements about CIS are untrue.

19. Why is hardware asset inventory critical for the success of security operations?

    A. Critical processes such as software asset and software licensing depends upon accurate asset inventory.

    B. Critical processes such as vulnerability management, event management, and antimalware depend upon accurate asset inventory.

    C. Vulnerability scans need to cover all hardware assets so that all assets are scanned.

    D. Penetration tests need to cover all hardware assets so that all assets are scanned.

    ☑ **B**. Vulnerability management, event visibility, and malware control are among the most critical security operations processes. When these processes are effective, the chances of a successful attack diminish significantly. When asset inventory processes are ineffective, it is possible that there will be assets that are not scanned for vulnerabilities, monitored for events, or protected by antimalware. Intruders are able to identify these assets, which makes asset inventory a critically important activity in information security.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because software inventory, while important for security operations, is not as important as vulnerability management, event management, and malware control. **C** and **D** are incorrect because vulnerability management and penetration tests, while important, are only a portion of critical activities that depend upon effective asset management.

20. What are the most important security-related criteria for system classification?

A. Data sensitivity

B. Data sensitivity and operational criticality

C. Operational criticality

D. Location

☑ **B**. Generally, the operational criticality of a system and the sensitivity of information stored in or processed by the system are the two most important criteria that determine a system's classification.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because data sensitivity alone does not take into account operational criticality. **C** is incorrect because operational criticality alone does not take into account data sensitivity. **D** is incorrect because location alone does not take into account operational criticality or data sensitivity.

21. A new CISO in a financial service organization is working to get asset inventory processes under control. The organization uses on-premises and IaaS-based virtualization services. What approach will most effectively identify all assets in use?

A. Perform discovery scans on all networks.

B. Obtain a list of all assets from the patch management platform.

C. Obtain a list of all assets from the security event and information management (SIEM) system.

D. Count all of the servers in each data center.

☑ **A**. Although none of these approaches is ideal, performing discovery scans on all networks is the best first step. Even so, it will be necessary to consult with network engineers to ensure that discovery scans will scan all known networks in on-premises and IaaS environments. Other helpful steps include interviewing system engineers to understand virtual machine management systems and obtain inventory information from them.

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because patch management systems may not be covering all assets in the organization's environment. **C** is incorrect because the SIEM may not be receiving log data from all assets in the organization's environment. **D** is incorrect because the organization is using virtualization technology, as well as IaaS-based platforms; counting servers in an on-premises data center will fail to discover virtual assets and IaaS-based assets.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

22. Which of the following security-based metrics is most likely to provide value when reported to management?

A. Number of firewall packets dropped per server per day

B. Number of persons who have completed security awareness training

C. Number of phishing messages blocked per month

D. Percent of production servers that have been patched within SLA

☑ **D**. Of the choices listed, this metric will provide the most value and meaning to management, because this helps to reveal the security posture of production servers that support the business.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because the number of packets dropped by the firewall does not provide any business value to management. **B** is incorrect because, although it does provide some value to management, this is not as good an answer as **D**. **C** is incorrect because the number of phishing messages blocked does not provide much business value to management.

23. Ravila, a CISO, reports security-related metrics to executive management. The trend for the past several months for the metric "Percent of patches applied within SLA for servers supporting manufacturing" is 100 percent, 99.5 percent, 100 percent, 100 percent, 99.2 percent, and 74.5 percent. What action should Ravila take with regards to these metrics?

A. Explain that risk levels have dropped correspondingly.

B. No action is required because this is normal for patch management processes.

C. Investigate the cause of the reduction in patching and report to management.

D. Wait until the next month to see if the metric returns to normal.

☑ **C**. As patching is an important activity, and because the servers support critical business operations, this sudden drop in patch coverage needs to be investigated immediately and corrected as quickly as possible.

☒ **A**, **B**, and **D** are incorrect. **A** is incorrect because a reduction in risk levels would not result in a decrease in patching. **B** is incorrect because the reduction in patch coverage is *not* a normal event. **D** is incorrect because it would be unwise to "wait and see" regarding such an important activity as server patching.

24. Duncan is the CISO in a large electric utility. Duncan received an advisory that describes a serious flaw in Intel CPUs that permits an attacker to take control of an affected system. Knowing that much of the utility's industrial control system (ICS) is Intel-based, what should Duncan do next?

A. Report the situation to executive management.

B. Create a new entry in the risk register.

C. Analyze the situation to understand business impact.

D. Declare a security incident.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

☑ **C.** Though it's tempting to notify executive management immediately, without first understanding any potential business impact, there's little to tell. For this reason, the best first step is to analyze the matter so that any business impact can be determined.

☒ **A**, **B**, and **D** are incorrect. **A** is incorrect because the impact is not yet known. **B** is incorrect because it is not the best answer. After understanding the matter, it may indeed be prudent to create a risk register entry, particularly if the matter is complicated and likely to persist for some time. **D** is incorrect because the impact of the advisory on the organization is not yet known. In some incident response plans, however, organizations may use advisories like this as a trigger for emergency analysis to take place.

25. Duncan is the CISO in a large electric utility. Duncan received an advisory that describes a serious flaw in Intel CPUs that permits an attacker to take control of an affected system. After analyzing the advisory, Duncan realizes that many of the ICS devices in the environment are vulnerable. Knowing that much of the utility's industrial control system (ICS) is Intel-based, what should Duncan do next?

A. Create a new entry in the risk register.

B. Report the situation to executive management.

C. Create a new entry in the vulnerability register.

D. Declare a security incident.

☑ **B**. Because the CISO has analyzed the advisory, the impact to the organization can be known. This matter should be reported to executive management, along with an explanation of business impact and a remediation plan.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because this matter has greater urgency than the risk management lifecycle is likely to provide. If, however, it is determined that there is no easy or quick fix, a risk register entry might be warranted. **C** is incorrect because it may be necessary to create many entries instead of a single entry. There may be many different types of devices that are affected by the advisory, necessitating an entry for each time, or an entry for each device, depending upon how the organization manages its vulnerabilities. **D** is incorrect because most organizations' incident response plans do not address vulnerabilities, but actual threat events.

26. An internal audit examination of the employee termination process determined that in 20 percent of employee terminations, one or more terminated employee user accounts were not locked or removed. The internal audit department also found that routine monthly user access reviews identified 100 percent of missed account closures, resulting in those user accounts being closed no more than 60 days after users were terminated. What corrective actions, if any, are warranted?

A. Increase user access review process frequency to twice per week.

B. Increase user access review process frequency to weekly.

C. No action is necessary since monthly user access review process is effective.

D. Improve the user termination process to reduce the number of missed account closures.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

☑ **D**. The rate that user terminations are not performed properly is too high. Increasing the frequency of user access reviews will likely take too much time. The best remedy is to find ways of improving the user termination process. Since the "miss" rate is 20 percent, it is assumed that all processes are manual.

☒ **A**, **B**, and **C** are incorrect. **A** and **B** are incorrect because the user access review process likely takes too much effort. Since the "miss" rate is 20 percent, it is assumed that all processes are manual. **C** is incorrect, since the "miss" rate of 20 percent would be considered too high in most organizations. An acceptable rate would be under 2 percent.

27. To optimize security operations processes, the CISO in an organization wants to establish an asset classification scheme. The organization has no data classification program. How should the CISO proceed?

A. Establish an asset classification scheme based upon operational criticality.

B. Establish an asset classification scheme based upon operational criticality and data classification.

C. First establish a data classification scheme and then an asset classification scheme based on data classification.

D. Treat all assets equally until a data classification program has been established.

☑ **A**. Even in the absence of a data classification program, an asset classification program can be developed. In such a case, asset classification cannot be based on data classification, but assets can be classified according to business operational criticality. For example, assets can be mapped to a business impact analysis (BIA) to determine which assets are the most critical to the business.

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because there is no data classification scheme upon which to base an asset classification scheme. **C** is incorrect because it can take a great deal of time to develop a data classification scheme and map data to assets. It is assumed that the CISO wants to establish the asset classification scheme quickly. **D** is incorrect because there should be an opportunity to classify assets according to operational criticality. If, however, there is little or no sense of business process priority and criticality, then, yes, it might be premature to develop an asset classification scheme.

28. A CISO in a U.S.-based healthcare organization is considering implementation of a data classification program. What criteria should be considered for classifying information?

A. Sensitivity, in scope for HIPAA, in scope for HITECH.

B. Monetary value, operational criticality, sensitivity.

C. Information system, storage, business owner.

D. Data at rest, data in motion, data in transit.

LEARN MORE BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

☑ **B**. Monetary value, operational criticality, and sensitivity are typical considerations for data classification. Some organizations may have additional considerations, such as intellectual property.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because these are not the best criteria. **C** is incorrect because these considerations are not the best criteria. **D** is incorrect because these are not classification considerations, but data-handling use cases.

29. The Good Doctor healthcare organization has initiated its data management program. One of the early activities is a data discovery project to learn about the extent of sensitive data in unstructured data stores. What is the best method for conducting this data discovery?

A. Implement passive DLP tools on servers and endpoints.

B. Implement intrusive DLP tools on servers and endpoints.

C. Manually examine a randomly chosen set of files to see if they contain sensitive data.

D. Run a data discovery tool against file servers and SharePoint servers.

☑ **D**. The best first activity is to run special-purpose data discovery tools against all unstructured data stores such as file servers, SharePoint servers, and cloud provider data stores. This will help the organization better understand the extent of sensitive data in these systems. Results from this activity can be used to determine what next steps are appropriate.

☒ **A**, **B**, and **C** are incorrect. **A** and **B** are incorrect because these are more intrusive and time-consuming options that may or may not be needed. **C** is incorrect because random sampling may miss significant instances, and this option may require excessive time.

30. What is typically the greatest challenge when implementing a data classification program?

A. Difficulty with industry regulators

B. Understanding the types of data in use

C. Training end users on data handling procedures

D. Implementing and tuning DLP agents on servers and endpoints

☑ **C**. The most difficult challenge associated with implementing a data classification program is ensuring that workers understand and are willing to comply with data handling procedures. By comparison, automation is simpler primarily because it is deterministic.

☒ **A**, **B**, and **D** are incorrect. **A** is incorrect because regulators are not typically as concerned with data classification as they are with the protection of relevant information. **B** is incorrect because, although it can be a challenge understanding the data in use in an organization, user compliance is typically the biggest challenge. **D** is incorrect because implementing and tuning agents are not usually as challenging as end user behavior training.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

31. Russ, a security manager at a small online retailer, is completing a self-assessment questionnaire for PCI-DSS compliance. In studying the questionnaire, Russ has noted that his organization is not in compliance with all requirements. No auditor will be verifying the accuracy of the questionnaire. What is Russ's best course of action?

A. Complete the form truthfully and notify senior management of the exceptions.

B. Complete the form truthfully and submit it to authorities.

C. Mark each control as compliant and submit it to authorities.

D. Mark each control as compliant and notify senior management that he must be truthful on the next such submission.

☑ **A**. Security professionals, particularly those who have industry certifications that have a code of conduct (including ISACA's CISM certification), must be truthful, even when there may be personal, professional, or organizational consequences. In this situation, the form must be completed accurately, even though this means that the organization may have some short-term compliance issues with authorities.

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because executive management should also be made aware of the compliance issue. **C** and **D** are incorrect because it would be unethical to falsify answers on the questionnaire.

32. Russ, a security manager at a small online retailer, learned recently about the European General Data Protection Regulation (GDPR). The retailer has customers all over the world. The organization has outsourced its online catalog, order acceptance, and payment functions to a cloud-based e-commerce platform. Russ is unaware of any efforts that the retailer may have made to be compliant with GDPR. What should Russ do about this?

A. Ask senior management or the legal department about this matter.

B. Assume that the organization is compliant with GDPR.

C. Nothing, because the cloud-based e-commerce platform is required to be GDPR compliant.

D. Contact the cloud-based e-commerce platform and confirm its compliance to GDPR.

☑ **A**. A responsible security manager would always reach out to the legal department or another member of senior management to inquire about the organization's state of compliance to a law or regulation.

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because it is unwise to assume that others in an organization have all matters taken care of. **C** is incorrect because the retailer itself must be GDPR compliant, regardless of whether any part of its operations is outsourced. **D** is incorrect because the organization itself must be GDPR compliant. That said, the outsourcing organization must also be GDPR compliant.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

33. Russ, a security leader at a global online retailer, is developing a system classification plan. Systems are classified as High, Moderate, or Low, depending upon operational criticality, data sensitivity, and exposure to threats. In a given environment, how should servers that support (such as DNS servers, time servers) High, Moderate, and Low production servers be classified?

A. Support servers should be classified as High, since some servers they support are High.

B. Support servers should be classified as Low, since they do not perform critical transactions, nor do they contain sensitive data.

C. Support servers should be classified at the same level as the lowest-level servers they support.

D. Support servers should be classified at the same level as the highest-level servers they support.

☑ **D**. The best option is to classify support servers at the same level as the highest-rated servers they support. For instance, if support servers provide support to servers that are rated Medium, then the support servers should be rated as Medium. This will ensure that the support servers are protected (whether for security, resilience, or both) at the same levels as the servers they support.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because the question does not specify the classification level of servers they're supporting. **B** is incorrect because it would be imprudent to classify support servers as Low. It would be better to classify them at the same level as the highest-rated servers they support. **C** is incorrect because the support servers might be supporting higher-rated servers.

34. Russ, a security leader at a global online retailer, is designing a facilities classification plan to provide more consistency and purpose for physical security controls at the organization's worldwide business and processing locations. What criteria should be used to classify facilities for this purpose?

A. Sensitivity of data stored or accessed there

B. Sensitivity of data stored or accessed there and criticality of operations performed there

C. Criticality of operations performed there

D. Size of facilities, and whether there are regulations requiring facilities protection

☑ **B**. Facilities classification is typically established based on two main criteria: sensitivity of information stored at, or accessed at, a location and operational criticality of activities being performed there. For example, a work facility would be classified as High if data classified as High was stored there, or if personnel who worked there routinely accessed data classified as High. A work facility could also be classified as High if critical operations were performed there, such as a hosting facility or a call center.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because facilities classification should be determined by more than just the sensitivity of data stored or accessed there. **C** is incorrect because facilities classification should be based on more than just the criticality of operations performed there. **D** is incorrect because data classification and operational criticality should also be considerations for facilities classification.

35. Which of the following is *not* a valid method for assigning asset value?

A. Net present value

B. Replacement cost

C. Repair cost

D. Book value

☑ **C**. Repair cost is *not* a valid method for assigning asset valuation. Valid methods include replacement cost, book value, net present value, redeployment cost, creation cost, reacquisition cost, and consequential financial cost.

☒ **A**, **B**, and **D** are incorrect. These *are* valid methods for assigning asset value.

36. Dylan is an executive security consultant who is assessing a client organization for compliance to various applicable information security and privacy regulations. Dylan has identified compliance issues and recommends that these issues be documented in the client organization's business. How should these issues be documented?

A. Separate entries for each regulation should be made in the organization's risk register.

B. A single entry should be made in the organization's risk register.

C. Separate entries for each regulation should be made in the organization's security incident log.

D. A single entry should be made in the organization's security incident log.

☑ **B**. The best way to document these findings is to create a single risk register entry for the matter. There could be dozens of similar issues that have common remedies, making it impractical to create potentially dozens of similar entries.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because there could be numerous similar entries that would create unnecessary clutter in the risk register. **C** and **D** are incorrect because the security incident log is not the best place to record this matter.

37. For disaster recovery purposes, why is book value *not* a preferred method for determining the value of assets?

A. Information assets have no book value.

B. Book value may vary based on location if a recovery site is located elsewhere.

C. Some assets may not be tracked for depreciation.

D. The cost to replace damaged or destroyed assets could exceed book value.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

☑ **D**. For disaster recovery purposes, organizations should use replacement or redeployment cost versus book value for asset value. If assets are damaged or destroyed in a disaster, they must be replaced; costs for replacements may be much higher than book value.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because this question is not specifically about information assets. **B** is incorrect because this is not a true statement. **C** is incorrect because this statement is not relevant.

38. A security analyst has identified a critical server that is missing an important security-related operating system patch. What has the security analyst identified?

A. A vulnerability

B. A threat

C. A risk

D. An incident

☑ **A**. The security analysist has identified a vulnerability, which is a weakness that could more easily permit one or more types of threats to occur.

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because the missing patch is not a threat, but a vulnerability that could permit a threat to occur. **C** is incorrect because this is not the best answer. **D** is incorrect because the missing patch is not an incident, although it may permit an incident to occur.

39. A security analyst has identified a new technique that cybercriminals are using to break into server operating systems. What has the security analyst identified?

A. A vulnerability

B. A threat

C. A risk

D. An incident

☑ **B**. The security analysis has identified a threat that, if realized, could result in an intrusion into the organization's systems.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because these techniques are not a vulnerability, but a threat. **C** is incorrect because this is not the best answer. **D** is incorrect because the new technique is not an incident, although it might be possible for an incident to occur because of the threat.

40. Threat actors consist of all of the following *except* which one?

A. Trojans

B. Hacktivists

C. Cybercriminal organizations

D. Employees

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

☑ **A**. Trojans are threats, but they are not threat actors. Threat actors consist of external parties such as hackers, cybercriminal organizations, hacktivists, and more; internal users are also considered threat actors in the context of "insider threat."

☒ **B**, **C**, and **D** are incorrect because hacktivists, employees, and cybercriminals are all considered threat actors.

41. While deliberating an item in an organization's risk register, members of the cybersecurity steering committee have decided that the organization should discontinue a new feature in its online social media platform. This decision is an example of what?

A. Risk transfer

B. Risk acceptance

C. Risk mitigation

D. Risk avoidance

☑ **D**. Risk avoidance is one of four risk treatment options. In risk avoidance, the activity associated with an identified risk is discontinued.

☒ **A**, **B**, and **C** are incorrect. Risk acceptance, risk mitigation, and risk transfer are not the correct terms associated with the organization's decision to discontinue the business activity discussed here.

42. NotPetya is an example of what?

A. Threat

B. Spyware

C. Mass-mailing worm

D. Password-cracking tool

☑ **A**. NotPetya is a threat. More specifically, NotPetya is malware that resembles ransomware but lacks the ability to decrypt data; thus, it is considered by many to be destructware, or software that destroys data files.

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because NotPetya is not spyware. **C** is incorrect because NotPetya is not a mass-mailing worm. **D** is incorrect because NotPetya is not a password cracker.

43. Randi, a security architect, is seeking ways to improve a defense-in-depth to defend against ransomware. Randi's organization employs advanced antimalware on all endpoints and antivirus software on its e-mail servers. Endpoints also have an IPS capability that functions while endpoints are onsite or remote. What other solutions should Randi consider to improve defenses against ransomware?

A. Data replication

B. Spam and phishing e-mail filtering

C. File integrity monitoring

D. Firewalls

LEARN MORE BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

☑ **B**. The next solution that should be considered is a solution that will block all incoming spam and phishing e-mail messages from reaching end users. This will provide a better defense-in-depth for ransomware since several other good controls are in place.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because data replication is not an adequate defense against ransomware, because files encrypted by ransomware are likely to be replicated onto backup file stores. Instead, offline backup such as magnetic tape or e-vaulting should be used. **C** is incorrect because file integrity monitoring (FIM) is generally not chosen as a defense against ransomware. **D** is incorrect because firewalls are not an effective defense against ransomware, unless they also have an IPS component that can detect and block command-and-control traffic.

44. Which European law enforces users' rights to privacy?

A. GLBA

B. GDPR

C. 95/46/EC

D. SB-1386

☑ **B**. GDPR, or the European General Data Protection Regulation, which took effect in 2018, provides several means to improve privacy for European residents.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because GLBA is a U.S. law that requires financial services organizations to protect information about its customers. **C** is incorrect because 95/46/EC, otherwise known as the European Privacy Directive, is the former European privacy law that has been superseded by GDPR. **D** is incorrect because SB-1386 is the original data breach disclosure law in the state of California.

45. Which mechanism does GDPR provide for multinational organizations to make internal transfers of PII?

A. Model clauses

B. Privacy Shield

C. Safe Harbor

D. Binding corporate rules

☑ **D**. Binding corporate rules were established by European privacy laws that permit multinational organizations to perform internal transfers of sensitive information. Typically this is applied to internal human resources information.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because model clauses are used between organizations to legally obligate them to comply with GDPR and other privacy regulations. **B** is incorrect because Privacy Shield is used by organizations to register their obligation to comply with GDPR. **C** is incorrect because Safe Harbor is the now-defunct means for organizations to register their obligation to comply with the former European privacy directive, 95/46/EC.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

46. Which mechanism provides the legal framework for the transfer of information from Europe to the United States?

A. Model clauses

B. Privacy Shield

C. Safe Harbor

D. Binding corporate rules

☑ **B**. The E.U.-U.S. Privacy Shield is the new legal framework for regulating the flow of information from Europe to the United States. Privacy Shield supersedes Safe Harbor, which was invalidated in 2015.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because model clauses are a set of legal language used in legal agreements between organizations regarding the protection of PII of European residents. **C** is incorrect because Safe Harbor was invalidated in 2015. **D** is incorrect as binding corporate rules are used for the internal transfer of PII within a multinational organization.

47. What language is used in legal agreements between organizations regarding the protection of personally identifiable information?

A. Model clauses

B. Privacy Shield

C. Safe Harbor

D. Binding corporate rules

☑ **A**. Model clauses are used in legal contracts between organizations regarding the protection of PII of European citizens. Model clauses are a set of specific language included in privacy regulations such as the former European Privacy Directive and the current Global Data Privacy Regulation (GDPR).

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because Privacy Shield is a legal framework for the protection of PII, but it does not include language used in contracts between organizations. **C** is incorrect because Safe Harbor is the former legal framework that is superseded by Privacy Shield. **D** is incorrect because binding corporate rules are the legal framework for the internal transfer of sensitive information in multinational companies.

48. Which mechanism was formally used as the legal framework for the transfer of information from Europe to the United States?

A. Model clauses

B. Privacy Shield

C. Safe Harbor

D. Binding corporate rules

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

☑ **C**. International Safe Harbor Privacy Principles, known primarily as Safe Harbor, is the former framework for the legal transfer of European PII to the United States. Safe Harbor was invalidated in 2015 by the European Court of Justice.

☒ **A**, **B**, and **D** are incorrect. **A** is incorrect because model clauses are legal agreement templates used for agreements between organizations. **B** is incorrect because Privacy Shield is the functional replacement for Safe Harbor. **D** is incorrect because binding corporate rules are used in the context of intracompany data transfers of PII.

49. The internal audit department in a public company recently audited key controls in the vulnerability management process and found that the control "Production servers will be patched within 30 days of receipt of critical patches" fails 30 percent of the time. What finding should the internal audit make?

A. A new control is needed for vulnerability management.

B. The control is ineffective and needs to be corrected.

C. The control should be changed from 30 days to 45 days.

D. The control should be changed from 30 days to 21 days.

☑ **B**. There is a control in place that is not effective. The best remedy is to fix the existing control, which is still reasonable and appropriate.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because creating an additional control should not be considered until the existing control is fixed. **C** and **D** are incorrect because the SLA for critical patches does not necessarily need to be changed.

50. The internal audit department in an organization recently audited the control "User accounts for terminated workers shall be locked or removed within 48 hours of termination" and found that user accounts for terminated workers are not locked or removed 20 percent of the time. What recommendation should internal audit make?

A. Change the timeframe in the control from 48 hours to 7 days.

B. Add a new compensating control for monthly review of terminated user accounts.

C. Add more staff to the team that manages user accounts.

D. No changes are needed since 20 percent is an acceptable failure rate.

☑ **B**. A compensating control in the form of a periodic access review is the best answer. Periodic access reviews are common and used for this purpose.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because seven days is far too long for user accounts to be active after a worker is terminated. **C** is incorrect because staffing levels are not necessarily the cause of this control failure. **D** is incorrect because 20 percent is considered too high a failure rate for a terminated user account access control.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

51. Upon examining the change control process in a SaaS provider organization, a new security manager has discovered that the change control process lacks a security impact procedure. What should the security management recommend for this matter?

A. Systems impacted by a change should be scanned before and after changes are made.

B. A post-change security review should be added to the change control process.

C. No change is needed because security is not needed in change control processes.

D. Add a security impact procedure to the change control process so that the security impact of each proposed change can be identified.

☑ **D**. The best remedy is the addition of a security impact procedure that is performed for each proposed change. This will help to identify any security-related issues associated with a proposed change that can be discussed prior to the change being made. This is preferable to the alternative: accepting a change that may have one or more security issues that may increase the risk of a security incident.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because not all security-related issues will be manifested in a vulnerability scan. **B** is incorrect because a security review should be performed prior to a change being made so that an organization can consider modifying the nature of the change so that there is no increase in risk. **C** is incorrect because security *is* an important consideration in a change control process.

52. A SaaS provider performs penetration tests on its services once per year, and many findings are identified each time. The organization's CISO wants to make changes so that penetration test results will improve. The CISO should recommend all of the following changes *except* which one?

A. Add a security review of all proposed software changes into the SDLC.

B. Introduce safe coding training for all software developers.

C. Increase the frequency of penetration tests from annually to quarterly.

D. Add the inclusion of security and privacy requirements into the SDLC.

☑ **C**. Increasing the frequency of penetration tests is not likely to get to the root cause of the problem, which is the creation of too many security-related software defects.

☒ **A**, **B**, and **D** are incorrect. **A** is incorrect because the addition of a security review for proposed changes is likely to reveal issues that can be corrected prior to development. **B** is incorrect because safe coding training can help developers better understand coding practices that will result in fewer security defects. **D** is incorrect because the addition of security and privacy requirements will help better define the nature of new and changed features.

LEARN MORE BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

☑ **C**. International Safe Harbor Privacy Principles, known primarily as Safe Harbor, is the former framework for the legal transfer of European PII to the United States. Safe Harbor was invalidated in 2015 by the European Court of Justice.

☒ **A**, **B**, and **D** are incorrect. **A** is incorrect because model clauses are legal agreement templates used for agreements between organizations. **B** is incorrect because Privacy Shield is the functional replacement for Safe Harbor. **D** is incorrect because binding corporate rules are used in the context of intracompany data transfers of PII.

49. The internal audit department in a public company recently audited key controls in the vulnerability management process and found that the control "Production servers will be patched within 30 days of receipt of critical patches" fails 30 percent of the time. What finding should the internal audit make?

A. A new control is needed for vulnerability management.

B. The control is ineffective and needs to be corrected.

C. The control should be changed from 30 days to 45 days.

D. The control should be changed from 30 days to 21 days.

☑ **B**. There is a control in place that is not effective. The best remedy is to fix the existing control, which is still reasonable and appropriate.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because creating an additional control should not be considered until the existing control is fixed. **C** and **D** are incorrect because the SLA for critical patches does not necessarily need to be changed.

50. The internal audit department in an organization recently audited the control "User accounts for terminated workers shall be locked or removed within 48 hours of termination" and found that user accounts for terminated workers are not locked or removed 20 percent of the time. What recommendation should internal audit make?

A. Change the timeframe in the control from 48 hours to 7 days.

B. Add a new compensating control for monthly review of terminated user accounts.

C. Add more staff to the team that manages user accounts.

D. No changes are needed since 20 percent is an acceptable failure rate.

☑ **B**. A compensating control in the form of a periodic access review is the best answer. Periodic access reviews are common and used for this purpose.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because seven days is far too long for user accounts to be active after a worker is terminated. **C** is incorrect because staffing levels are not necessarily the cause of this control failure. **D** is incorrect because 20 percent is considered too high a failure rate for a terminated user account access control.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

51. Upon examining the change control process in a SaaS provider organization, a new security manager has discovered that the change control process lacks a security impact procedure. What should the security management recommend for this matter?

A. Systems impacted by a change should be scanned before and after changes are made.

B. A post-change security review should be added to the change control process.

C. No change is needed because security is not needed in change control processes.

D. Add a security impact procedure to the change control process so that the security impact of each proposed change can be identified.

☑ **D**. The best remedy is the addition of a security impact procedure that is performed for each proposed change. This will help to identify any security-related issues associated with a proposed change that can be discussed prior to the change being made. This is preferable to the alternative: accepting a change that may have one or more security issues that may increase the risk of a security incident.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because not all security-related issues will be manifested in a vulnerability scan. **B** is incorrect because a security review should be performed prior to a change being made so that an organization can consider modifying the nature of the change so that there is no increase in risk. **C** is incorrect because security *is* an important consideration in a change control process.

52. A SaaS provider performs penetration tests on its services once per year, and many findings are identified each time. The organization's CISO wants to make changes so that penetration test results will improve. The CISO should recommend all of the following changes *except* which one?

A. Add a security review of all proposed software changes into the SDLC.

B. Introduce safe coding training for all software developers.

C. Increase the frequency of penetration tests from annually to quarterly.

D. Add the inclusion of security and privacy requirements into the SDLC.

☑ **C**. Increasing the frequency of penetration tests is not likely to get to the root cause of the problem, which is the creation of too many security-related software defects.

☒ **A**, **B**, and **D** are incorrect. **A** is incorrect because the addition of a security review for proposed changes is likely to reveal issues that can be corrected prior to development. **B** is incorrect because safe coding training can help developers better understand coding practices that will result in fewer security defects. **D** is incorrect because the addition of security and privacy requirements will help better define the nature of new and changed features.

LEARN MORE BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

53. A SaaS provider performs penetration tests on its services once per year, and many findings are identified each time. What is the best way to report this matter to executive management?

A. Develop a KRI that reports the trend of security defects over time.

B. Penetration test reports should be distributed to executive management so that they can have a better understanding of the problem.

C. The executive summary section of penetration test reports should be distributed to executive management.

D. Report the number of defects found to executive management.

☑ **A**. A key risk indicator (KRI) should be developed that illustrates the risk that security defects make on the organization. An example KRI for this situation could read, "Number of critical software defects introduced into SAAS Product."

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because penetration test reports are quite detailed and technical, and they provide little, if any, business insight to an executive. **C** is incorrect because even an executive summary section in a penetration test report is unlikely to express business risk in a meaningful way. **D** is incorrect because the number of defects alone is not a good risk indicator.

54. A SaaS provider performs penetration tests on its services once per year, and many findings are identified each time. What is the best KRI that would highlight risks to executives?

A. Number of software vulnerabilities that exist on production SaaS applications

B. Number of days that critical software vulnerabilities exist on production SaaS applications

C. Number of vulnerability scans performed on production SaaS applications

D. Names of developers who introduced the greatest number of security defects onto production SaaS applications

☑ **B**. The total number of days that unmitigated software defects existed on production applications is the best risk indicator, particularly when tracked over a period of time.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because the number of vulnerabilities alone does not sufficiently convey risk; a better depiction of risk is the number of days that unpatched vulnerabilities were present on production systems. **C** is incorrect because the number of scans does not provide an indication of risk. **D** is incorrect because a list of offenders is not a key risk indicator.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

55. The security leader at a SaaS provider has noticed that the number of security defects in the SaaS application is gradually climbing over time to unacceptable levels. What is the best first step the security leader should take?

A. Contact the software development leader and report that more security defects are being created.

B. Initiate the procurement process for a web application firewall.

C. Initiate a low-severity security incident.

D. Create a new risk register entry that describes the problem along with potential fixes.

☑ **D**. When there is a disturbing trend developing, such as an increase in the number of security vulnerabilities being identified, creating an entry in the risk register is the best first step. This will facilitate action in the organization's risk management process that will enable business and technology leaders to discuss the matter and make decisions to manage the risk.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because this is not the best first choice. Contacting the development leader is, however, a prudent move so that the development leader will not feel blindsided by later proceedings. **B** is incorrect because a WAF may not be the best solution here; besides, this represents a unilateral decision on the part of the security leader, when a better approach would be a discussion with stakeholders. **C** is incorrect because a situation like this is not commonly regarded as a security incident.

56. Why is the KRI "Number of days that critical software vulnerabilities exist on production SaaS applications" considered a leading risk indicator?

A. This is the first KRI that executives are likely to pay attention to.

B. This KRI provides a depiction of the probability of a security incident through the exploitation of vulnerabilities. The risk of an incident is elevated with each successive day that unpatched vulnerabilities exist.

C. Critical software vulnerabilities are the leading cause of security incidents.

D. The KRI indicates that critical software vulnerabilities are the most likely cause of a future incident.

☑ **B**. A KRI is a leading risk indicator because it portends the likelihood of a future event. The KRI in this question points to the likelihood of a security breach that occurs through the exploitation of a defect in an organization's Internet-facing software application.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because leading risk indicators are so-named because they help predict the likelihood of future events. **C** is incorrect because the meaning of a leading risk indicator is related to the likelihood of a specific future event. The fact that the KRI in this question is related to a leading cause of incidents is coincidental. **D** is incorrect because the KRI does not attempt to identify the most likely cause of a future incident.

LEARN MORE BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

57. Which is the best method for reporting risk matters to senior management?

A. Sending after-action reviews of security incidents

B. Sending the outcomes of risk treatment decisions

C. Periodic briefing on the contents of the risk register

D. Sending memos each time a new risk is identified

☑ **C.** The best method available here is to provide a summary briefing on the contents of the risk register. Providing a summary overview of the items of the risk register will enable the leadership team to focus on the key areas or emerging risks that need their attention. This will help senior management better understand the entire catalog of unmanaged risks in the organization.

☒ **A**, **B**, and **D** are incorrect. **A** is incorrect because risks often exist, apart from security incidents. **B** is incorrect because senior management should participate in risk treatment decisions, not merely be informed about them (implying that others are making those decisions). **D** is incorrect because sending memos is unstructured, and memos may not always be read. Further, a briefing from the risk register is much better, because this is an interactive event where senior management can ask questions about risks in the risk register.

58. Janice has worked in the Telco Company for many years and is now the CISO. For several years, Janice has recognized that the engineering organization contacts information security just prior to the release of new products and features so that security can be added in at the end. Now that Janice is the CISO, what is the best long-range solution to this problem?

A. Introduce security at the conceptual, requirements, and design steps in the product development process.

B. Train engineering in the use of vulnerability scanning tools so that they can find and fix vulnerabilities on their own.

C. Add security requirements to other requirements that are developed in product development projects.

D. There is no problem to fix: it is appropriate for engineering to contact security prior to product release to add in necessary security controls.

☑ **A**. The best long-term solution is the introduction of appropriate security activities throughout the product development lifecycle, starting at the conceptual stage where new products and features are initially discussed. Security steps at the requirements and design stages will help ensure that products are secure by design.

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because vulnerability scanning will fail to identify many types of security problems. **C** is incorrect because adding security requirements alone, while helpful, is not the best choice. **D** is incorrect because responsible organizations ensure that their products are secure by design.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

59. Janice has worked in the Telco Company for many years and is now the CISO. For several years, Janice has recognized that the engineering organization contacts information security just prior to the release of new products and features so that security can be added in at the end. Now that Janice is the CISO, what is the best first step for Janice to take?

A. Initiate a low-severity security incident.

B. Create a new risk register entry that describes the problem along with potential fixes.

C. Initiate a high-severity security incident.

D. Write a memo to the leader of the engineering organization requesting that security be added to the product development lifecycle.

☑ **B**. Creation of a risk register entry is the best first step. Presuming that a cross-functional cybersecurity council exists, the next step will be discussion of the matter that will lead to an eventual decision.

☒ **A**, **C**, and **D** are incorrect. **A** and **C** are incorrect because initiation of a security incident is not an appropriate response. **D** is incorrect because a wider conversation should be conducted by cybersecurity steering committee members.

60. The term "insider threat" includes all of the following *except* which one?

A. End users who are ignorant and make unwise decisions

B. Employees who have a grudge against their employer

C. Customers who attempt to break into systems while onsite

D. End users who are doing the right thing but make mistakes

☑ **C**. Customers, even while onsite, are not considered insiders.

☒ **A**, **B**, and **D** are incorrect. Each of these is considered an insider threat.

61. Examples of employees gone rogue include all of the following *except* which one?

A. A developer who inserts a time bomb in application source code

B. A securities trader who makes unauthorized trades resulting in huge losses

C. An engineer who locks co-workers out of the network because they are not competent

D. A systems engineer who applies security patches that cause applications to malfunction

☑ **D**. The systems engineer who applies patches to fix feature or security defects is the best choice, because there is little or no sign of malice. In this example, the change control process should be improved so that there is an opportunity to test software applications in a nonproduction environment prior to applying patches to production.

☒ **A**, **B**, and **C** are incorrect. Each of these is an example of an employee who has gone rogue and is consequently harming the organization.

LEARN MORE BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

62. Janice, a new CISO in a healthcare delivery organization, has discovered that virtually all employees are local administrators on their laptop/desktop computers. This is an example of what?

A. Insider threat

B. Vulnerability

C. Threat

D. Incident

☑ **B**. The matter of end users being local administrators means that they have administrative control of the computers they use, namely their laptop and/or desktop computers. This means they can install software and security patches and change the configuration of the operating system. This also means that malware introduced by the user onto the system will probably be able to run with administrative privileges, which may result in significantly more harm to the system and the organization.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because this configuration setting is not, by itself, an insider threat. However, an insider threat situation can be made worse through end users having local administrative privileges. **C** is incorrect because this is not a threat, but a vulnerability (these terms are often misused). **D** is incorrect because this is not an incident. However, an incident is somewhat more likely to occur and more likely to have greater impact because end users have local administrative privileges.

63. An end user in an organization opened an attachment in e-mail, which resulted in ransomware running on the end user's workstation. This is an example of what?

A. Incident

B. Vulnerability

C. Threat

D. Insider threat

☑ **A**. Ransomware executing on an end user's workstation is considered an incident. It may have been allowed to execute because of one or more vulnerabilities.

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because a vulnerability is a configuration setting or a software defect that can, if exploited, result in an incident. **C** is incorrect because ransomware, by itself, is considered a threat, but ransomware executing on a system is considered an incident. **D** is incorrect because this is not considered an insider threat. However, users having poor judgment (which may include clicking on phishing messages) is considered an insider threat.

64. What is the purpose of the third-party risk management process?

A. Identify risks that can be transferred to third parties.

B. Identify a party responsible for a security breach.

C. Identify a party that can perform risk assessments.

D. Identify and treat risks associated with the use of third-party services.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

☑ **D**. Third-party risk management encompasses processes and procedures for identifying risks associated with third-party service providers and suppliers; assessments of third parties enable management to make decisions regarding whether to do business with specific third parties and under what conditions.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because third-party risk management is not related to risk transfer. **B** is incorrect because third-party risk management is not involved in security breach response and investigation. **C** is incorrect because third-party risk management is not related to the process of performing internal risk assessments.

65. What is the correct sequence of events when onboarding a third-party service provider?

A. Contract negotiation, examine services, identify risks, risk treatment

B. Examine services, identify risks, risk treatment, contract negotiation

C. Examine services, contract negotiation, identify risks, risk treatment

D. Examine services, identify risks, risk treatment

☑ **B**. The best sequence here is to examine the services offered by the third party, identify risks associated with doing service with the third party, make decisions about what to do about these risks, and enter into contract negotiations.

☒ **A**, **C**, and **D** are incorrect. **A** and **C** are incorrect because contract negotiation should not take place prior to identifying risks that may need to be addressed in a contract. **D** is incorrect because contract negotiation is not included.

66. A campaign by a cybercriminal to perform reconnaissance on a target organization and develop specialized tools to build a long-term presence in the organization's environment is known as what?

A. Watering hole attack

B. Hacktivism

C. Advanced persistent campaign (APC)

D. Advanced persistent threat (APT)

☑ **D**. A long-term campaign of patient reconnaissance, development of tools, and establishment of a long-term quiet presence inside an organization's environment is known as an advanced persistent threat (APT). It is "advanced" on account of the reconnaissance and the development of an intrusion strategy with specialized tools; it is "persistent" by design, so that the intruder can maintain a long-term presence in the environment; it is a "threat" because the criminal actor is performing all of this to reach a long-term objective, whether the acquisition or destruction of sensitive information or the disruption of the organization's operations.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because a watering hole attack is an attack on an organization via a compromised website that will automatically download malware onto visitors' systems. **B** is incorrect because hacktivism refers to an ideology wherein an attacker seeks to expose or disrupt an organization for ideological reasons. **C** is incorrect because the term "advanced persistent campaign" is not in use.

67. Joel, a CISO in a manufacturing company, has identified a new cybersecurity-related risk to the business and is discussing it privately with the chief risk officer (CRO). The CRO has asked Joel not to put this risk in the risk register. What form of risk treatment does this represent?

A. This is not risk treatment, but the avoidance of managing the risk altogether.

B. This is risk avoidance, where the organization elects to avoid the risk altogether.

C. This is risk transfer, as the organization has implicitly transferred this risk to insurance.

D. This is risk acceptance, as the organization is accepting the risk as-is.

☑ **A**. The deliberate "burying" of a risk is not risk treatment, but the refusal to deal with the risk altogether. Although there may be legitimate reasons for this action, based on the information here, there is an appearance of negligence on the part of the CRO.

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because risk avoidance is a formal decision wherein the organization will discontinue the activity that manifests the identified risk. **C** is incorrect because there is no indication in this question that cyber insurance will assume this risk. **D** is incorrect because formal risk acceptance involves the use of the risk management lifecycle that includes the risk being recorded in the risk ledger, followed by analysis and a risk treatment decision.

68. Which of the following factors in risk analysis is the most difficult to determine?

A. Exposure factor

B. Single-loss expectancy

C. Event probability

D. Event impact

☑ **C**. Event probability is the most difficult of these values to determine accurately, particularly for high-impact events. Because event probability is so difficult to determine, much risk analysis work performed is qualitative in nature.

☒ **A**, **B**, and **D** are incorrect. **A** is incorrect because exposure factor (which is calculated as a percentage of an asset's value) is relatively easy to determine. **B** is incorrect because single-loss expectancy (which is calculated as asset value times exposure factor) is relatively easy to determine. **D** is incorrect because event impact (formally known as event cost) is not altogether difficult to determine.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

69. An estimate on the number of times that a threat might occur in a given year is known as what?

A. Annualized loss expectancy (ALE)

B. Annualized rate of occurrence (ARO)

C. Exposure factor (EF)

D. Annualized exposure factor (AEF)

☑ **B**. Annualized rate of occurrence (ARO) is defined as an estimate of the number of times that a threat will occur per year.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because annualized loss expectancy (ALE) is defined as the annualized rate of occurrence (ARO) times the single loss expectancy (SLE). **C** is incorrect as exposure factor (EF) is the loss that represents a percentage of an asset's value (because in some cases, an asset is not completely destroyed). **D** is incorrect because there is no such term is annualized exposure factor (AEF).

70. Which is the best method for prioritizing risks and risk treatment?

A. Threat event probability times asset value, from highest to lowest

B. Threat event probability, followed by asset value

C. Professional judgment

D. A combination of threat event probability, asset value, and professional judgment

☑ **D**. The best method for prioritizing risks and risk treatment is to examine the probability of event occurrence (difficult though that may be), asset value, and impact to the organization. Professional judgment plays a big role as well because factors such as business reputation are difficult to quantify.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because this approach allows no room for professional judgment. **B** is incorrect because there is no logical sequence based on these two items that are measured differently. **C** is incorrect because professional judgment alone risks the failure to consider high-value assets, high impact, and high probability of occurrence.

71. Joel is a security manager in a large manufacturing company. The company uses primarily Microsoft, Cisco, and Oracle products. Joel subscribes to security bulletins from these three vendors. Which of the following statements best describes the adequacy of these advisory sources?

A. Joel should also subscribe to nonvendor security sources such as US-CERT and InfraGard.

B. Joel's security advisory sources are adequate.

C. Joel should discontinue vendor sources and subscribe to nonvendor security sources such as US-CERT and InfraGard.

D. Joel should focus on threat hunting in the dark web.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

☑ **A**. The best set of security advisories includes those from all IT product vendors, as well as a number of nonvendor sources such as US-CERT and InfraGard.

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because Joel should also have at least one good nonvendor source such as US-CERT. **C** is incorrect because it is important to continue to receive vendor advisories. **D** is incorrect because "threat hunting on the dark web" is not a real activity.

72. The primary advantage of automatic controls versus manual controls includes all of the following *except* which one?

A. Automatic controls are generally more reliable than manual controls.

B. Automatic controls are less expensive than manual controls.

C. Automatic controls are generally more consistent than manual controls.

D. Automatic controls generally perform better in audits than manual controls.

☑ **B**. Automatic controls are not necessarily less expensive than manual controls; in some cases, they may be considerably more expensive than manual controls.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because automated controls are typically more reliable and accurate than manual controls. **C** is incorrect because automated controls are typically more consistent than manual controls. **D** is incorrect because automated controls generally perform better in audits.

73. Which of the following statements about PCI-DSS compliance is true?

A. Only organizations that store, transfer, or process more than 6 million credit card numbers are required to undergo an annual PCI audit.

B. Service providers are not required to submit an attestation of compliance (AOC) annually.

C. Merchants that process fewer than 15,000 credit card transactions are not required to submit an attestation of compliance (AOC).

D. All organizations that store, transfer, or process credit card data are required to submit an attestation of compliance (AOC) annually.

☑ **D**. All organizations that store, process, or transmit credit card data are required to submit an attestation of compliance (AOC) annually to their acquiring bank, processing bank, or card brand.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because some organizations that process fewer credit card numbers are also required to undergo annual PCI audits—for example, organizations that have suffered a breach may be required to undergo audits. **B** is incorrect because service providers are required to submit attestations of compliance (AOC) annually. **C** is incorrect because all merchants are required to submit attestations of compliance (AOC).

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

74. A security leader wants to commission an outside company to assess the organization's performance against the NIST SP800-53 control framework to see which controls the organization is operating properly and which controls require improvement. What kind of an assessment does the security leader need to commission?

A. Controls risk assessment

B. Controls maturity assessment

C. Controls gap assessment

D. Risk assessment

☑ **C**. The organization needs to commission a controls gap assessment, which will reveal which controls are being operated properly and which ones require improvement of some kind.

☒ **A**, **B**, and **D** are incorrect. **A** is incorrect because a risk assessment will not provide the desired results. **B** is incorrect because a maturity assessment will not provide the desired results. **D** is incorrect because a risk assessment will not provide the desired results.

75. An organization needs to better understand how well organized its operations are from a controls point of view. What kind of an assessment will best reveal this?

A. Controls risk assessment

B. Controls maturity assessment

C. Controls gap assessment

D. Risk assessment

☑ **B**. A controls maturity assessment will reveal, control by control, the level of organization and consistency of each control in the organization.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because a controls risk assessment will not provide the desired results. **C** is incorrect because a controls gap assessment will not provide the desired results. **D** is incorrect because a risk assessment will not provide the desired results.

76. An organization needs to better understand which of its controls are more important than others. What kind of an assessment will best reveal this?

A. Controls risk assessment

B. Controls maturity assessment

C. Controls gap assessment

D. Risk assessment

☑ **A**. A controls risk assessment will reveal which controls have greater risk associated with them. This will help the organization better understand which controls warrant greater attention and scrutiny.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because a controls maturity assessment will not provide the desired results. **C** is incorrect because a controls gap assessment will not provide the desired results. **D** is incorrect because a risk assessment will not provide the desired results.

77. An organization needs to better understand whether its control framework is adequately protecting the organization from known and unknown hazards. What kind of an assessment will best reveal this?

A. Controls risk assessment

B. Controls maturity assessment

C. Controls gap assessment

D. Risk assessment

☑ **D**. A risk assessment will best help the organization understand the entire array of risks and potential impacts facing the organization and whether its control framework is adequately covering them.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because a controls risk assessment (the next best choice) will not provide the desired results. **B** is incorrect because a controls maturity assessment will not provide the desired results. **C** is incorrect because a controls gap assessment will not provide the desired results.

78. An organization recently suffered a significant security incident. The organization was surprised by the incident and believed that this kind of an event would not occur. To avoid a similar event in the future, what should the organization do next?

A. Commission an enterprise-wide risk assessment.

B. Commission a controls maturity assessment.

C. Commission an internal and external penetration test.

D. Commission a controls gap assessment.

☑ **A**. An enterprise-wide risk assessment is the best option here so that risks of all kinds can be identified and remedies suggested for mitigating them.

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because it's possible that there are missing controls; a controls maturity assessment takes too narrow a view here and focuses only on existing controls, when the problem might be controls that are nonexistent. **C** is incorrect because the nature of the incident is unknown and may not be related to technical vulnerabilities that a penetration test would reveal (for example, it may have been phishing or fraud). **D** is incorrect because a controls gap assessment takes too narrow a view here and focuses only on existing controls, when the problem might be controls that are nonexistent.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

79. Stephen is a security leader for a SaaS company that provides file storage services to corporate clients. Stephen is examining proposed contract language from a prospective customer that is requiring the SaaS company implement "best practices" for protecting customer information. How should Stephen respond to this contract language?

A. Stephen should accept the contract language as-is.

B. Stephen should not accept a customer's contract but instead use his company's contract language.

C. Stephen should change the language from "best practices" to "industry-standard practices."

D. Stephen should remove the security-related language as it is unnecessary for a SaaS environment.

☑ **C.** The term "best practices" is good to impose on others but bad to accept from others. "Best practices" in this case implies that Stephen's company will use the best available processes and tools that are superior to all others. Instead, a phrase such as "industry-standard practices" should be used.

☒ **A**, **B**, and **D** are incorrect. **A** is incorrect because few companies can afford to truly implement "best practices" controls, particularly a SaaS company that stores information. **B** is incorrect because it is commonplace to accept a customer's contract (just as it is commonplace to use one's own). **D** is incorrect because complete removal of the security language will likely be unacceptable by the customer.

80. Security analysts in the SOC have noticed that the organization's firewall is being scanned by a port scanner in a hostile country. Security analysts have notified the security manager. How should the security manager respond to this matter?

A. Declare a high-severity security event.

B. Declare a low-severity security event.

C. Take no action.

D. Direct the SOC to blackhole the scan's originating IP address.

☑ **D.** The best course of action is to blackhole the IP address that is the origination of the port scan. However, even this may not be necessary because a port scan is not, by itself, a serious matter. However, it may represent reconnaissance by an intruder that is targeting the organization.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because a port scan is not a high-severity security matter. **B** is incorrect because this is not the best answer; however, some organizations might consider a port scan a low-level security incident and respond in some way, such as blackholing the IP address. **C** is incorrect because taking no action at all is not the best course of action.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

81. A security leader recently commissioned a controls maturity assessment and has received the final report. Control maturity in the assessment is classified as "Initial," "Managed," "Defined," "Quantitatively Managed," and "Optimized." What maturity scale was used in this maturity assessment?

A. Organizational Project Maturity Model

B. Open Source Maturity Model

C. Capability Maturity Model

D. Capability Maturity Model Integrated

☑ **D**. The maturity model used for this assessment was the Capability Maturity Model Integrated.

☒ **A**, **B**, and **C** are incorrect. The maturity levels in the question do not correspond to any of these other maturity models.

82. Security analysts in the SOC have noticed a large volume of phishing e-mails that are originating from a single "from" address. Security analysts have notified the security manager. How should the security manager respond to the matter?

A. Declare a high-level security incident.

B. Block all incoming e-mail from that address at the e-mail server or spam filter.

C. Issue an advisory to all employees to be on the lookout for suspicious messages and to disregard them.

D. Blackhole the originating IP address.

☑ **B**. Of the choices available, the best one is to block any new incoming e-mail messages from the offending e-mail address. A better solution would be the use of a system that would do this automatically, as well as retrieve any offending messages already delivered to some users before the message was recognized as harmful.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because this is not the best choice. However, depending on the nature of the threat (which is not revealed in this question), if the phishing is known to carry a malicious payload known to infect user machines successfully in the organization, then perhaps a high-severity incident is the right course of action. **C** is incorrect because this is not the best choice. However, in the absence of antiphishing controls, this may be the organization's best choice. **D** is incorrect because this is not the best choice; the adversary may be able to continue sending e-mails from different servers.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

83. The corporate controller in an organization recently received an e-mail from the CEO with instructions to wire a large amount of money to an offshore bank account that is part of secret merger negotiations. How should the corporate controller respond?

A. Contact the CEO and ask for confirmation.

B. Wire the money as directed.

C. Reply to the e-mail and ask for confirmation.

D. Direct the wire transfer clerk to wire the money as directed.

☑ **A.** The best course of action is to contact the CEO directly, via phone or e-mail, asking for confirmation of the directive. On the surface, this appears to be a case of business e-mail compromise (BEC).

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because this may be a case of business e-mail compromise (BEC) that could result in large financial losses. **C** is incorrect because this may be a case of business e-mail compromise. A better response would be to initiate a new e-mail to the CEO; better yet would be a phone call. **D** is incorrect because this appears to be a case of business e-mail compromise (BEC) that could result in large financial losses.

84. An organization's information security department conducts quarterly user access reviews of the financial accounting system. Who is the best person to approve users' continued access to roles in the system?

A. Security manager

B. IT manager

C. Corporate controller

D. Users' respective managers

☑ **C.** The best person to approve ongoing user access in an application is a business unit leader or department head, or someone in the business responsible for the business process(es) supported by the information system.

☒ **A**, **B**, and **D** are incorrect. **A** is incorrect because the security manager is not going to be as familiar with finance department operations to know which persons should continue to have access to roles. **B** is incorrect because the IT manager is not going to be as familiar with finance department operations to know which persons should continue to have access to roles. **D** is incorrect because users' managers are not going to be as familiar with finance department operations to know which persons should continue to have access to roles.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

85. All of the following are possible techniques for setting the value of information in a database *except* which one?

A. Recovery cost

B. Replacement cost

C. Lost revenue

D. Book value

☑ **D**. Book value is the least likely method to be used to assign value to information in a database. Book value is generally used for hardware assets only.

☒ **A**, **B**, and **C** are incorrect. Recovery cost, replacement cost, and lost revenue are all feasible methods for assigning value to information in a database.

86. For disaster recovery scenarios, which of the following methods for setting the value of computer equipment is most appropriate?

A. Recovery cost

B. Replacement cost

C. Lost revenue

D. Book value

☑ **B**. Replacement cost may be best suited for disaster recovery scenarios. In a disaster situation, computer equipment may need to be replaced rather than repaired.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because recovery cost is not usually associated with computer equipment, but instead with information. **C** is incorrect because this is not the best method. If in cases where revenue derived from computer equipment is greater than its replacement value, this would underscore the need for rapid replacement or use of an alternative processing center. **D** is incorrect because it may be difficult to replace lost assets if only book value is available to obtain replacements.

87. A security leader in a SaaS services organization has recently commissioned a controls maturity assessment. The consultants who performed the assessment used the CMMI model for rating individual control maturity. The assessment report rated most controls from 2.5 to 3.5 on a scale of 1 to 5. How should the security leader interpret these results?

A. Acceptable: the maturity scores are acceptable and align with those of other software companies.

B. Unacceptable: develop a strategy to improve control maturity to 4.5–5.0 over the next three to four years.

C. Unacceptable: develop a strategy to improve control maturity to 3.4–4.5 over the next three to four years.

D. Irrelevant: too little is known to make a determination of long-term maturity targets.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

☑ **A**. These results are acceptable, and they may even be interpreted as pretty good. The maturity of security controls in a SaaS or software company is generally in the 2.5–3.5 range.

☒ **B**, **C**, and **D** are incorrect. **B** is incorrect because few organizations aspire to bring their control maturity to the 4.5–5.0 range. **C** is incorrect because few software companies aspire to bring their control maturity to the 3.5–4.5 range. **D** is incorrect because this is not the best answer. That said, the question did not specify the industry or type of software in use.

88. In a mature third-party risk management (TPRM) program, how often are third parties typically assessed?

A. At the time of onboarding and annually thereafter

B. At the time of onboarding

C. At the time of onboarding and annually thereafter if the third party is rated as high risk

D. At the time of onboarding and later on if the third party has a security incident

☑ **C**. Better organizations' TPRM programs assess all third parties at the time of onboarding. High-risk third parties are assessed annually thereafter; medium-risk third parties might be assessed every two to three years, and low-risk third parties might not be reassessed at all.

☒ **A**, **B**, and **D** are incorrect. **A** is incorrect because not all third parties warrant reassessment. **B** is incorrect because assessing third parties only at the time of onboarding is considered insufficient, particularly for medium- and high-risk third parties. **D** is incorrect because high-risk third parties should be assessed annually.

89. David, a security analyst in a financial services firm, has requested the Expense Management Company, a service provider, to furnish him with a SOC1 audit report. The Expense Management Company furnished David with a SOC1 audit report for the hosting center where Expense Management Company servers are located. How should David respond?

A. File the report and consider the Expense Management Company as assessed.

B. Analyze the report for significant findings.

C. Thank them for the report.

D. Thank them for the report and request a SOC1 audit report for the Expense Management Company itself.

☑ **D**. The SOC1 report that the Expense Management Company provided is not for its business, but instead for its hosting provider. Most of the time this is insufficient, as a SOC1 report is needed also for the company itself.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because little is still known about the Expense Management Company controls. **B** is incorrect because little is still known about the Expense Management Company controls. **C** is incorrect because little is still known about the Expense Management Company controls.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

90. A healthcare delivery organization has a complete inventory of third-party service providers and keeps good records on initial and follow-up assessments. What information should be reported to management?

A. Metrics related to the number of third-party assessments that are performed

B. A risk dashboard that indicates patterns and trends of risks associated with third parties

C. Metrics related to the number of third-party assessments, along with their results

D. Status on whether there are sufficient resources to perform third-party risk assessments

☑ **B**. The best thing to report to management is a risk dashboard that shows them which third parties have the highest risks or greatest potential impact to the organization, as well as the trends of risk over time.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because this does not portray risk. **C** is incorrect because this does not portray risk as well as a risk dashboard. **D** is incorrect because this does not directly portray risk. This is, however, an important item to report on so that management knows whether there are sufficient resources to manage third-party risk effectively.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill

Mc Graw Hill

CCSP®
Certified Cloud Security Professional

69. An estimate on the number of times that a threat might occur in a given year is known as what?

A. Annualized loss expectancy (ALE)

B. Annualized rate of occurrence (ARO)

C. Exposure factor (EF)

D. Annualized exposure factor (AEF)

☑ **B**. Annualized rate of occurrence (ARO) is defined as an estimate of the number of times that a threat will occur per year.

☒ **A**, **C**, and **D** are incorrect. **A** is incorrect because annualized loss expectancy (ALE) is defined as the annualized rate of occurrence (ARO) times the single loss expectancy (SLE). **C** is incorrect as exposure factor (EF) is the loss that represents a percentage of an asset's value (because in some cases, an asset is not completely destroyed). **D** is incorrect because there is no such term is annualized exposure factor (AEF).

70. Which is the best method for prioritizing risks and risk treatment?

A. Threat event probability times asset value, from highest to lowest

B. Threat event probability, followed by asset value

C. Professional judgment

D. A combination of threat event probability, asset value, and professional judgment

☑ **D**. The best method for prioritizing risks and risk treatment is to examine the probability of event occurrence (difficult though that may be), asset value, and impact to the organization. Professional judgment plays a big role as well because factors such as business reputation are difficult to quantify.

☒ **A**, **B**, and **C** are incorrect. **A** is incorrect because this approach allows no room for professional judgment. **B** is incorrect because there is no logical sequence based on these two items that are measured differently. **C** is incorrect because professional judgment alone risks the failure to consider high-value assets, high impact, and high probability of occurrence.

71. Joel is a security manager in a large manufacturing company. The company uses primarily Microsoft, Cisco, and Oracle products. Joel subscribes to security bulletins from these three vendors. Which of the following statements best describes the adequacy of these advisory sources?

A. Joel should also subscribe to nonvendor security sources such as US-CERT and InfraGard.

B. Joel's security advisory sources are adequate.

C. Joel should discontinue vendor sources and subscribe to nonvendor security sources such as US-CERT and InfraGard.

D. Joel should focus on threat hunting in the dark web.

LEARN MORE

BUY NOW

Because learning changes everything.®

©2020 McGraw-Hill