Answer Sheet

Q1) You are part of a security team investigating a compromised service account key. You need to audit which new resources were created by the service account. What should you do?

● Query Data Access logs.

✓ Query Admin Activity logs.

Explanation:-Admin Activity audit logs contain log entries for API calls or other administrative actions that modify the configuration or metadata of resources. For example, these logs record when users create VM instances or change Identity and Access Management permissions.
This is exactly what you want to see. "What resources were created by the SA?"
Refer - https://cloud.google.com/logging/docs/audit#admin-activity

● Query Access Transparency logs.

● Query Stackdriver Monitoring Workspace.

Q2) You have an application where the frontend is deployed on a managed instance group in subnet A and the data layer is stored on a mysql Compute Engine virtual machine (VM) in subnet B on the same VPC. Subnet A and Subnet B hold several other Compute Engine VMs. You only want to allow these application frontend to access the data in the application's mysql instance on port 3306. What should you do?

● Configure an ingress firewall rule that allows communication from the src IP range of subnet A to the tag "data-tag" that is applied to the mysql Compute Engine VM on port 3306.

✓ Configure an ingress firewall rule that allows communication from the frontend's unique service account to the unique service account of the mysql Compute Engine VM on port 3306.

Explanation:-Can definitely be an option and will meet the requirements. Service accounts are special Google accounts that belong to your application or service running on a VM and can be used to authenticate the application or service for the resources it needs to access. This is the best option out of B and D.
Refer - https://cloud.google.com/blog/products/gcp/three-ways-to-configure-robust-firewall-rules and check how to create robust firewall rules.

● Configure a network tag "fe-tag" to be applied to all instances in subnet A and a network tag "data-tag" to be applied to all instances in subnet B. Then configure an egress firewall rule that allows communication from Compute Engine VMs tagged with data-tag to destination Compute Engine VMs tagged fe-tag.

● Configure a network tag "fe-tag" to be applied to all instances in subnet A and a network tag "data-tag" to be applied to all instances in subnet B. Then configure an ingress firewall rule that allows communication from Compute Engine VMs tagged with fe-tag to destination Compute Engine VMs tagged with data-tag.

Q3) You are the security admin of your company. You have 3,000 objects in your Cloud Storage bucket. You do not want to manage access to each object individually. You also do not want the uploader of an object to always have full control of the object. However, you want to use Cloud Audit Logs to manage access to your bucket. What should you do?

● Set up an ACL with OWNER permission to a scope of allUsers.

● Set up an ACL with READER permission to a scope of allUsers.

● Set up a default bucket ACL and manage access for users using IAM.

✓ Set up Uniform bucket-level access on the Cloud Storage bucket and manage access for users using IAM.

Explanation:-Uniform (recommended): Uniform bucket-level access allows you to use Identity and Access Management (IAM) alone to manage permissions. IAM applies permissions to all the objects contained inside the bucket or groups of objects with common name prefixes. IAM also allows you to use features that are not available when working with ACLs, such as IAM Conditions and Cloud Audit Logs.
Refer - https://cloud.google.com/storage/docs/access-control

Q4) A customer needs an alternative to storing their plain text secrets in their source-code management (SCM) system. How should the customer achieve this using Google Cloud Platform?

● Use Cloud Source Repositories, and store secrets in Cloud SQL.

✓ Encrypt the secrets with a Customer-Managed Encryption Key (CMEK), and store them in Cloud Storage.

Explanation:-Currently the best option to achieve this is by using Secret Manager but out of all the options listed here the only option that makes sense is B. Let me know in the QA section if you don't agree with me.

● Run the Cloud Data Loss Prevention API to scan the secrets, and store them in Cloud SQL.

● Deploy the SCM to a Compute Engine VM with local SSDs, and enable preemptible VMs.

Q5) You need to provide a corporate user account in Google Cloud for each of your developers and operational staff who need direct access to GCP resources. Corporate policy requires you to maintain the user identity in a third-party identity management provider and leverage single sign-on. You learn that a significant number of users are using their corporate domain email addresses for personal Google accounts, and you need to follow Google recommended practices to convert existing unmanaged users to managed accounts. Which two actions should you take? (Choose two.)

✓ Use Google Cloud Directory Sync to synchronize your local identity management system to Cloud Identity.

Explanation:-Requirement is third-party identity management provider and leverage single sign-on.
Refer - https://cloud.google.com/architecture/identity/assessing-existing-user-accounts
(Use the transfer tool for unmanaged users to identify consumer accounts that use an email address that matches one of the domains you've added to Cloud Identity or G Suite.)
For more details refer - https://cloud.google.com/architecture/identity/migrating-consumer-accounts#initiating_a_transfer

● Use the Google Admin console to view which managed users are using a personal account for their recovery email.

● Add users to your managed Google account and force users to change the email addresses associated with their personal accounts.

✓ Use the Transfer Tool for Unmanaged Users (TTUU) to find users with conflicting accounts and ask them to transfer their personal Google accounts.

Explanation:-Requirement is third-party identity management provider and leverage single sign-on.
Refer - https://cloud.google.com/architecture/identity/assessing-existing-user-accounts
(Use the transfer tool for unmanaged users to identify consumer accounts that use an email address that matches one of the domains you've added to Cloud Identity or G Suite.)

For more details refer - https://cloud.google.com/architecture/identity/migrating-consumer-accounts#initiating_a_transfer

**Q6) A customer needs an alternative to storing their plain text secrets in their source-code management (SCM) system.
How should the customer achieve this using Google Cloud Platform?**

● Use Cloud Source Repositories, and store secrets in Cloud SQL.
✓ Encrypt the secrets with a Customer-Managed Encryption Key (CMEK), and store them in Cloud Storage.
● Run the Cloud Data Loss Prevention API to scan the secrets, and store them in Cloud SQL.
● Deploy the SCM to a Compute Engine VM with local SSDs, and enable preemptible VMs.

**Q7) Your team wants to centrally manage GCP IAM permissions from their on-premises Active Directory Service. Your team wants to manage permissions by AD group membership.
What should your team do to meet these requirements?**

● Set up Cloud Directory Sync to sync groups, and set IAM permissions on the groups.
✓ Set up SAML 2.0 Single Sign-On (SSO), and assign IAM permissions to the groups.
Explanation:-Reference: https://cloud.google.com/blog/products/identity-security/using-your-existing-identity-management-system-with-google-cloud-platform
● Use the Cloud Identity and Access Management API to create groups and IAM permissions from Active Directory.
● Use the Admin SDK to create groups and assign IAM permissions from Active Directory.

**Q8) Your team needs to obtain a unified log view of all development cloud projects in your SIEM. The development projects are under the NONPROD organization folder with the test and pre-production projects. The development projects share the ABC-BILLING billing account with the rest of the organization.
Which logging export strategy should you use to meet the requirements?**

● 1. Export logs to a Cloud Pub/Sub topic with folders/NONPROD parent and includeChildren property set to True in a dedicated SIEM project.
2. Subscribe SIEM to the topic.
✓ 1. Create a Cloud Storage sink with billingAccounts/ABC-BILLING parent and includeChildren property set to False in a dedicated SIEM project.
2. Process Cloud Storage objects in SIEM.
● 1. Export logs in each dev project to a Cloud Pub/Sub topic in a dedicated SIEM project.
2. Subscribe SIEM to the topic.
● 1. Create a Cloud Storage sink with a publicly shared Cloud Storage bucket in each project.
2. Process Cloud Storage objects in SIEM.

**Q9) In order to meet PCI DSS requirements, a customer wants to ensure that all outbound traffic is authorized.
Which two cloud offerings meet this requirement without additional compensating controls? (Choose two.)**

✓ App Engine
Explanation:-Reference: https://cloud.google.com/solutions/pci-dss-compliance-in-gcp
● Cloud Functions
✓ Compute Engine
Explanation:-Reference: https://cloud.google.com/solutions/pci-dss-compliance-in-gcp
● Google Kubernetes Engine

**Q10) Your company is using Cloud Dataproc for its Spark and Hadoop jobs. You want to be able to create, rotate, and destroy symmetric encryption keys used for the persistent disks used by Cloud Dataproc. Keys can be stored in the cloud.
What should you do?**

✓ Use the Cloud Key Management Service to manage the data encryption key (DEK).
● Use the Cloud Key Management Service to manage the key encryption key (KEK).
● Use customer-supplied encryption keys to manage the data encryption key (DEK).
● Use customer-supplied encryption keys to manage the key encryption key (KEK).

Q11)

You are creating an internal App Engine application that needs to access a user's Google Drive on the user's behalf.

Your company does not want to rely on the current user's credentials. It also wants to follow Google-recommended practices.

What should you do?

● Create a new service account, and grant it G Suite domain-wide delegation. Have the application use it to impersonate the user.
● Use a dedicated G Suite Admin account, and authenticate the application's operations with these G Suite credentials.
● Create a new Service account, and add all application users to a Google Group. Give this group the role of Service Account User.
✓ Create a new Service account, and give all application users the role of Service Account User.

Q12)

A customer wants to move their sensitive workloads to a Compute Engine-based cluster using Managed Instance Groups (MIGs).

The jobs are bursty and must be completed quickly. They have a requirement to be able to manage and rotate the encryption keys.

Which boot disk encryption solution should you use on the cluster to meet this customer's requirements?

● Encryption by default
● Customer-supplied encryption keys (CSEK)

✓ Customer-managed encryption keys (CMEK) using Cloud Key Management Service (KMS)
● Pre-encrypting files before transferring to Google Cloud Platform (GCP) for analysis

Q13)

Your company is using Cloud Dataproc for its Spark and Hadoop jobs.

You want to be able to create, rotate, and destroy symmetric encryption keys used for the persistent disks used by Cloud Dataproc.

Keys can be stored in the cloud.

What should you do?

● Use customer-supplied encryption keys to manage the key encryption key (KEK).
● Use customer-supplied encryption keys to manage the data encryption key (DEK).
● Use the Cloud Key Management Service to manage the key encryption key (KEK).
✓ Use the Cloud Key Management Service to manage the data encryption key (DEK).

Q14)

You are a member of the security team at an organization.

Your team has a single GCP project with credit card payment processing systems alongside web applications and data processing systems.

You want to reduce the scope of systems subject to PCI audit standards.

What should you do?

✓ Use VPN for all connections between your office and cloud environments.
● Move the cardholder data environment into a separate GCP project.
● Use only applications certified compliant with PA-DSS.
● Use multi-factor authentication for admin access to the web application.

Q15)

A customer's internal security team must manage its own encryption keys for encrypting data on Cloud Storage and decides to use customer-supplied encryption keys (CSEK).

How should the team complete this task?

✓ Encrypt the object, then use the gsutil command line tool or the Google Cloud Platform Console to upload the object to Cloud Storage.
● Generate an encryption key in the Google Cloud Platform Console, and upload an object to Cloud Storage using the specified key.
● Use the gsutil command line tool to upload the object to Cloud Storage, and specify the location of the encryption key.
● Upload the encryption key to a Cloud Storage bucket, and then upload the object to the same bucket.

Q16)

A customer has 300 engineers.

The company wants to grant different levels of access and efficiently manage IAM permissions between users in the development and production environment projects.

Which two steps should the company take to meet these requirements? (Choose two.)

● Create projects for each environment, and grant IAM rights to each engineering user.
✓ Create an Organizational Policy constraint for each folder environment.
● Create a Google Group for the Engineering team, and assign permissions at the folder level.
✓ Create a folder for each development and production environment.
● Create a project with multiple VPC networks for each environment.

Q17)

You want to evaluate GCP for PCI compliance. You need to identify Google's inherent controls.

Which document should you review to find the information?

✓ PCI SSC Cloud Computing Guidelines
Explanation:-Evaluating compliance is to know which service / product compliess to standard, and to which extent. PCI Data Security Standard compliance guide helps you learn how to implement the Payment Card Industry Data Security Standard (PCI DSS) for your business on Google Cloud. link - https://cloud.google.com/solutions/pci-dss-compliance-in-gcp
● PCI DSS Requirements and Security Assessment Procedures
● Google Cloud Platform: Customer Responsibility Matrix
● Product documentation for Compute Engine

Q18)

Your company runs a website that will store PII on Google Cloud Platform.

To comply with data privacy regulations, this data can only be stored for a specific amount of time and must be fully deleted after this specific period.

Data that has not yet reached the time period should not be deleted. You want to automate the process of complying with this

regulation.

What should you do?

● Store the data in a single BigTable table and set an expiration time on the column families.
● Store the data in a single Cloud Storage bucket and configure the bucket's Time to Live.
✓ Store the data in a single BigQuery table and set the appropriate table expiration time.
● Store the data in a single Persistent Disk, and delete the disk at expiration time.

Q19)

A DevOps team will create a new container to run on Google Kubernetes Engine.

As the application will be internet-facing, they want to minimize the attack surface of the container.

What should they do?

● Use a Continuous Delivery tool to deploy the application.
● Delete non-used versions from Container Registry.
✓ Build small containers using small base images.
● Use Cloud Build to build the container images.

Q20)

While migrating your organization's infrastructure to GCP, a large number of users will need to access GCP Console.

The Identity Management team already has a well-established way to manage your users and want to keep using your existing Active Directory or LDAP server along with the existing SSO password.

What should you do?

● Users sign in using OpenID (OIDC) compatible IdP, receive an authentication token, then use that token to log in to the GCP Console.
● Users sign in directly to the GCP Console using the credentials from your on-premises Kerberos compliant identity provider.
✓ Use Google Cloud Directory Sync to synchronize the data in Google domain with your existing Active Directory or LDAP server.
● Manually synchronize the data in Google domain with your existing Active Directory or LDAP server.

Q21)

Your company has deployed an application on Compute Engine. The application is accessible by clients on port 587.

You need to balance the load between the different instances running the application. The connection should be secured using TLS, and terminated by the Load Balancer.

What type of Load Balancing should you use?

● TCP Proxy Load Balancing
● HTTP(S) Load Balancing
● Network Load Balancing
✓ SSL Proxy Load Balancing

Q22) What is the correct definition for Tier II of the Uptime Institute Data Center Site Infrastructure Tier Standard Topology?

● Concurrently Maintainable Site Infrastructure
● Fault-Tolerant Site Infrastructure
● Basic Site Infrastructure
✓ Redundant Site Infrastructure Capacity Components
Explanation:-The Uptime Institute's Data Center Site Infrastructure Tier Classification -
Tier I – Basic Data Center Site Infrastructure
Tier II – Redundant Site Infrastructure Capacity Components
Tier III – Concurrently Maintainable Site Infrastructure
Tier IV – Fault Tolerant Site Infrastructure
Link - https://www.just.edu.jo/~natheer/INCS775/TierStandards.pdf

Q23) Which of the following is the recommended operating range for temperature and humidity in a data center?

✓ Between 64° F and 81° F and 40 percent and 60 percent relative humidity
Explanation:-The recommended operating range for temperature and humidity in a data center is: 18-27°C / 64-80°F. Link - https://serverscheck.com/sensors/temperature_best_practices.asp
● Between 64° F and 84° F and 30 percent and 60 percent relative humidity
● Between 60° F and 85° F and 40 percent and 60 percent relative humidity
● Between 62° F and 81° F and 40 percent and 65 percent relative humidity

Q24) Which of the following are supported authentication methods for iSCSI? (Choose two.)

✓ SRP
Explanation:-A number of authentication methods are supported with iSCSI:
Kerberos: A network authentication protocol designed to provide strong authentication for client/server applications by using secret key cryptography. The Kerberos protocol uses strong cryptography so that a client can prove its identity to server (and vice versa) across an insecure network connection. After a client and server use Kerberos to prove their identity, they can encrypt all their communications to ensure privacy and data

integrity as they go about their business.
SRP: A secure password-based authentication and key-exchange protocol that exchanges a cryptographically strong secret as a by-product of successful authentication. This enables the two parties to communicate securely.
SPKM1/2: Provides authentication, key establishment, data integrity, and data confidentiality in an online distributed application environment using a publickey infrastructure. The use of a public-key infrastructure allows digital signatures supporting nonrepudiation to be employed for message exchanges.
CHAP: Used to periodically verify the identity of the peer using a three-way handshake. This is done upon initial link establishment and may be repeated anytime after the link has been established.

● TLS
✓ Kerberos

Explanation:-A number of authentication methods are supported with iSCSI:
Kerberos: A network authentication protocol designed to provide strong authentication for client/server applications by using secret key cryptography. The Kerberos protocol uses strong cryptography so that a client can prove its identity to server (and vice versa) across an insecure network connection. After a client and server use Kerberos to prove their identity, they can encrypt all their communications to ensure privacy and data integrity as they go about their business.
SRP: A secure password-based authentication and key-exchange protocol that exchanges a cryptographically strong secret as a by-product of successful authentication. This enables the two parties to communicate securely.
SPKM1/2: Provides authentication, key establishment, data integrity, and data confidentiality in an online distributed application environment using a publickey infrastructure. The use of a public-key infrastructure allows digital signatures supporting nonrepudiation to be employed for message exchanges.
CHAP: Used to periodically verify the identity of the peer using a three-way handshake. This is done upon initial link establishment and may be repeated anytime after the link has been established.

● L2TP

Q25)

A large financial institution is moving its Big Data analytics to Google Cloud Platform.

They want to have maximum control over the encryption process of data stored at rest in BigQuery.

What technique should the institution use?

✓ Customer-managed encryption keys (CMEK).
● Use a Cloud Hardware Security Module (Cloud HSM).
● Use Cloud Storage as a federated Data Source.
● Customer-supplied encryption keys (CSEK).

Q26)

A company is deploying their application on Google Cloud Platform.

Company policy requires long-term data to be stored using a solution that can automatically replicate data over at least two geographic places.

Which Storage solution are they allowed to use?

● Compute Engine Persistent Disk
● Compute Engine SSD Disk
● Cloud BigQuery
✓ Cloud Bigtable

Q27)

A large e-retailer is moving to Google Cloud Platform with its ecommerce website.

The company wants to ensure payment information is encrypted between the customer's browser and GCP when the customers checkout online.

What should they do?

● Configure the firewall to allow outbound traffic on port 443, and block all other outbound traffic.
● Configure the firewall to allow inbound traffic on port 443, and block all other inbound traffic.
● Configure an SSL Certificate on a Network TCP Load Balancer and require encryption.
✓ Configure an SSL Certificate on an L7 Load Balancer and require encryption.

Q28) When does an XSS flaw occur?

● Whenever an application takes untrusted data and sends it to a web browser with proper validation or escaping
● Whenever an application takes trusted data and sends it to a web browser with proper validation or escaping
✓ Whenever an application takes untrusted data and sends it to a web browser without proper validation or escaping

Explanation:-XSS flaws occur whenever an application takes untrusted data and sends it to a web browser without proper validation or escaping. XSS allows attackers to execute scripts in the victim’s browser, which can hijack user sessions, deface websites, or redirect the user to malicious sites.

● Whenever an application takes trusted data and sends it to a web browser without proper validation or escaping

---

Q29)

Applications often require access to "secrets" - small pieces of sensitive data at build or run time.

The administrator managing these secrets on GCP wants to keep a track of "who did what, where, and when?" within their GCP projects.

Which two log streams would provide the information that the administrator is looking for? (Choose two.)

● VPC Flow logs
✓ Data Access logs
● System Event logs
✓ Admin Activity logs
● Agent logs

Q30)

You are in charge of migrating a legacy application from your company datacenters to GCP before the current maintenance contract expires.

You do not know what ports the application is using and no documentation is available for you to check.

You want to complete the migration without putting your environment at risk. What should you do?

● Refactor the application into a micro-services architecture in a GKE cluster. Disable all traffic from outside the cluster using Firewall Rules. Use VPC Flow logs to determine what traffic should be allowed for the application to work properly.
● Refactor the application into a micro-services architecture hosted in Cloud Functions in an isolated project. Disable all traffic from outside your project using Firewall Rules. Use VPC Flow logs to determine what traffic should be allowed for the application to work properly.
● Migrate the application into an isolated project using a "Lift & Shift" approach in a custom network. Disable all traffic within the VPC and look at the Firewall logs to determine what traffic should be allowed for the application to work properly.
✓ Migrate the application into an isolated project using a "Lift & Shift" approach. Enable all internal TCP traffic using VPC Firewall rules. Use VPC Flow logs to determine what traffic should be allowed for the application to work properly.
Explanation:-
The goal of a migration like this is to provide a more nimble and scalable environment for individual features of the site, where the features can be more easily managed and updated than if they are part of the monolithic platform. Running in such an environment leads to faster improvements on each migrated feature, providing users with value along the way. Link - https://cloud.google.com/solutions/migrating-a-monolithic-app-to-microservices-gke

Q31) You are a member of the security team at an organization. Your team has a single GCP project with credit card payment processing systems alongside web applications and data processing systems. You want to reduce the scope of systems subject to PCI audit standards.
What should you do?

● Use multi-factor authentication for admin access to the web application.
● Use only applications certified compliant with PA-DSS.
● Move the cardholder data environment into a separate GCP project
✓ Use VPN for all connections between your office and cloud environments.
Explanation:-Reference: https://cloud.google.com/solutions/pci-dss-compliance-in-gcp

Q32) A retail customer allows users to upload comments and product reviews. The customer needs to make sure the text does not include sensitive data before the comments or reviews are published.
Which Google Cloud Service should be used to achieve this?

● Cloud Key Management Service
● Cloud Data Loss Prevention API
● BigQuery
✓ Cloud Security Scanner

Q33) A website design company recently migrated all customer sites to App Engine. Some sites are still in progress and should only be visible to customers and company employees from any location. Which solution will restrict access to the in-progress sites?

● Upload an .htaccess file containing the customer and employee user accounts to App Engine.
● Create an App Engine firewall rule that allows access from the customer and employee networks and denies all other traffic.
✓ Enable Cloud Identity-Aware Proxy (IAP), and allow access to a Google Group that contains the customer and employee user accounts.
Explanation:-Refer - https://cloud.google.com/iap/docs/concepts-overview#when_to_use_iap
● Use Cloud VPN to create a VPN connection between the relevant on-premises networks and the company's GCP Virtual Private Cloud (VPC) network.

Q34) A company's application is deployed with a user-managed Service Account key. You want to use Google-recommended practices to rotate the key. What should you do?

● Open Cloud Shell and run gcloud iam service-accounts enable-auto-rotate --iam-account=IAM_ACCOUNT.
● Open Cloud Shell and run gcloud iam service-accounts keys rotate --iam-account=IAM_ACCOUNT --key=NEW_KEY.
✓ Create a new key, and use the new key in the application. Delete the old key from the Service Account.
Explanation:-Refer - https://cloud.google.com/iam/docs/understanding-service-accounts#managing_service_account_keys
● Create a new key, and use the new key in the application. Store the old key on the system as a backup key.

Q35) In a federated environment, who is the relying party, and what does it do?

● The relying party is the service provider; it consumes the tokens that the customer generates.
✓ The relying party is the service provider; it consumes the tokens that the identity provider generates.
Explanation:-In a federated environment, there is an identity provider and a relying party. The identity provider holds all the identities and generates a token for known users. The relying party is the service provider and consumes these tokens.

● The relying party is the identity provider; it consumes the tokens that the service provider generates.
● The relying party is the customer; he consumes the tokens that the identity provider generates.

Q36) Select the page of Configuration Manager in GCDS, to select the type of object you want to synchronize

✓ General Settings page
Explanation:-On the General Settings page, specify what you intend to synchronize from your LDAP server. Link- https://support.google.com/a/answer/6162412?hl=en#zippy=%2Cspecify-your-general-settings%2Cdefine-your-ldap-settings%2Cdefine-your-google-domain-settings
● Configuration page
● Org Units page
● LDAP Configuration page

Q37) Identify the page of Configuration Manager in GCDS, to specify how your LDAP organizational units correspond to organizational units in your Google domain ?

● General Settings page
● Configuration page
✓ Org Units page
● LDAP Configuration page

Q38) Select the correct levels to incorporate logical design for data separation

✓ Compute nodes and network
Explanation:-Logical design for data separation needs to be incorporated at all the following levels:
Management plane
Compute nodes
Storage nodes
Control plane
Network
Link - https://resources.infosecinstitute.com/certification/ccsp-domain-5-operations/
● Storage nodes and application
● Control plane and session
● Management plane and presentation

Q39) What are the two biggest challenges associated with the use of IPSec in cloud computing environments?

● Auditability and governance
✓ Configuration management and performance
Explanation:-The two key challenges with the deployment and use of IPSec follow:
Configuration management: The use of IPSec is optional, and as such, many endpoint devices connecting to cloud infrastructure do not have IPSec support enabled and configured. If IPSec is not enabled on the endpoint, then depending on the configuration choices made on the server side of the IPSec solution, the endpoint may not be able to connect and complete a transaction if it does not support IPSec. CSPs may not have the proper visibility on the customer endpoints or the server infrastructure to understand IPSec configurations. As a result, the ability to ensure the use of IPSec to secure network traffic may be limited.
Performance: The use of IPSec imposes a performance penalty on the systems deploying the technology. Although the impact to the performance of an average system is small, it is the cumulative effect of IPSec across an enterprise architecture, end to end, that must be evaluated prior to implementation.
● Access control and patch management
● Training customers on how to use IPSec and documentation

Q40) What does the concept of nondestructive testing mean in the context of a vulnerability assessment?

● Known vulnerabilities are not exploited during the vulnerability assessment.
● Detected vulnerabilities are not exploited after the vulnerability assessment.
✓ Detected vulnerabilities are not exploited during the vulnerability assessment.
Explanation:-During a vulnerability assessment, the cloud environment is tested for known vulnerabilities. Detected vulnerabilities are not exploited during a vulnerability assessment (nondestructive testing) and may require further validation to detect false positives.
● Known vulnerabilities are not exploited before the vulnerability assessment.

Q41) What should configuration management always be tied to?

● IT service management
✓ Change management
Explanation:-The need to tie configuration management to change management is because change management has to approve any changes to all production systems prior to them taking place. In other words, there should never be a change that is allowed to take place to a Configuration Item (CI) in a production system unless change management has approved the change first
● Financial management
● Business relationship management

Q42) What are the objectives of change management? (Choose all that apply.)

✓ Ensure that changes are recorded and evaluated.
Explanation:-The objective of change management in this context is to ensure that standardized methods and procedures are used for efficient and prompt handling of all changes to control IT infrastructure, in order to minimize the number and impact of any related incidents upon service. -

Wikipedia

Change management has several objectives:

Respond to a customer's changing business requirements while maximizing value and reducing incidents, disruption, and rework.

Respond to business and IT requests for change that aligns services with business needs.

Ensure that changes are recorded and evaluated.

Ensure that authorized changes are prioritized, planned, tested, implemented, documented, and reviewed in a controlled manner.

Ensure that all changes to CIs are recorded in the configuration management system.

Optimize overall business risk; it is often correct to minimize business risk, but sometimes it is appropriate to knowingly accept a risk because of the potential benefit.

✓ Respond to a customer's changing business requirements while maximizing value and reducing incidents, disruption, and rework.

Explanation:-The objective of change management in this context is to ensure that standardized methods and procedures are used for efficient and prompt handling of all changes to control IT infrastructure, in order to minimize the number and impact of any related incidents upon service. -

Wikipedia

Change management has several objectives:

Respond to a customer's changing business requirements while maximizing value and reducing incidents, disruption, and rework.

Respond to business and IT requests for change that aligns services with business needs.

Ensure that changes are recorded and evaluated.

Ensure that authorized changes are prioritized, planned, tested, implemented, documented, and reviewed in a controlled manner.

Ensure that all changes to CIs are recorded in the configuration management system.

Optimize overall business risk; it is often correct to minimize business risk, but sometimes it is appropriate to knowingly accept a risk because of the potential benefit.

● Respond to business and IT requests for change that will disassociate services with business needs.

● Ensure that all changes are prioritized, planned, tested, implemented, documented, and reviewed in a controlled manner.

Q43) What does an audit scope statement provide to a cloud service customer or organization?

● The outcome of the audit, as well as any findings that need to be addressed

● A list of all the security controls to be audited

✓ The required level of information for the client or organization subject to the audit to fully understand (and agree) with the scope, focus, and type of assessment being performed

Explanation:-An audit scope statement provides the required level of information for the client or organization subject to the audit to fully understand (and agree with) the scope, focus, and type of assessment being performed. Typically, an audit scope statement includes the following:

General statement of focus and objectives

Scope of audit (including exclusions)

Type of audit (certification, attestation, and so on)

Security assessment requirements

Assessment criteria (including ratings)

Acceptance criteria

Deliverables

Classification (confidential, highly confidential, secret, top secret, public, and so on)

The audit scope statement can also catalog the circulation list, along with key individuals associated with the audit.

● The credentials of the auditors, as well as the projected cost of the audit

Q44) Which of the following should be carried out first when seeking to perform a gap analysis?

● Identify potential risks.

✓ Obtain management support.

Explanation:-Numerous stages are carried out prior to commencing a gap analysis review. Although they can vary depending on the review, common stages include the following:

1. Obtain management support from the right managers.
2. Define the scope and objectives.
3. Plan an assessment schedule.
4. Agree on a plan.
5. Conduct information gathering exercises
6. Interview key personnel.
7. Review supporting documentation.
8. Verify the information obtained.
9. Identify any potential risks.
10. Document the findings.
11. Develop a report and recommendations.
12. Present the report.
13. Sign off and accept the report.

The objective of a gap analysis is to identify and report on any gaps or risks that may affect the AIC of key information assets. The value of such an assessment is often determined based on what you did not know or for an independent resource to communicate to relevant management or senior personnel such risks, as opposed to internal resources saying what you need or should be doing.

● Define scope and objectives.

● Conduct information gathering.

Q45) What is the first international set of privacy controls in the cloud?

● ISO/IEC 27032

● ISO/IEC 27005

● ISO/IEC 27002

✓ ISO/IEC 27018

Explanation:-ISO/IEC 27018 addresses the privacy aspects of cloud computing for consumers. ISO 27018 is the first international set of privacy controls in the cloud. ISO 27018 was published by the ISO on July 30, 2014, as a new component of the ISO 27001 standard. ISO 27018 sets forth a code of practice for protection of PII in public clouds acting as PII processors. CSPs adopting ISO/IEC 27018 must operate under five key principles:

Consent: CSPs must not use the personal data they receive for advertising and marketing unless expressly instructed to do so by the customers. In

addition, a customer should be able to employ the service without having to consent to the use of her personal data for advertising or marketing.
Control: Customers have explicit control over how CSPs are to use their information.
Transparency: CSPs must inform customers about items such as where their data resides. CSPs also need to disclose to customers the use of any subcontractors who will be used to process PII.
Communication: CSPs should keep clear records about any incident and their response to it, and they should notify customers.
Independent and yearly audit: To remain compliant, the CSP must subject
itself to yearly third-party reviews. This allows the customer to rely upon the findings to support her own regulatory obligations.
Trust is key for consumers leveraging the cloud; therefore, vendors of cloud services are working toward adopting the stringent privacy principles outlined in ISO 27018.

Q46) What is domain A.16 of the ISO 27001:2013 standard?

● Security Policy Management
● Organizational Asset Management
● System Security Management
✓ Security Incident Management
Explanation:-The following domains make up the ISO 27001:2013, the most widely used global standard for ISMS implementations:
A.5—Security Policy Management
A.6—Corporate Security Management
A.7—Personnel Security Management
A.8—Organizational Asset Management
A.9—Information Access Management
A.10—Cryptography Policy Management
A.11—Physical Security Management
A.12—Operational Security Management
A.13—Network Security Management
A.14—System Security Management
A.15—Supplier Relationship Management
A.16—Security Incident Management
A.17—Security Continuity Management
A.18—Security Compliance Management

Q47) What is a data custodian responsible for?

● Data content, context, and associated business rules
● Logging and alerts for all data
✓ The safe custody, transport, storage of data, and implementation of business rules
Explanation:-The following are key roles associated with data management:
Data subject: This is an individual who is the focus of personal data.
Data controller: This is a person who either alone or jointly with other persons determines the purposes for which and the manner in which any personal data is processed.
Data processor: In relation to personal data, this is any person other than an employee of the data controller who processes the data on behalf of the data controller.
Data stewards: These people are commonly responsible for data content, context, and associated business rules.
Data custodians: These people are responsible for the safe custody, transport, data storage, and implementation of business rules.
Data owners: These people hold legal rights and complete control over a single piece or set of data elements. Data owners can also define distribution and associated policies.
● Customer access and alerts for all data

Q48) What is typically not included in an SLA?

● Change management process to be used
✓ Pricing for the services to be covered by the SLA
Explanation:-Within an SLA, the following contents and topics should be covered as a minimum:
Availability (for example, 99.99 percent of services and data)
Performance (for example, expected response times versus maximum response times)
Security and privacy of the data (for example, encrypting all stored and transmitted data)
Logging and reporting (for example, audit trails of all access and the ability to report on key requirements and indicators)
Disaster recovery expectations (for example, worse-case recovery commitment, RTO, MPTD)
Location of the data (for example, ability to meet requirements or consistent with local legislation)
Data format and structure (for example, data retrievable from provider in readable and intelligent format)
Portability of the data (for example, ability to move data to a different provider or to multiple providers)
Identification and problem resolution (for example, help desk/service desk, call center, or ticketing system)
Change-management process (for example, updates or new services)
Dispute-mediation process (for example, escalation process and consequences)
Exit strategy with expectations on the provider to ensure smooth transition
● Availability of the services to be covered by the SLA
● Dispute mediation process to be used

Q49) What are the four elements that a data retention policy should define?

● Retention periods, data access methods, data security, and data retrieval procedures
● Retention periods, data formats, data security, and data destruction procedures
● Retention periods, data formats, data security, and data communication procedures
✓ Retention periods, data formats, data security, and data retrieval procedures
Explanation:-A data retention policy is an organization's established protocol for retaining information for operational or regulatory compliance needs. The objectives of a data retention policy are to keep important information for future use or reference, to organize information so it can be

searched and accessed at a later date, and to dispose of information that is no longer needed. The policy balances the legal, regulation, and business data archival requirements against data storage costs, complexity, and other data considerations. A good data retention policy should define the following:
Retention periods
Data formats
Data security
Data retrieval procedures for the enterprise

Q50) What are the three things that you must understand before you can determine the necessary controls to deploy for data protection in a cloud environment?

● Actors, policies, and procedures
✓ Function, location, and actors
Explanation:-To determine the necessary controls to be deployed, you must first understand the following:
Functions of the data
Locations of the data
Actors upon the data
Once you understand and document these three items, you can design the appropriate controls and apply them to the system to safeguard data and control access to it. These controls can be of a preventive, detective (monitoring), or corrective nature.
● Management, provisioning, and location
● Lifecycle, function, and cost

Q51) Which of the following are storage types used with an IaaS solution?

● Volume and block
● Structured and object
● Unstructured and ephemeral
✓ Volume and object
Explanation:-IaaS uses the following storage types:
Volume storage: A virtual hard drive that can be attached to a VM instance and be used to host data within a file system. Volumes attached to IaaS instances behave just like a physical drive or an array does. Examples include VMware VMFS, Amazon EBS, Rackspace RAID, and OpenStack Cinder.
Object storage: Object storage is like a file share accessed via APIs or a web interface. Examples include Amazon S3 and Rackspace cloud files.
