Vendor: HashiCorp

Exam Code:VAULT-ASSOCIATE

Exam Name: HashiCorp Certified: Vault Associate (002)

Version: Demo

QUESTION 1

When using Integrated Storage, which of the following should you do to recover from possible data loss?

A. Failover to a standby node

B. Use snapshot

C. Use audit logs

D. Use server logs

Correct Answer: B

Integrated Storage is a Raft-based storage backend that allows Vault to store its data internally without relying on an external storage system. It also enables Vault to run in high availability mode with automatic leader election and failover. However, Integrated Storage is not immune to data loss or corruption due to hardware failures, network partitions, or human errors. Therefore, it is recommended to use the snapshot feature to backup and restore the Vault data periodically or on demand. A snapshot is a point-in-time capture of the entire Vault data, including the encrypted secrets, the configuration, and the metadata. Snapshots can be taken and restored using the vault operator raft snapshot command or the sys/ storage/raft/snapshot API endpoint. Snapshots are encrypted and can only be restored with a quorum of unseal keys or recovery keys. Snapshots are also portable and can be used to migrate data between different Vault clusters or storage backends. References:
https://developer.hashicorp.com/vault/docs/concepts/integrated- storage1,
https://developer.hashicorp.com/vault/docs/commands/operator/raft/snapshot2,
https://developer.hashicorp.com/vault/api-docs/system/ storage/raft/snapshot3

QUESTION 2

A web application uses Vault\\'s transit secrets engine to encrypt data in-transit. If an attacker intercepts the data in transit which of the following statements are true? Choose two correct answers.

A. You can rotate the encryption key so that the attacker won\\'t be able to decrypt the data

B. The keys can be rotated and min_decryption_version moved forward to ensure this data cannot be decrypted

C. The Vault administrator would need to seal the Vault server immediately

D. Even if the attacker was able to access the raw data, they would only have encrypted bits (TLS in transit)

Correct Answer: BD

A web application that uses Vault\\'s transit secrets engine to encrypt data in- transit can benefit from the following security features: Even if the attacker was able to access the raw data, they would only have encrypted bits (TLS in transit). This means that the attacker would need to obtain the encryption key from Vault in order to decrypt the data, which is protected by Vault\\'s authentication and authorization mechanisms. The transit secrets engine does not store the data sent to it, so the attacker cannot access the data from Vault either. The keys can be rotated and min_decryption_version moved forward to ensure this data cannot be decrypted. This means that the web application can periodically change the encryption key used to encrypt the data, and set a minimum decryption version for the key, which prevents older versions of the key from being used to decrypt the data. This way, even if the attacker somehow obtained an old version of the key, they would not be able to decrypt the data that was encrypted with a newer version of the key. The other statements are not true, because: You cannot rotate the encryption key so that the attacker won\\'t be able to decrypt the data. Rotating the key alone does not prevent the attacker from decrypting the data, as they may

still have access to the old version of the key that was used to encrypt the data. You need to also move the min_decryption_version forward to invalidate the old version of the key. The Vault administrator would not need to seal the Vault server immediately. Sealing the Vault server would make it inaccessible to both the attacker and the legitimate users, and would require unsealing it with the unseal keys or the recovery keys. Sealing the Vault server is a last resort option in case of a severe compromise or emergency, and is not necessary in this scenario, as the attacker does not have access to the encryption key or the data in Vault. References: Transit Secrets Engines | Vault | HashiCorp Developer, Encryption as a service: transit secrets engine | Vault | HashiCorp Developer

---

QUESTION 3

Which of the following describes usage of an identity group?

A. Limit the policies that would otherwise apply to an entity in the group

B. When they want to revoke the credentials for a whole set of entities simultaneously

C. Audit token usage

D. Consistently apply the same set of policies to a collection of entities

Correct Answer: D

An identity group is a collection of entities that share some common attributes. An identity group can have one or more policies attached to it, which are inherited by all the members of the group. An identity group can also have subgroups, which can further refine the policies and attributes for a subset of entities. One of the use cases of an identity group is to consistently apply the same set of policies to a collection of entities. For example, an organization may have different teams or departments, such as engineering, sales, or marketing. Each team may have its own identity group, with policies that grant access to the secrets and resources that are relevant to their work. By creating an identity group for each team, the organization can ensure that the entities belonging to each team have the same level of access and permissions, regardless of which authentication method they use to log in to Vault. References: Identity: entities and groups | Vault | HashiCorp Developer, vault_identity_group | Resources | hashicorp/vault | Terraform | Terraform Registry

---

QUESTION 4

To give a role the ability to display or output all of the end points under the /secrets/apps/* end point it would need to have which capability set?

A. update

B. read

C. sudo

D. list

E. None of the above

Correct Answer: C

To give a role the ability to display or output all of the end points under the /secrets/apps/* end point, it would need to have the list capability set. The list capability allows a role to perform any operation on any path in Vault, including reading, writing, deleting, and listing. The list capability is required for roles that need to access sensitive data or

perform administrative tasks in Vault. The other capabilities are not relevant for this scenario, as they only allow specific operations on specific paths or secrets engines. References: Policies | Vault | HashiCorp Developer, token capabilities - Command | Vault | HashiCorp Developer

---

QUESTION 5

Your organization has an initiative to reduce and ultimately remove the use of long lived

A. 509 certificates. Which secrets engine will best support this use case?

B. PKI

C. Key/Value secrets engine version 2, with TTL defined

D. Cloud KMS

E. Transit

Correct Answer: A

The PKI secrets engine is designed to support the use case of reducing and ultimately removing the use of long lived X.509 certificates. The PKI secrets engine can generate dynamic X.509 certificates on demand, with short time-to-live (TTL) and automatic revocation. This eliminates the need for manual processes of generating, signing, and rotating certificates, and reduces the risk of certificate compromise or misuse. The PKI secrets engine can also act as a certificate authority (CA) or an intermediate CA, and can integrate with external CAs or CRLs. The PKI secrets engine can issue certificates for various purposes, such as TLS, SSH, code signing, email encryption, etc. References: https://developer.hashicorp.com/vault/docs/secrets/pki1, https://developer.hashicorp.com/vault/tutorials/getting-started/getting-started-dynamic- secrets

---

QUESTION 6

Which of the following are replication methods available in Vault Enterprise? Choose two correct answers.

A. Cluster sharding

B. Namespaces

C. Performance Replication

D. Disaster Recovery Replication

Correct Answer: CD

The replication methods available in Vault Enterprise are performance replication and disaster recovery replication. These methods allow critical data to be replicated across clusters to support horizontally scaling and disaster recovery workloads. Performance replication enables a primary cluster to replicate data to one or more secondary clusters, which can handle client requests and improve performance and availability. Performance replication replicates most Vault data, such as secrets, policies, auth methods, and leases, but not tokens. Performance secondaries generate their own tokens and leases, which are not replicated back to the primary. Performance replication also supports filtering, which allows selective replication of data based on namespaces or paths. Disaster recovery replication enables a primary cluster to replicate data to one or more secondary clusters, which act as standby clusters in case of a failure or outage of the primary. Disaster recovery replication replicates all Vault data, including tokens and leases, and maintains the same configuration and state as the primary. Disaster recovery secondaries do not handle client requests, but they can

be promoted to a primary in a disaster recovery scenario. References: Replication - Vault Enterprise | Vault | HashiCorp Developer, Performance Replication - Vault Enterprise | Vault | HashiCorp Developer, Disaster Recovery Replication - Vault Enterprise | Vault | HashiCorp Developer

---

QUESTION 7

What is a benefit of response wrapping?

A. Log every use of a secret

B. Load balanc secret generation across a Vault cluster

C. Provide error recovery to a secret so it is not corrupted in transit

D. Ensure that only a single party can ever unwrap the token and see what\\'s inside

Correct Answer: D

Response wrapping is a feature that allows Vault to take the response it would have sent to a client and instead insert it into the cubbyhole of a single-use token, returning that token instead. The client can then unwrap the token and retrieve the original response. Response wrapping has several benefits, such as providing cover, malfeasance detection, and lifetime limitation for the secret data. One of the benefits is to ensure that only a single party can ever unwrap the token and see what\\'s inside, as the token can be used only once and cannot be unwrapped by anyone else, even the root user or the creator of the token. This provides a way to securely distribute secrets to the intended recipients and detect any tampering or interception along the way5. The other options are not benefits of response wrapping: Log every use of a secret: Response wrapping does not log every use of a secret, as the secret is not directly exposed to the client or the network. However, Vault does log the creation and deletion of the response-wrapping token, and the client can use the audit device to log the unwrapping operation6. Load balance secret generation across a Vault cluster: Response wrapping does not load balance secret generation across a Vault cluster, as the secret is generated by the Vault server that receives the request and the response- wrapping token is bound to that server. However, Vault does support high availability and replication modes that can distribute the load and improve the performance of the cluster7. Provide error recovery to a secret so it is not corrupted in transit: Response wrapping does not provide error recovery to a secret so it is not corrupted in transit, as the secret is encrypted and stored in the cubbyhole of the token and cannot be modified or corrupted by anyone. However, if the token is lost or expired, the secret cannot be recovered either, so the client should have a backup or retry mechanism to handle such cases. References:
5(https://developer.hashicorp.com/vault/docs/concepts/response- wrapping),
6(https://developer.hashicorp.com/vault/docs/secrets), 7(https://developer.hashi corp.com/vault/docs/secrets),
(https://developer.hashicorp.com/vault/ tutorials/secrets- management/cubbyhole-response-wrapping)

QUESTION 8

You have a 2GB Base64 binary large object (blob) that needs to be encrypted. Which of the following best describes the transit secrets engine?

A. A data key encrypts the blob locally, and the same key decrypts the blob locally.

B. To process such a large blob. Vault will temporarily store it in the storage backend.

C. Vault will store the blob permanently. Be sure to run Vault on a compute optimized machine

D. The transit engine is not a good solution for binaries of this size.

Correct Answer: D

The transit secrets engine is not a good solution for binaries of this size, because it is designed to handle cryptographic functions on data in-transit, not data at-rest. The transit secrets engine does not store any data sent to it, so it would require sending the entire 2GB blob to Vault for encryption or decryption, which would be inefficient and impractical. A better solution would be to use the transit secrets engine to generate a data key, which is a high-entropy key that can be used to encrypt or decrypt data locally. The data key can be returned in plaintext or wrapped by another key, depending on the use case. This way, the transit secrets engine only handles the encryption or decryption of the data key, not the data itself, and the data can be stored in any primary data store. References: Transit - Secrets Engines | Vault | HashiCorp Developer, Encryption as a service: transit secrets engine | Vault | HashiCorp Developer

---

QUESTION 9

Which Vault secret engine may be used to build your own internal certificate authority?

A. Transit

B. PKI

C. PostgreSQL D. Generic

Correct Answer: B

The Vault secret engine that can be used to build your own internal certificate authority is the PKI secret engine. The PKI secret engine generates dynamic X.509 certificates on-demand, without requiring manual processes of generating private keys and CSRs, submitting to a CA, and waiting for verification and signing. The PKI secret engine can act as a root CA or an intermediate CA, and can issue certificates for various purposes, such as TLS, code signing, email encryption, etc. The PKI secret engine can also manage the certificate lifecycle, such as rotation, revocation, renewal, and CRL generation. The PKI secret engine can also integrate with external CAs, such as Venafi or Entrust, to delegate the certificate issuance and management. References: PKI - Secrets Engines | Vault | HashiCorp Developer, Build Your Own Certificate Authority (CA) | Vault - HashiCorp Learn

---

QUESTION 10

When an auth method is disabled all users authenticated via that method lose access.

A. True

B. False

Correct Answer: A

The statement is true. When an auth method is disabled, all users authenticated via that method lose access. This is because the tokens issued by the auth method are automatically revoked when the auth method is disabled. This prevents the users from performing any operation in Vault using the revoked tokens. To regain access, the users have to authenticate again using a different auth method that is enabled and has the appropriate policies attached. References: Auth Methods | Vault | HashiCorp Developer, auth disable - Command | Vault | HashiCorp Developer

---

QUESTION 11

What environment variable overrides the CLI\\'s default Vault server address?

A. VAULT_ADDR

B. VAULT_HTTP_ADORESS

C. VAULT_ADDRESS

D. VAULT _HTTPS_ ADDRESS

Correct Answer: B

The environment variable VAULT_ADDR overrides the CLI\\'s default Vault server address. The VAULT_ADDR environment variable specifies the address of the Vault server that is used to communicate with Vault from other applications or processes. By setting this variable, you can avoid hard-coding the Vault server address in your code or configuration files, and you can also use different addresses for different environments or scenarios. For example, you can use a local development server for testing purposes, and a production server for deploying your application. References: Commands (CLI) | Vault | HashiCorp Developer, Vault Agent - secrets as environment variables | Vault | HashiCorp Developer

---

QUESTION 12

Which of the following describes the Vault\\'s auth method component?

A. It verifies a client against an internal or external system, and generates a token with the appropriate policies attached

B. It verifies a client against an internal or external system, and generates a token with root policy

C. It is responsible for durable storage of client tokens

D. It dynamically generates a unique set of secrets with appropriate permissions attached

Correct Answer: A

The Vault\\'s auth method component is the component that performs authentication and assigns identity and policies to a client. It verifies a client against an internal or external system, and generates a token with the appropriate policies attached. The token can then be used to access the secrets and resources that are authorized by the policies. Vault supports various auth methods, such as userpass, ldap, aws, kubernetes, etc., that can integrate with different identity providers and systems. The auth method component can also handle token renewal and revocation, as well as identity grouping and aliasing. References: Auth Methods | Vault | HashiCorp Developer, Authentication - Concepts | Vault | HashiCorp Developer