Vendor: CyberArk

Exam Code:SECRET-SEN

Exam Name: CyberArk Sentry - Secrets Manager

Version: Demo

QUESTION 1

When installing the Vault Conjur Synchronizer, you see this error:

Forbidden

Logon Token is Empty ?Cannot logon

Unauthorized

What must you ensure to remediate the issue?

A. This admin user must not be logged in to other sessions during the Vault Conjur Synchronizer installation process.

B. You specified the correct url for Conjur and it is listed as a SAN on that url\\'s certificate.

C. You correctly URI encoded the url in the installation script.

D. You ran powershell as Administrator and there is sufficient space on the server on which you are running the installation.

Correct Answer: A

This error occurs when the Vault Conjur Synchronizer installation script tries to log in to the Vault using the admin user credentials, but the admin user is already logged in to other sessions. The Vault has a limit on the number of concurrent sessions per user, and the default value is one. Therefore, the installation script fails to authenticate the admin user and returns the error message: Forbidden Logon Token is Empty - Cannot logon Unauthorized. To remediate the issue, the admin user must log out of any other sessions before running the installation script, or increase the limit on the number of concurrent sessions per user in the Vault configuration file12. References: = Troubleshoot CyberArk Vault Synchronizer 1, Error: Forbidden Logon Token is Empty - Cannot logon Unauthorized Vault.ini File Parameters 2, ConcurrentSessionsPerUser

---

QUESTION 2

DRAG DROP

You want to allow retrieval of a secret with the CCP. The safe and the required secrets already exist.

Assuming the CCP is installed, arrange the steps in the correct sequence.

Select and Place:

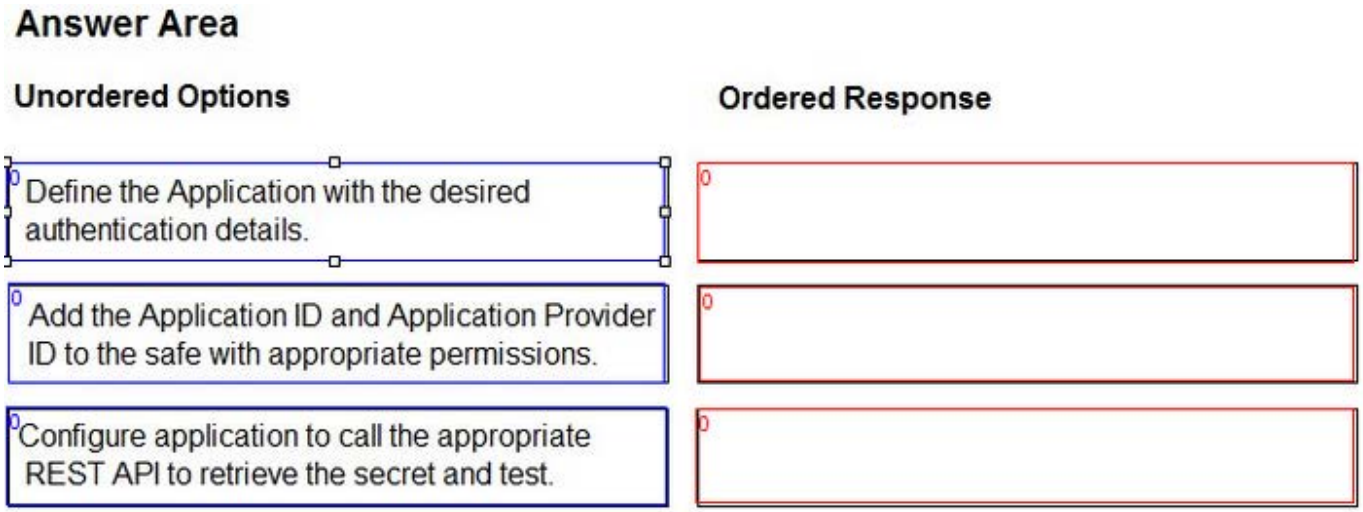

Correct Answer:

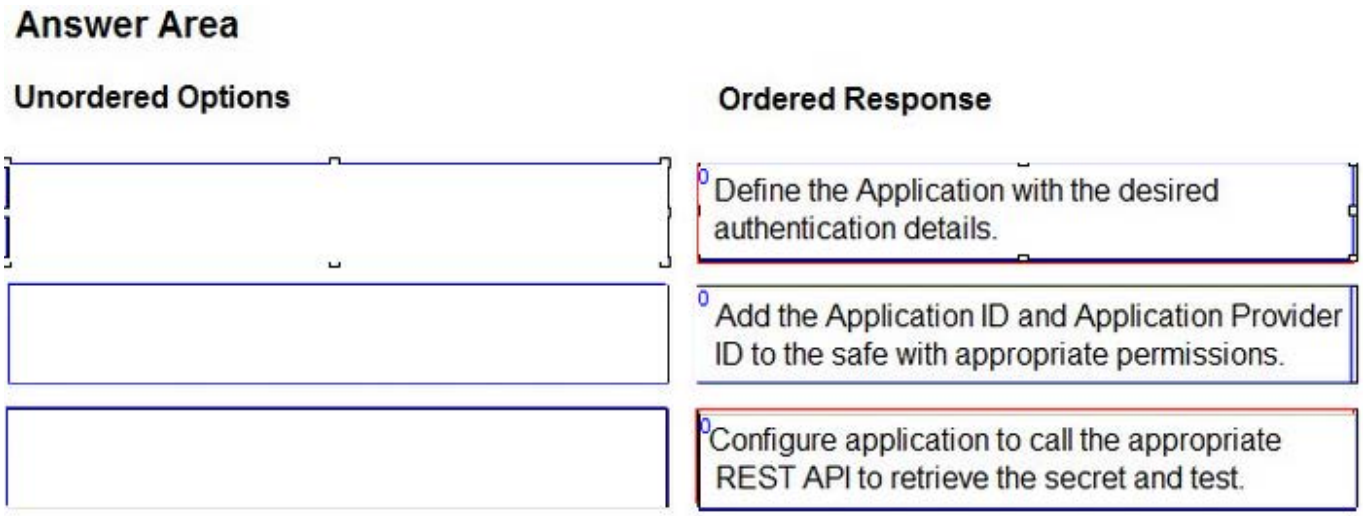

The correct order of the steps is: Define the Application with the desired authentication details Add the Application ID and Application Provider ID to the safe with appropriate permissions Configure application to call the appropriate REST API to retrieve the secret and test To allow an application to retrieve a secret with the CCP, the following steps are required: Define the Application with the desired authentication details: This step involves creating an Application object in the Vault with a unique Application ID and an Application Provider ID. The Application Provider ID is used to identify the CCP instance that will serve the request. The Application object also defines the authentication method and parameters that the application will use to connect to the CCP, such as certificate, password, or AppRole. Add the Application ID and Application Provider ID to the safe with appropriate permissions: This step involves granting the Application object the necessary permissions to access the safe and the secret that it needs. The Application ID and the Application Provider ID are added as members of the safe with at least List and Retrieve permissions. The secret name or ID can also be specified as a restriction to limit the access to a specific secret within the safe. Configure application to call the appropriate REST API to retrieve the secret and test: This step involves configuring the application to send a REST API request to the CCP endpoint with the required parameters, such as the Application ID, the Application Provider ID, the safe name, and the secret name or ID. The application should also provide the authentication credentials or token that match the method defined in the Application object. The application should receive a JSON response from the CCP with the secret value and other metadata. The application should test the connection and the secret retrieval before deploying to production. References: CyberArk Secrets Manager Sentry - Secrets Manager - Sample Items and Study Guide Sentry - Secrets Secrets Management Essentials for Developers

---

QUESTION 3

An application owner reports that their application is suddenly receiving an incorrect password. CPM logs show the password was recently changed, but the value currently being retrieved by the application is a different value. The Vault Conjur Synchronizer service is running.

What is the most likely cause of this issue?

A. The Vault Conjur Synchronizer is not configured with the DR Vault IP address and there has been a failover event.

B. Dual Accounts are in use, but after the CPM changed the password for the Inactive account, it accidentally updated the password for the Active account instead.

C. The CPM is writing password changes to the Primary Vault while the Vault Conjur Synchronizer is configured to replicate from the DR Vault.

D. The application has been configured to retrieve the wrong password.

Correct Answer: C

This is the most likely cause of this issue because it creates a discrepancy between the passwords stored in the Primary Vault and the DR Vault, which affects the Vault Conjur Synchronizer service (Synchronizer) and the application. The

Synchronizer is a service that synchronizes secrets from the CyberArk Vault to the Conjur database. The application is a client that retrieves secrets from the Conjur database using the Conjur REST API. The CPM is a component that

manages the lifecycle of the passwords stored in the CyberArk Vault, such as changing, verifying, and reconciling them. If the CPM is writing password changes to the Primary Vault while the Synchronizer is configured to replicate from the

DR Vault, the following scenario may occur:

The CPM changes the password for an account in the Primary Vault and updates the password value in the Vault database.

The Synchronizer does not detect the password change in the DR Vault, as the DR Vault database has not been updated yet with the new password value. The Synchronizer does not sync the new password value to the Conjur database, as

it assumes that the password value in the DR Vault database is the latest and correct one.

The application requests the password value from the Conjur database and receives the old password value, which is different from the new password value in the Primary Vault database.

The application tries to use the old password value to access the target platform or device and fails, as the target platform or device expects the new password value. This answer is based on the CyberArk Secrets Manager documentation1

and the CyberArk Secrets Manager training course2.

QUESTION 4

DRAG DROP

You are installing a Credential Provider on a Linux host. Arrange the installation steps in the correct sequence.

Select and Place:

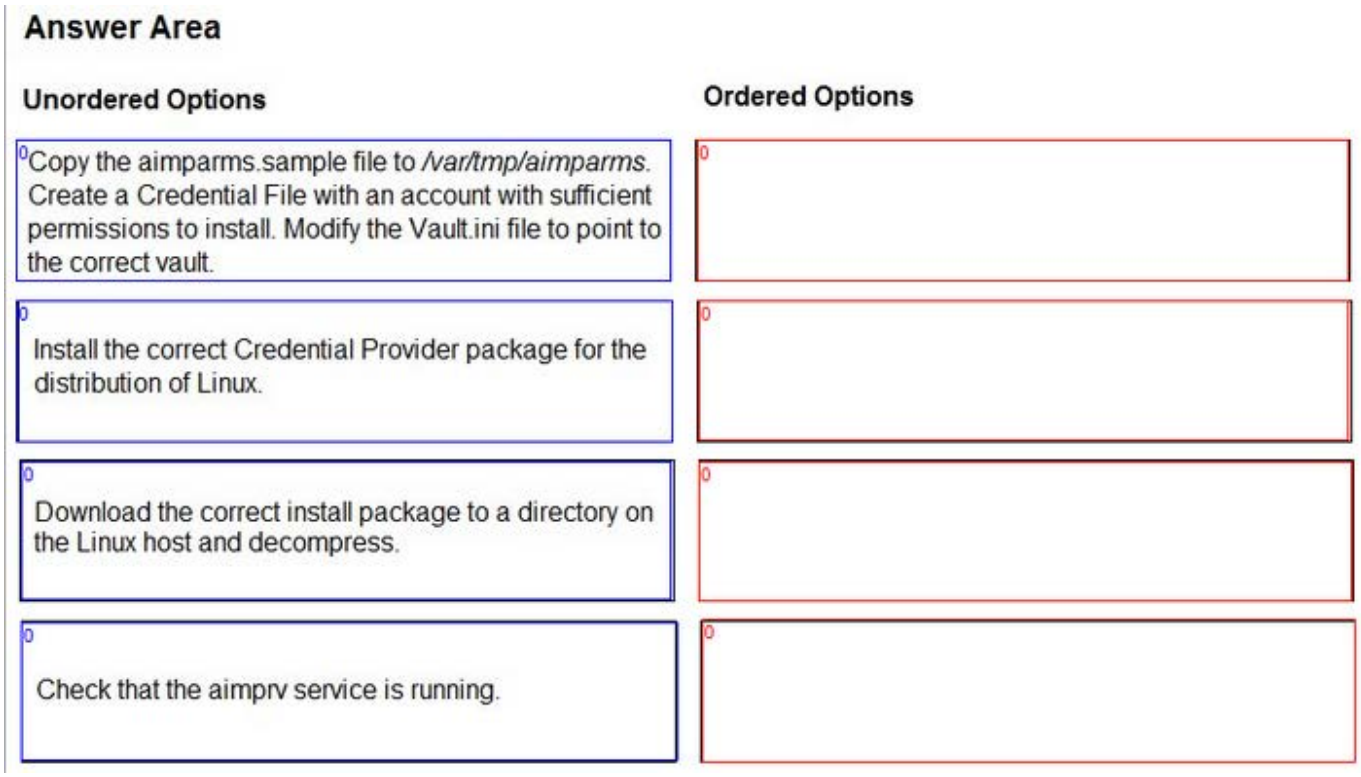

Correct Answer:

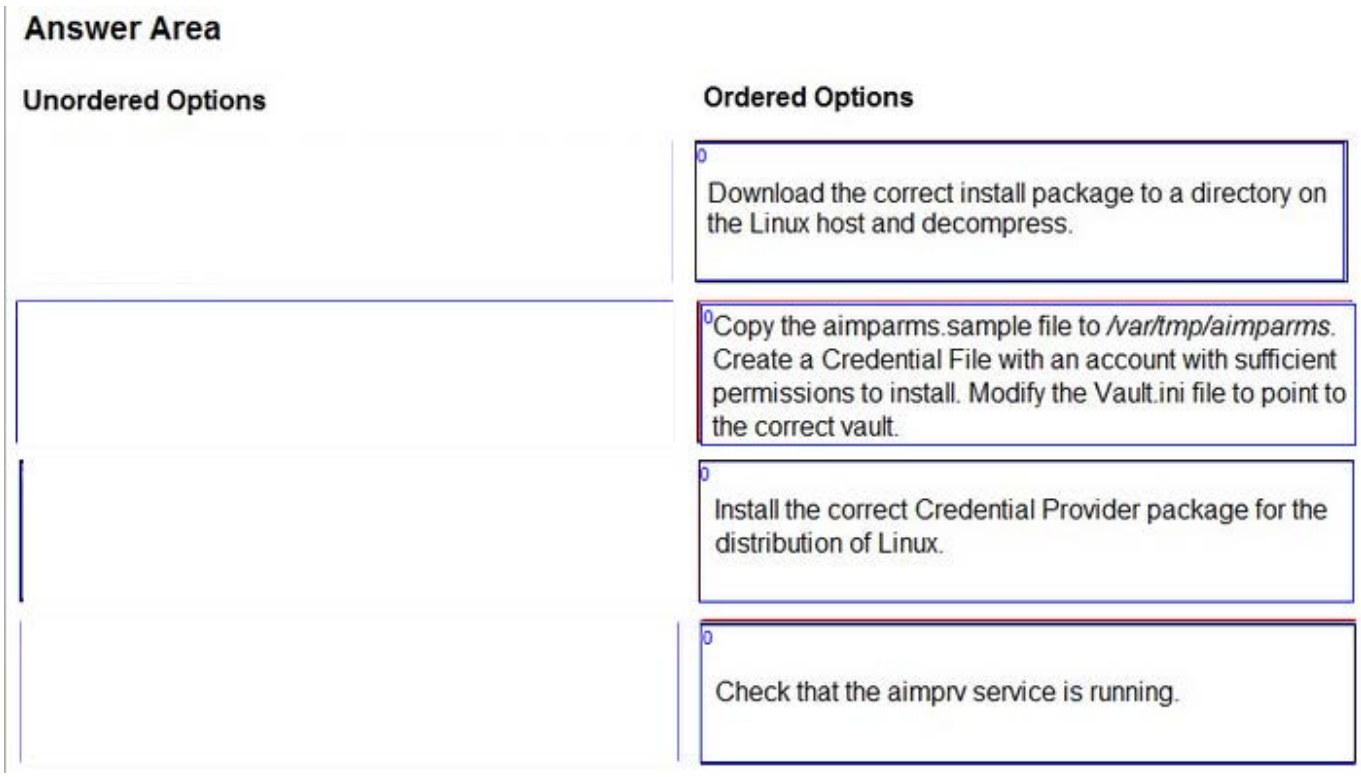

The correct sequence of installation steps for a Credential Provider on a Linux host is as follows:

Download the correct install package to a directory on the Linux host and decompress1.

Copy the aimparms.sample file to /var/tmp/aimparms. Create a Credential File with an account with sufficient permissions to install. Modify the Vault.ini file to point to the correct vault2.

Install the correct Credential Provider package for the distribution of Linux using the command: rpm -ivh CARKaim-..rpm2. Check that the aimprv service is running using the command: service aimprv

status2.

References: 1: Download the Credential Provider 2: Install Credential Provider on Linux / AIX

QUESTION 5

What is the correct process to upgrade the CCP Web Service?

A. Run "sudo yum update aimprv" from the CLI.

B. Double-click the Credential Provider installer executable and select upgrade.

C. Double-click the AimWebService.msi and select upgrade.

D. Uninstall and reinstall the CCP Web Service.

Correct Answer: D

The correct process to upgrade the CCP Web Service is D. Uninstall and reinstall the CCP Web Service. The CCP Web Service is a component of the CyberArk Central Credential Provider (CCP) that enables applications to retrieve secrets from the CyberArk Vault using REST API calls. To upgrade the CCP Web Service, you need to first uninstall the existing CCP Web Service from the Windows Server Manager or the Control Panel, and then reinstall the CCP Web Service using the latest installation package from the CyberArk website. The installation package contains both the Credential Provider and the CCP Web Service components, and you need to run the AimWebService.msi file to install the CCP Web Service. You also need to make sure that the CCP Web Service has the correct configuration and permissions, and that the CyberArk CRL (Certificate Revocation List) is open from the CCP server. The other options are not correct processes to upgrade the CCP Web Service. Running "sudo yum update aimprv" from the CLI is a command to update the Credential Provider on Linux, not the CCP Web Service on Windows. Double-clicking the Credential Provider installer executable and selecting upgrade is a process to upgrade the Credential Provider on Windows, not the CCP Web Service. Double-clicking the AimWebService.msi and selecting upgrade is not a valid option, as the CCP Web Service does not support an upgrade option, and you need to uninstall it first before reinstalling it. References: Upgrade the Central Credential Provider (CCP) - CyberArk, Section "Upgrade the Central Credential Provider (CCP)" Central Credential Provider web service configuration - CyberArk, Section "Central Credential Provider web service configuration"

QUESTION 6

When working with Credential Providers in a Privileged Cloud setting, what is a special consideration?

A. If there are installation issues, troubleshooting may need to involve the Privileged Cloud support team.

B. Credential Providers are not supported in a Privileged Cloud setting.

C. The AWS Cloud account number must be defined in the file main appprovider.conf.. found in the AppProviderConf Safe.

D. Debug logging for Credential Providers deployed in a Privileged Cloud setting can inadvertently exhaust available

disk space.

Correct Answer: A

Credential Providers are tools that enable applications to securely retrieve credentials from CyberArk Secrets Manager without hard-coding or storing them in files. Credential Providers can be installed on application servers or on a central server that acts as a proxy for multiple applications. Credential Providers can integrate with Privileged Cloud, which is a cloud-based solution that provides privileged access management as a service. Privileged Cloud integrates with Secrets Manager Credential Providers to manage application credentials as privileged accounts within Privileged Cloud. When working with Credential Providers in a Privileged Cloud setting, a special consideration is that if there are installation issues, troubleshooting may need to involve the Privileged Cloud support team. This is because the installation of Credential Providers in a Privileged Cloud setting requires some additional steps and configurations that are performed by the Privileged Cloud support team. For example, the Privileged Cloud support team needs to configure the connection between Privileged Cloud and Credential Providers, and provide the necessary certificates and keys for secure communication. Therefore, if there are any problems or errors during the installation process, the Privileged Cloud support team may need to assist with the troubleshooting and resolution. The other options are not correct. Credential Providers are supported in a Privileged Cloud setting, as described in the Secrets Manager Credential Providers integration documentation1. The AWS Cloud account number does not need to be defined in the file main appprovider.conf.. found in the AppProviderConf Safe. This file is used to configure the Credential Provider settings, such as the Privileged Cloud URL, the application ID, and the SSL options. The AWS Cloud account number is not relevant for this file. Debug logging for Credential Providers deployed in a Privileged Cloud setting can be enabled or disabled by the Privileged Cloud support team, as described in the Credential Provider installation documentation2. Debug logging can help with troubleshooting and diagnostics, but it does not necessarily exhaust available disk space, as the log files can be rotated and archived. References: Secrets Manager Credential Providers integration; Credential Provider installation

---

QUESTION 7

After manually failing over to your disaster recovery site (Site B) for testing purposes, you need to failback to your primary site (Site A).

Which step is required?

A. Contact CyberArk for a new license file.

B. Reconfigure the Vault Conjur Synchronizer to point to the new Conjur Leader.

C. Generate a seed for the new Leader to be deployed in Site A.

D. Trigger autofailover to promote the Standby in Site A to Leader.

Correct Answer: C

According to the CyberArk Sentry Secrets Manager documentation1, the steps to failback to the primary site after a manual failover to the disaster recovery site are as follows: On the DR site, stop the Conjur Leader node using the command docker stop . On the primary site, generate a seed for the new Leader node using the command evoke seed leader . This will create a file named .tar in the current directory. On the primary site, copy the Leader seed file to the new Leader server using the command scp .tar :.tar On the new Leader server, create a new container using the same name as the one you just stopped, and load the Leader seed file using the command docker run --name -d --restart=always -v /var/log/conjur:/var/log/conjur -v /opt/conjur/backup:/opt/conjur/backup -p "443:443" -p "5432:5432" -p "1999:1999" cyberark/conjur:latest seed fetch .tar On the new Leader server, configure the Conjur Leader node using the command evoke configure leader -h -p On the new Leader server, reconfigure the Vault Conjur Synchronizer to point to the new Conjur Leader using the command evoke vault sync set On the DR site, generate a seed for the new Standby node using the command evoke seed standby . This will create a file named .tar in the current directory. On the DR site, copy the Standby seed file to the new Standby server using the command scp .tar :.tar On the new Standby

server, create a new container using the same name as the one you just stopped, and load the Standby seed file using the command docker run --name -d --restart=always -v /var/log/conjur:/var/log/ conjur -v /opt/conjur/backup:/opt/conjur/backup -p "443:443" -p "5432:5432" -p "1999:1999" cyberark/conjur:latest seed fetch .tar On the new Standby server, re-enroll the node to the cluster using the command evoke cluster enroll The other options are not correct, as they are either unnecessary or incorrect. Contacting CyberArk for a new license file is not required, as the license is valid for both sites. Reconfiguring the Vault Conjur Synchronizer to point to the new Conjur Leader is a step that should be done on the new Leader server, not on the DR site. Triggering autofailover to promote the Standby in Site A to Leader is not possible, as the Standby node is not aware of the manual failover and will not accept the promotion request.

QUESTION 8

A customer has 100 .NET applications and wants to use Summon to invoke the application and inject secrets at run time.

Which change to the NET application code might be necessary to enable this?

A. It must be changed to include the REST API calls necessary to retrieve the needed secrets from the CCP.

B. It must be changed to access secrets from a configuration file or environment variable.

C. No changes are needed as Summon brokers the connection between the application and the backend data source through impersonation.

D. It must be changed to include the host API key necessary for Summon to retrieve the needed secrets from a Follower

Correct Answer: B

Summon is a utility that allows applications to access secrets from a variety of trusted stores and export them as environment variables to a sub-process environment. Summon does not require any changes to the application code to retrieve secrets from the CyberArk Central Credential Provider (CCP), as it uses a provider plugin that handles the communication with the CCP. However, the application code must be able to access secrets from a configuration file or environment variable, as these are the methods that Summon uses to inject secrets into the application. Summon reads a secrets.yml file that defines the secrets that the application needs and maps them to environment variables. Then, Summon fetches the secrets from the CCP using the provider plugin and exports them as environment variables to the application sub-process. The application can then read the secrets from the environment variables as if they were hard-coded in the configuration file. References: Summon-inject secrets, .NET Application Password SDK

QUESTION 9

A customer requires high availability in its AWS cloud infrastructure.

What is the minimally viable Conjur deployment architecture to achieve this?

A. one Follower in each AZ. load balancer for the region

B. two Followers in each region, load balanced for the region

C. two Followers in each AZ. load balanced for the region

D. two Followers in each region, load balanced across all regions

Correct Answer: A

According to the CyberArk Sentry Secrets Manager documentation, Conjur is a secrets management solution that consists of a leader node and one or more follower nodes. The leader node is responsible for managing the secrets, policies,

and audit records, while the follower nodes are read-only replicas that can serve secrets requests from applications. To achieve high availability in AWS cloud infrastructure, the minimally viable Conjur deployment architecture is to have one

follower in each availability zone (AZ) and a load balancer for the region. This way, if one AZ fails, the applications can still access secrets from another AZ through the load balancer. Having two followers in each region, load balanced for the

region, is not enough to ensure high availability, as a regional outage can affect both followers. Having two followers in each AZ, load balanced for the region, is more than necessary, as one follower per AZ can handle the secrets requests.

Having two followers in each region, load balanced across all regions, is not feasible, as Conjur does not support cross-region replication.

References: 1: Conjur Architecture 2: Deploying Conjur on AWS

QUESTION 10

What is a possible Conjur node role change?

A. A Standby may be promoted to a Leader.

B. A Follower may be promoted to a Leader.

C. A Standby may be promoted to a Follower.

D. A Leader may be demoted to a Standby in the event of a failover.

Correct Answer: A

According to the CyberArk Sentry Secrets Manager documentation, Conjur is a secrets management solution that consists of a leader node and one or more follower nodes. The leader node is responsible for managing the secrets, policies,

and audit records, while the follower nodes are read-only replicas that can serve secrets requests from applications. Additionally, Conjur supports a standby node, which is a special type of follower node that can be promoted to a leader node

in case of a leader failure. A standby node is synchronized with the leader node and can take over its role in a disaster recovery scenario. A possible Conjur node role change is when a standby node is promoted to a leader node, either

manually or automatically, using the auto-failover feature. A follower node cannot be promoted to a leader node, as it does not have the same data and functionality as the leader node. A standby node cannot be promoted to a follower node,

as it already has the same capabilities as a follower node, plus the ability to become a leader node. A leader node cannot be demoted to a standby node in the event of a failover, as it would lose its data and functionality and would not be able

to resume its role as a leader node.

References:

1: Conjur Architecture

2: Deploying Conjur on AWS

3: Auto-failover

QUESTION 11

What is the most maintenance-free way to ensure a Conjur host\\'s access reflects any changes made to accounts in a safe in the CyberArk vault?

A. Write an automation script to update and load the host\\'s policy using PATCH/update.

B. Use yami anchor [and] and wildcard (*) syntax to maintain its list of permission grants.

C. Grant the consumers group/role created by the Synchronizer for the Safe to the host.

D. Use PVWA to add the Conjur host ID as a member of the Safe.

Correct Answer: C

The most maintenance-free way to ensure a Conjur host\\'s access reflects any changes made to accounts in a safe in the CyberArk vault is to grant the consumers group/role created by the Synchronizer for the Safe to the host. This means

that the host will inherit the read and execute permissions on all the secrets in the Safe from the consumers group/role, and will automatically get access to any new or updated secrets in the Safe without requiring any manual intervention or

policy changes. The consumers group/role is created by the Vault Conjur Synchronizer, which is a service that synchronizes secrets between the CyberArk vault and Conjur. The Synchronizer creates a policy branch for each Safe in Conjur,

and assigns the consumers group/role to have read and execute permissions on all the secrets in the Safe. The Synchronizer also creates a delegation policy for each Safe, which allows the Safe admins to grant permissions to other users,

hosts, groups, or layers12.

The other options are not the most maintenance-free ways to ensure a Conjur host\\'s access reflects any changes made to accounts in a safe in the CyberArk vault. Writing an automation script to update and load the host\\'s policy using

PATCH/update may work, but it requires additional effort and maintenance to ensure the script is always running and up to date with the changes in the Safe. Using yami anchor [and] and wildcard (*) syntax to maintain its list of permission

grants may simplify the policy writing, but it still requires manual editing and loading of the policy whenever a new secret is added or removed from the Safe. Using PVWA to add the Conjur host ID as a member of the Safe may not be

possible or advisable, as the PVWA is designed for managing human users and not Conjur hosts, and it may not have the necessary integration or authorization to do so3.

References:

Vault Conjur Synchronizer 1, Synchronizer Policy Structure Grant permissions on secrets 2, Grant role permissions on all secrets in a Safe Privileged Access Manager - Self-Hosted 3, Privileged Web Access (PVWA)

QUESTION 12

Refer to the exhibit.

"replication_status": { "pg_stat_replication": [{ "username": "conjur-follower.mycompany.local",
"application_name": "follower_conjur_follower_mycompany_local_c63e36c427c3".
"client_addr": "12.16.23.10", "backend_start": "2020-11-13 22:45:04 +0000" "state":
"streaming", "sent_lsn": "0/30021C8", "replay_lsn": "0/30021C8", "sync_priority":0,
"sync_state": "async", "sent_lsn_bytes": 50340296, "replay_lsn_bytes": 50340296,
"replication_lag_bytes":0 }], "pg_current_xlog_location": "0/30021C8",
"pg_current_xlog_location_bytes": 50340296}

How can you confirm that the Follower has a current copy of the database?

A. Compare the pgcurrentxlog_locationlocation from the Leader to the Follower you need to validate against.

B. Count the number of components in pgstartreplication and compare this to the total number of Followers in the deployment.

C. Validate that the Follower container ID matches the node in the info endpoint on the Leader.

D. Retrieve the credential from a test application on the Leader cluster; then retrieve against the Follower and compare if they are accurate.

Correct Answer: A

The exhibit shows a JSON object that contains the replication status of a database in a Secrets Manager cluster. Secrets Manager is a secrets management solution that securely stores and manages secrets and credentials used by applications, DevOps tools, and other systems. Secrets Manager can be deployed in a cluster mode, which consists of a Leader node and one or more Follower nodes. The Leader node is the primary node that handles all write operations and coordinates the replication of data to the Follower nodes. The Follower nodes are read-only nodes that replicate data from the Leader node and serve requests from clients and applications that need to retrieve secrets or perform other read-only operations. To confirm that the Follower has a current copy of the database, you can compare the pgcurrentxlog_locationlocation from the Leader to the Follower you need to validate against. The pgcurrentxlog_locationlocation is a property that indicates the current position of the write-ahead log (WAL) in the database. The WAL is a mechanism that records all changes made to the database in a sequential log file, before they are applied to the actual data files. The WAL ensures the durability and consistency of the database in case of a crash or a power failure. The WAL also enables the replication of data from the Leader node to the Follower nodes, by streaming the WAL records to the Follower nodes and applying them to their local databases. By comparing the pgcurrentxlog_locationlocation from the Leader to the Follower, you can determine how far behind the Follower is from the Leader in terms of the WAL records. If the pgcurrentxlog_locationlocation values are identical or very close, it means that the Follower has a current copy of the database, and that the replication is working properly. If the pgcurrentxlog_locationlocation values are different or far apart, it means that the Follower has an outdated copy of the database, and that there is a replication lag or a replication failure. In that case, you may need to troubleshoot the replication issue and resolve it as soon as possible. References: Secrets Manager Cluster Installation; Secrets Manager Cluster Configuration; Write-Ahead Logging - PostgreSQL Documentation