NCM-MCI-6.5 Q&As

Nutanix Certified Master - Multicloud Infrastructure (NCM-MCI)v6.5

QUESTION 1

CORRECT TEXT

Task 15

An administrator found a CentOS VM, Cent_Down, on the cluster with a corrupted network stack. To correct the issue, the VM will need to be restored from a previous snapshot to become reachable on the network again.

VM credentials:

Username: root

Password: nutanix/4u

Restore the VM and ensure it is reachable on the network by pinging 172.31.0.1 from the VM.

Power off the VM before proceeding.

A. Answer: See the for step by step solution.

Correct Answer: A

To restore the VM and ensure it is reachable on the network, you can follow these steps:

Log in to the Web Console of the cluster where the VM is running. Click on Virtual Machines on the left menu and find Cent_Down from the list. Click on the power icon to power off the VM.

Click on the snapshot icon next to the power icon to open the Snapshot Management window.

Select a snapshot from the list that was taken before the network stack was corrupted. You can use the date and time information to choose a suitable snapshot. Click on Restore VM and confirm the action in the dialog box. Wait for the

restore process to complete.

Click on the power icon again to power on the VM. Log in to the VM using SSH or console with the username and password provided. Run the command ping 172.31.0.1 to verify that the VM is reachable on the network. You should see a

reply from the destination IP address.

Go to VMS from the prism central gui

Select the VMand go to More -> Guest Shutdown

Go to Snapshots tab and revert to latest snapshot available power on vm and verify if ping is working

QUESTION 2

CORRECT TEXT Task 6 An administrator has requested the commands needed to configure traffic segmentation on an unconfigured node. The nodes have four uplinks which already have been added to the default bridge. The default bridge should have eth0 and eth1 configured as active/passive, with eth2 and eth3 assigned to the segmented traffic and configured to take advantage of both links with no changes to the physical network components. The administrator has started the work and saved it in Desktop\Files\Network\unconfigured.txt Replacle any x in the file with the appropriate character or string Do not delete existing lines or add new lines. Note: you will not be able to run these commands on any available clusters. Unconfigured.txt manage_ovs --bond_name brX-up --bond_mode xxxxxxxxxxx --interfaces ethX,ethX update_uplinks manage_ovs --bridge_name brX-up --interfaces ethX,ethX --bond_name bond1 -- bond_mode xxxxxxxxxxx update_uplinks

A. Answer: See the for step by step solution.

Correct Answer: A

To configure traffic segmentation on an unconfigured node, you need to run the following commands on the node: manage_ovs --bond_name br0-up --bond_mode active-backup --interfaces eth0,eth1 update_uplinks manage_ovs --bridge_name br0-up --interfaces eth2,eth3 --bond_name bond1 --bond_mode balance-slb update_uplinks These commands will create a bond named br0-up with eth0 and eth1 as active and passive interfaces, and assign it to the default bridge. Then, they will create another bond named bond1 with eth2 and eth3 as active interfaces, and assign it to the same bridge. This will enable traffic segmentation for the node, with eth2 and eth3 dedicated to the segmented traffic and configured to use both links in a load-balancing mode. I have replaced the x in the file Desktop\Files\Network \unconfigured.txt with the appropriate character or string for you. You can find the updated file in Desktop\Files\Network\configured.txt.

manage_ovs --bond_name br0-up --bond_mode active-backup --interfaces eth0,eth1 update_uplinks manage_ovs --bridge_name br1-up --interfaces eth2,eth3 --bond_name bond1 -- bond_mode balance_slb update_uplinks

https://portal.nutanix.com/page/documents/solutions/details?targetId=BP-2071-AHV- Networking:ovs-command-line-configuration.html

QUESTION 3

CORRECT TEXT

Task 2

An administrator needs to configure storage for a Citrix-based Virtual Desktop infrastructure.

Two VDI pools will be created

Non-persistent pool names MCS_Pool for tasks users using MCS Microsoft Windows 10 virtual Delivery Agents (VDAs)

Persistent pool named Persist_Pool with full-clone Microsoft Windows 10 VDAs for power users

20 GiB capacity must be guaranteed at the storage container level for all power user VDAs

The power user container should not be able to use more than 100 GiB

Storage capacity should be optimized for each desktop pool.

Configure the storage to meet these requirements. Any new object created should include the name of the pool(s) (MCS and/or Persist) that will use the object.

Do not include the pool name if the object will not be used by that pool.

Any additional licenses required by the solution will be added later.

A. Answer: See the for step by step solution.

Correct Answer: A

To configure the storage for the Citrix-based VDI, you can follow these steps:

Log in to Prism Central using the credentials provided. Go to Storage > Storage Pools and click on Create Storage Pool. Enter a name for the new storage pool, such as VDI_Storage_Pool, and select the disks to include in the pool. You can

choose any combination of SSDs and HDDs, but for optimal performance, you may prefer to use more SSDs than HDDs.

Click Save to create the storage pool.

Go to Storage > Containers and click on Create Container. Enter a name for the new container for the non-persistent pool, such as MCS_Pool_Container, and select the storage pool that you just created, VDI_Storage_Pool, as the source.

Under Advanced Settings, enable Deduplication and Compression to reduce the storage footprint of the non-persistent desktops. You can also enable Erasure Coding if you have enough nodes in your cluster and want to save more space.

These settings will help you optimize the storage capacity for the non-persistent pool.

Click Save to create the container.

Go to Storage > Containers and click on Create Container again. Enter a name for the new container for the persistent pool, such as Persist_Pool_Container, and select the same storage pool, VDI_Storage_Pool, as the source.

Under Advanced Settings, enable Capacity Reservation and enter 20 GiB as the reserved capacity. This will guarantee that 20 GiB of space is always available for the persistent desktops. You can also enter 100 GiB as the advertised

capacity to limit the maximum space that this container can use. These settings will help you control the storage allocation for the persistent pool.

Click Save to create the container.

Go to Storage > Datastores and click on Create Datastore. Enter a name for the new datastore for the non-persistent pool, such as MCS_Pool_Datastore, and select NFS as the datastore type. Select the container that you just created,

MCS_Pool_Container, as the source.

Click Save to create the datastore.

Go to Storage > Datastores and click on Create Datastore again. Enter a name for the new datastore for the persistent pool, such as Persist_Pool_Datastore, and select NFS as the datastore type. Select the container that you just created,

Persist_Pool_Container, as the source.

Click Save to create the datastore.

The datastores will be automatically mounted on all nodes in the cluster. You can verify this by going to Storage > Datastores and clicking on each datastore. You should see all nodes listed under Hosts.

You can now use Citrix Studio to create your VDI pools using MCS or full clones on these datastores. For more information on how to use Citrix Studio with Nutanix Acropolis, seeCitrix Virtual Apps and Desktops on NutanixorNutanix virtualization environments.

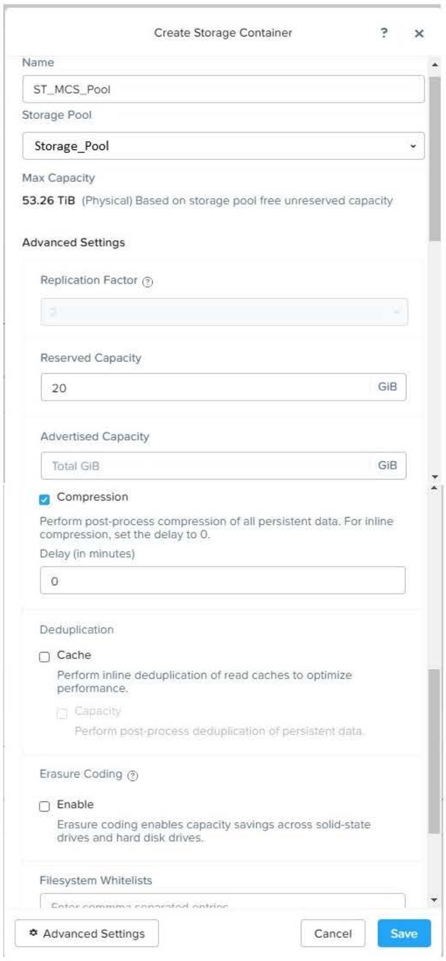

Create Storage Container

Name: ST_MCS_Pool

Storage Pool: Storage_Pool

Max Capacity

53.26 TiB (Physical) Based on storage pool free unreserved capacity

Advanced Settings

Replication Factor: 2

Reserved Capacity: 20 GiB

Advertised Capacity: Total GiB GiB

☑ Compression

Perform post-process compression of all persistent data. For inline compression, set the delay to 0.

Delay (in minutes): 0

Deduplication

□ Cache

Perform inline deduplication of read caches to optimize performance.

□ Capacity

Perform post-process deduplication of persistent data.

Erasure Coding

□ Enable

Erasure coding enables capacity savings across solid-state drives and hard disk drives.

Filesystem Whitelists

Enter comma separated entries

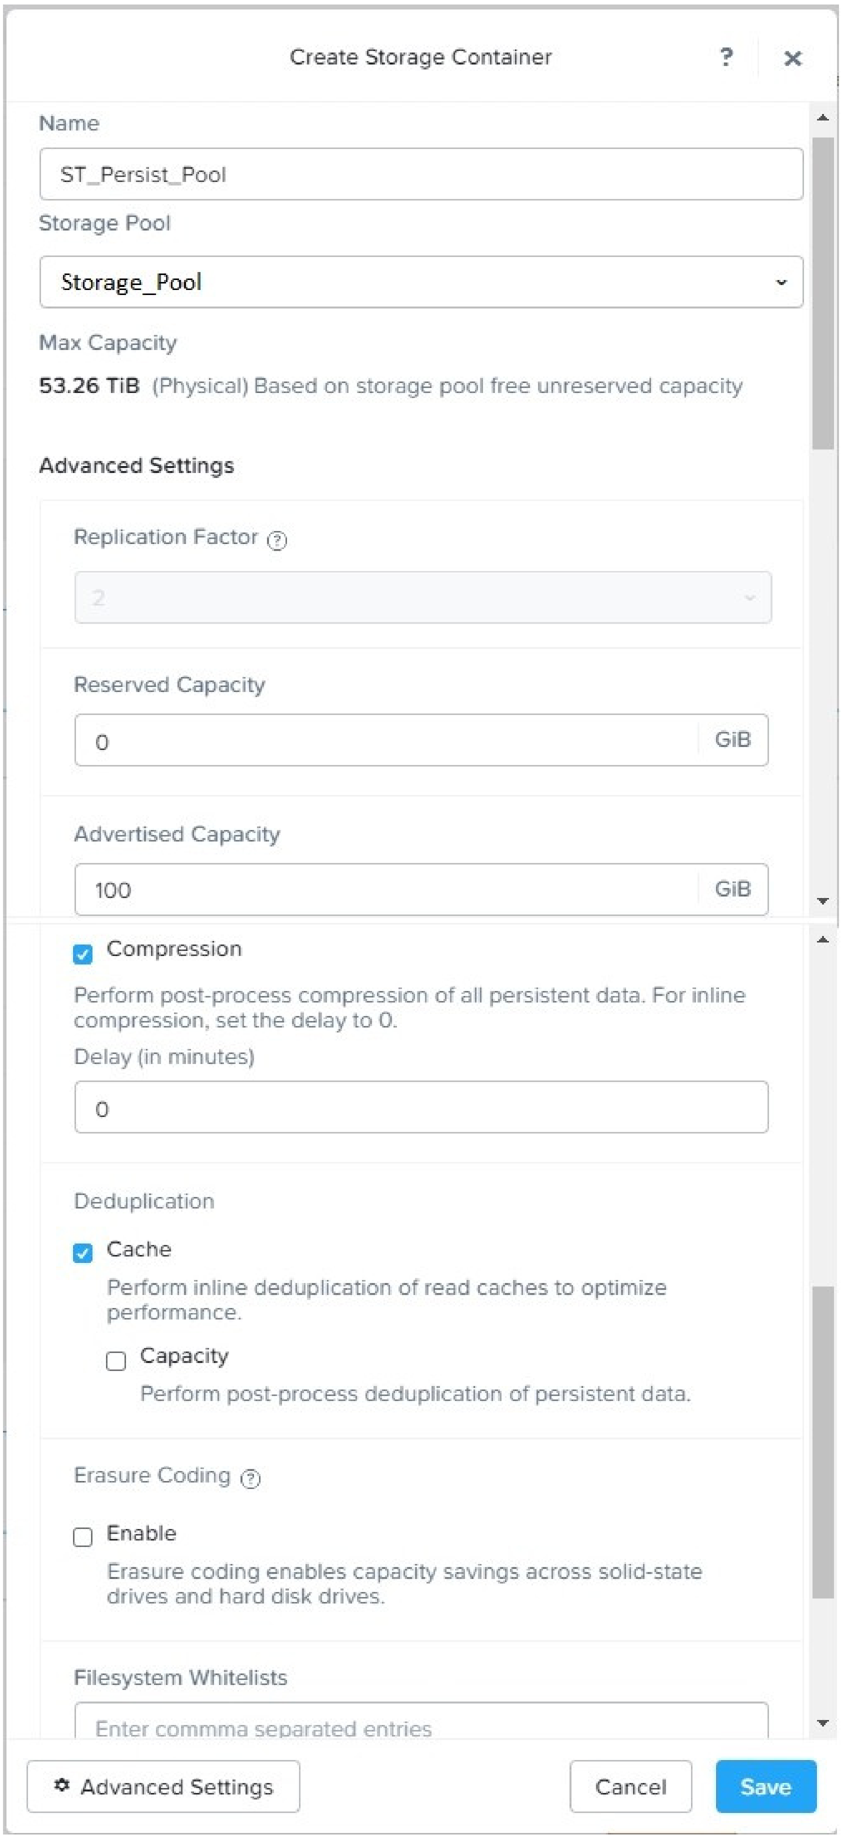

Create Storage Container

Name: ST_Persist_Pool

Storage Pool: Storage_Pool

Max Capacity

53.26 TiB (Physical) Based on storage pool free unreserved capacity

Advanced Settings

Replication Factor: 2

Reserved Capacity: 0 GiB

Advertised Capacity: 100 GiB

☑ Compression

Perform post-process compression of all persistent data. For inline compression, set the delay to 0.

Delay (in minutes): 0

Deduplication

☑ Cache

Perform inline deduplication of read caches to optimize performance.

☐ Capacity

Perform post-process deduplication of persistent data.

Erasure Coding

☐ Enable

Erasure coding enables capacity savings across solid-state drives and hard disk drives.

Filesystem Whitelists

Enter commma separated entries

https://portal.nutanix.com/page/documents/solutions/details?targetId=BP-2079-Citrix- Virtual-Apps-and-Desktops:bp-nutanix-storage-configuration.html

QUESTION 4

CORRECT TEXT

Task 7

An administrator has environment that will soon be upgraded to 6.5. In the meantime, they need to implement log and apply a security policy named Staging_Production, such that not VM in the Staging Environment can communicate with any

VM in the production Environment,

Configure the environment to satisfy this requirement.

Note: All other configurations not indicated must be left at their default values.

A. Answer: See the for step by step solution.

Correct Answer: A

To configure the environment to satisfy the requirement of implementing a security policy named Staging_Production, such that no VM in the Staging Environment can communicate with any VM in the production Environment, you need to do the following steps: Log in to Prism Central and go to Network > Security Policies > Create Security Policy. Enter Staging_Production as the name of the security policy and select Cluster A as the cluster. In the Scope section, select VMs as the entity type and add the VMs that belong to the Staging Environment and the Production Environment as the entities. You can use tags or categories to filter the VMs based on their environment. In the Rules section, create a new rule with the following settings: Direction: Bidirectional Protocol: Any Source: Staging Environment Destination: Production Environment Action: Deny Save the security policy and apply it to the cluster. This will create a security policy that will block any traffic between the VMs in the Staging Environment and the VMs in the Production Environment. You can verify that the security policy is working by trying to ping or access any VM in the Production Environment from any VM in the Staging Environment, or vice versa. You should not be able to do so.

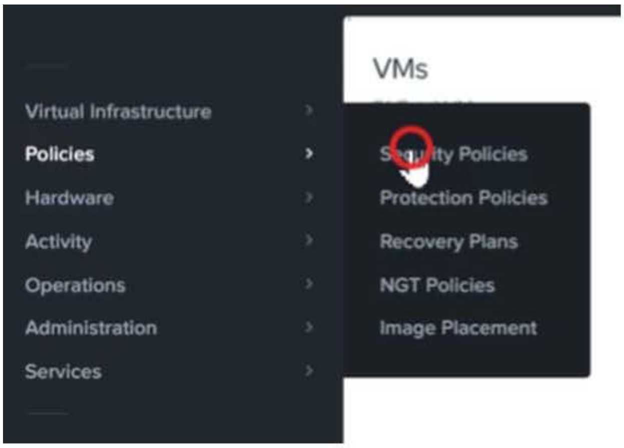

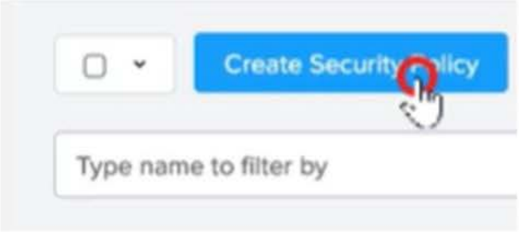

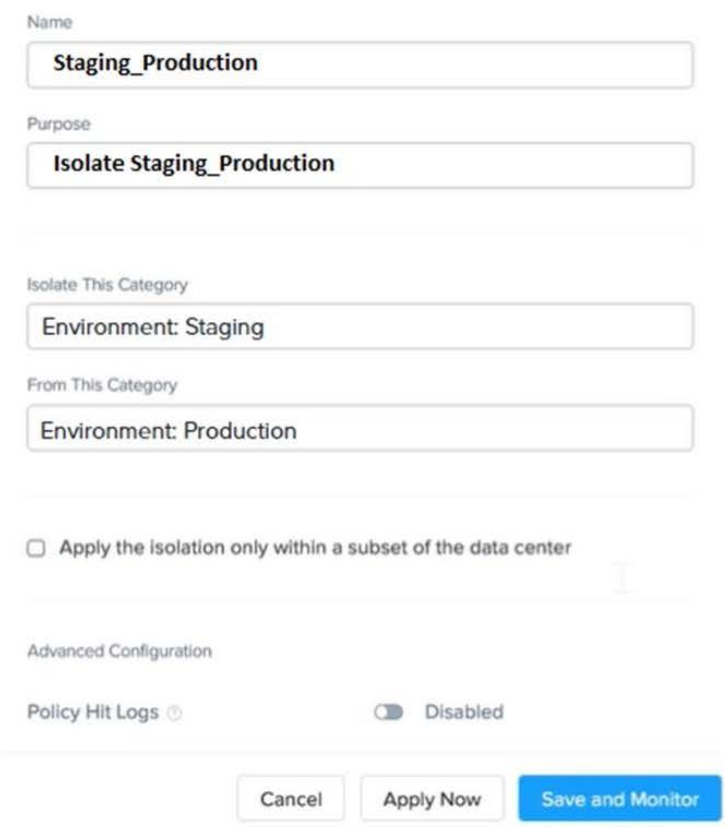

Name
Staging_Production

Purpose
Isolate Staging_Production

Isolate This Category
Environment: Staging

From This Category
Environment: Production

□ Apply the isolation only within a subset of the data center

Advanced Configuration

Policy Hit Logs Disabled

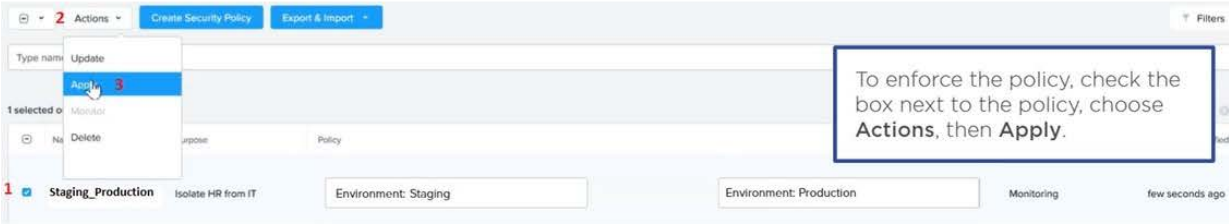

Staging_Production	Isolate HR from IT	Environment: Staging	Environment: Production	Monitoring	few seconds ago

To enforce the policy, check the box next to the policy, choose **Actions**, then **Apply**.

QUESTION 5

CORRECT TEXT

Task 8

Depending on the order you perform the exam items, the access information and credentials could change. Please refer to the other item performed on Cluster B if you have problems accessing the cluster.

The infosec team has requested that audit logs for API Requests and replication capabilities be enabled for all clusters for the top 4 severity levels and pushed to their syslog system using highest reliability possible. They have requested no other logs to be included.

Syslog configuration:

Syslog Name: Corp_syslog

Syslop IP: 34.69.43.123

Port: 514

Ensure the cluster is configured to meet these requirements.

A. Answer: See the for step by step solution.

Correct Answer: A

To configure the cluster to meet the requirements of the infosec team, you need to do the following steps:

Log in to Prism Central and go to Network > Syslog Servers > Configure Syslog Server. Enter Corp_syslog as the Server Name, 34.69.43.123 as the IP Address, and 514 as the Port. Select TCP as the Transport Protocol and enable RELP

(Reliable Logging Protocol). This will create a syslog server with the highest reliability possible. Click Edit against Data Sources and select Cluster B as the cluster. Select API Requests and Replication as the data sources and set the log level

to CRITICAL for both of them. This will enable audit logs for API requests and replication capabilities for the top 4 severity levels (EMERGENCY, ALERT, CRITICAL, and ERROR) and push them to the syslog server. Click Save.

Repeat step 2 for any other clusters that you want to configure with the same requirements.

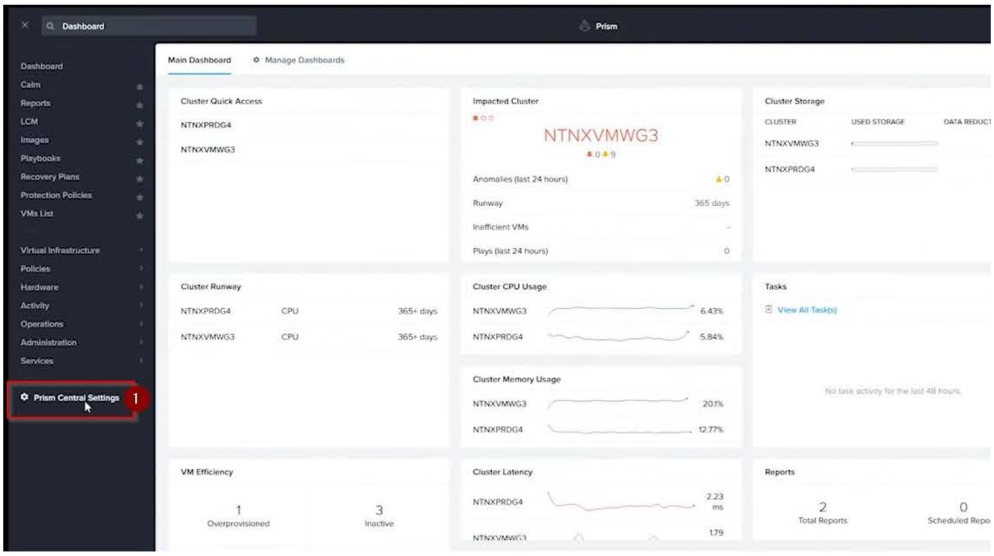

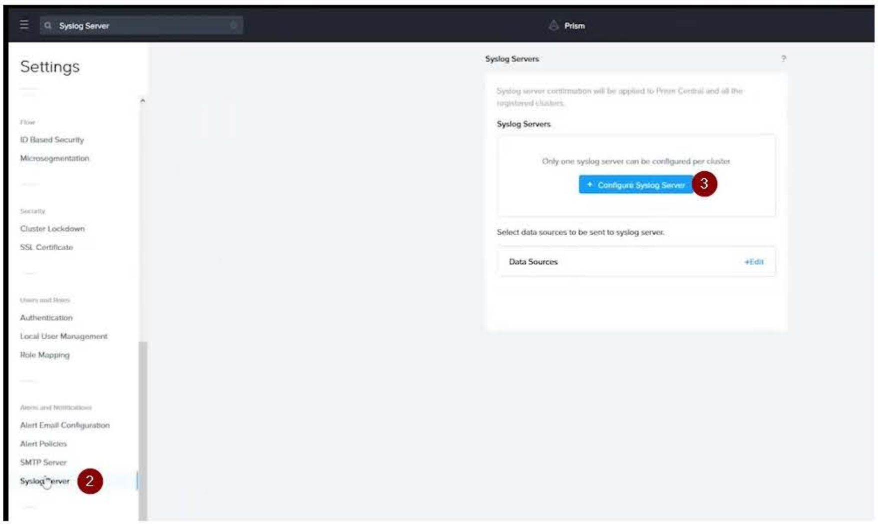

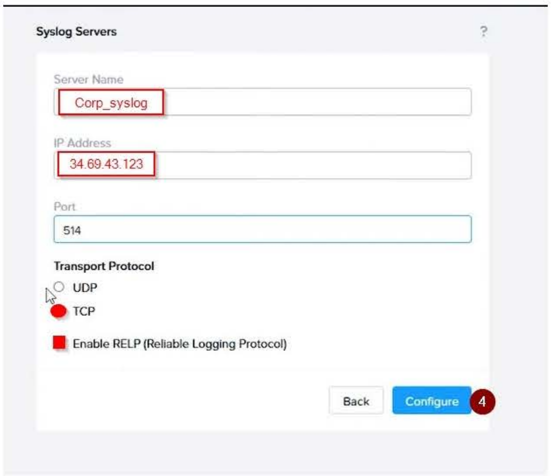
Syslog Servers

Server Name
Corp_syslog

IP Address
34.69.43.123

Port
514

Transport Protocol
○ UDP
● TCP

■ Enable RELP (Reliable Logging Protocol)

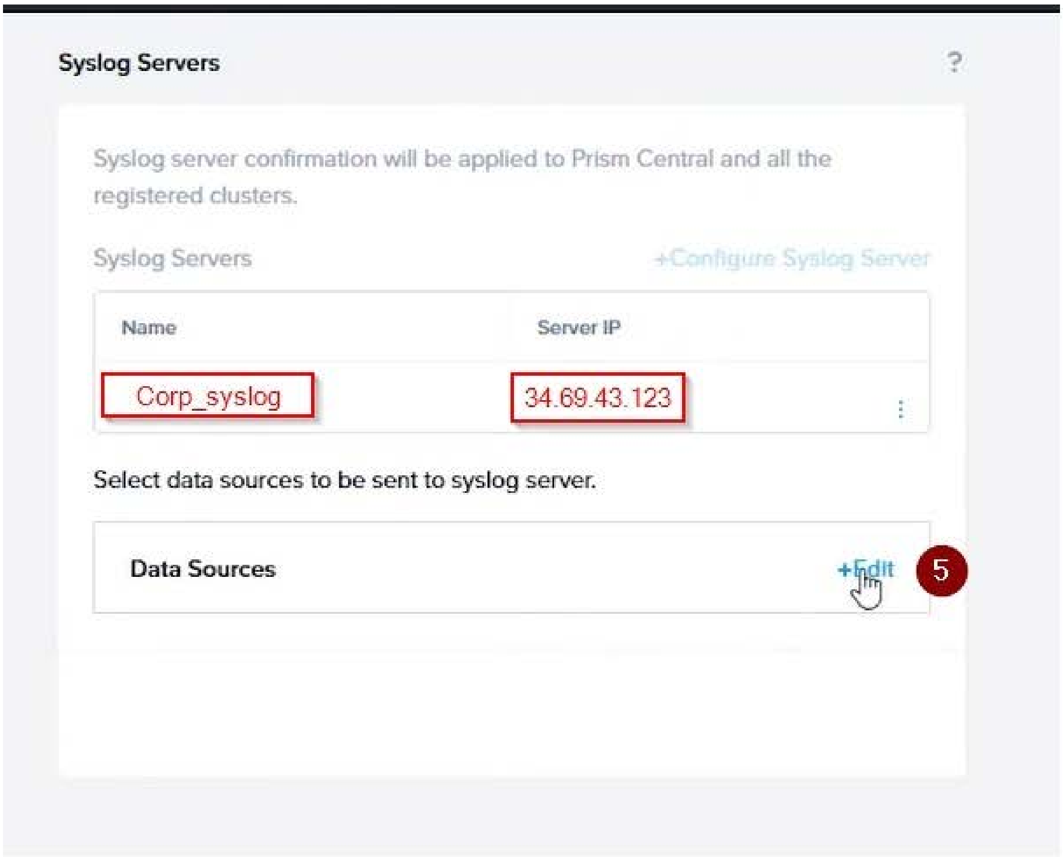

Syslog Servers ?

Syslog server confirmation will be applied to Prism Central and all the registered clusters.

Syslog Servers

Name	Server IP
Corp_syslog	34.69.43.123

Select data sources to be sent to syslog server.

Data Sources

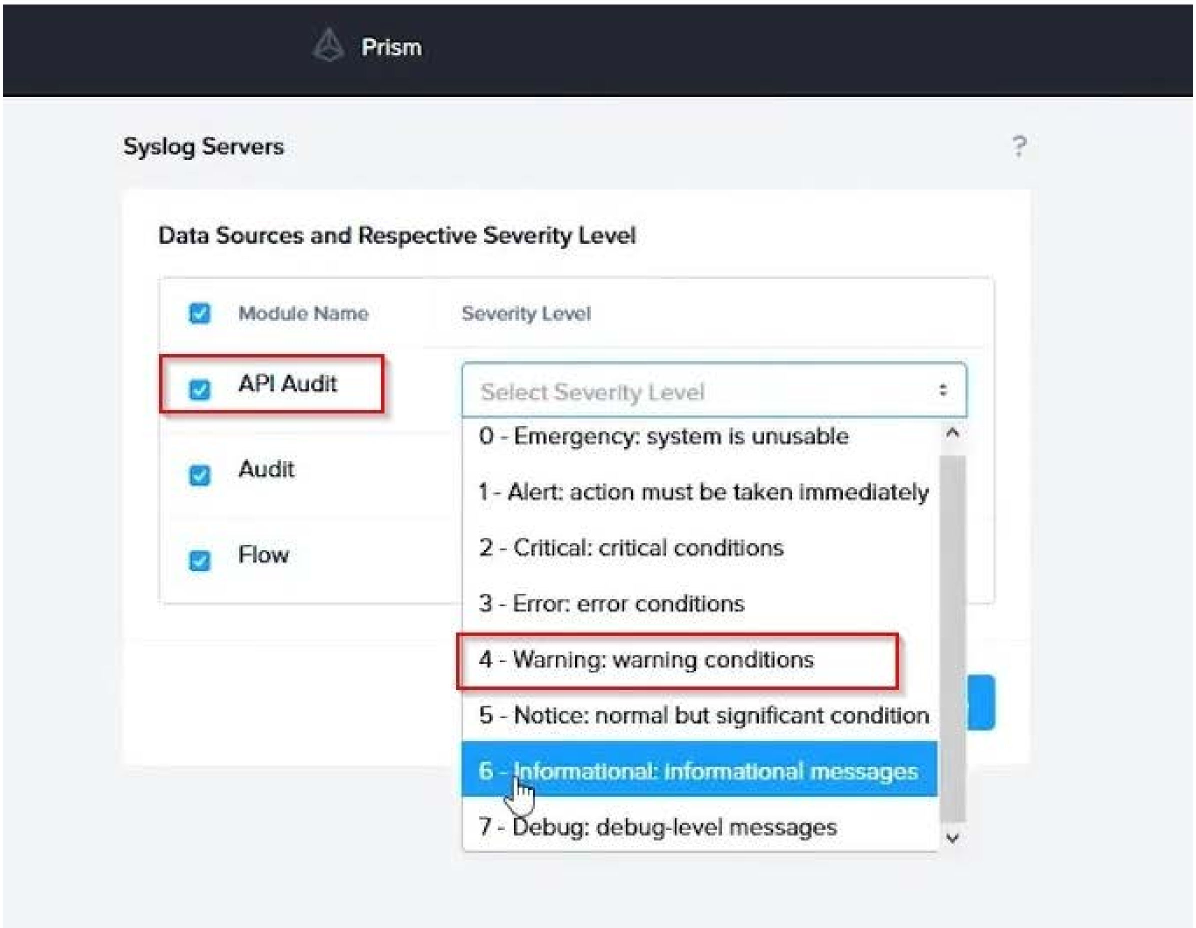

Syslog Servers ?

Data Sources and Respective Severity Level

☑ Module Name	Severity Level
☑ API Audit	Select Severity Level
☑ Audit	
☑ Flow	

0 - Emergency: system is unusable
1 - Alert: action must be taken immediately
2 - Critical: critical conditions
3 - Error: error conditions
4 - Warning: warning conditions
5 - Notice: normal but significant condition
6 - Informational: informational messages
7 - Debug: debug-level messages

To configure the Nutanix clusters to enable audit logs for API Requests and replication capabilities, and push them to the syslog system with the highest reliability possible, you can follow these steps:

Log in to the Nutanix Prism web console using your administrator credentials. Navigate to the "Settings" section or the configuration settings interface within Prism. Locate the "Syslog Configuration" or "Logging" option and click on it.

Configure the syslog settings as follows:

Syslog Name: Enter "Corp_syslog" as the name for the syslog configuration. Syslog IP: Set the IP address to "34.69.43.123", which is the IP address of the syslog system.

Port: Set the port to "514", which is the default port for syslog. Enable the option for highest reliability or persistent logging, if available. This ensures that logs are sent reliably and not lost in case of network interruptions.

Save the syslog configuration.

Enable Audit Logs for API Requests:

In the Nutanix Prism web console, navigate to the "Cluster" section or the cluster management interface.

Select the desired cluster where you want to enable audit logs. Locate the "Audit Configuration" or "Security Configuration" option and click on it. Look for the settings related to audit logs and API requests. Enable the audit logging feature and

select the top 4 severity levels to be logged.

Save the audit configuration.

Enable Audit Logs for Replication Capabilities:

In the Nutanix Prism web console, navigate to the "Cluster" section or the cluster management interface.

Select the desired cluster where you want to enable audit logs. Locate the "Audit Configuration" or "Security Configuration" option and click on it. Look for the settings related to audit logs and replication capabilities. Enable the audit logging

feature and select the top 4 severity levels to be logged.

Save the audit configuration.

After completing these steps, the Nutanix clusters will be configured to enable audit logs for API Requests and replication capabilities. The logs will be sent to the specified syslog system with the highest reliability possible.

```
ncli

 rsyslog-config set-status enable=false

 rsyslog-config add-server name=Corp_Syslog ip-address=34.69.43.123 port=514 network-protocol=tdp relp-enabled=false

 rsyslog-config add-module server-name= Corp_Syslog module-name=APLOS level=INFO

 rsyslog-config add-module server-name= Corp_Syslog module-name=CEREBRO level=INFO

 rsyslog-config set-status enable=true
```

https://portal.nutanix.com/page/documents/kbs/details?targetId=kA00e0000009CEECA2