

DUMPS ARENA

Nutanix Certified Professional (NCP) 5.10 Exam

Nutanix NCP-5.10

Version Demo

Total Demo Questions: 19

Total Premium Questions: 194

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Nutanix Enterprise Cloud Platform	20
Topic 2, Cluster Configuration	8
Topic 3, Storage Configuration	33
Topic 4, Network Configuration	27
Topic 5, VM Configuration	31
Topic 6, Cluster Management	51
Topic 7, Data Protection	15
Topic 8, Troubleshooting	9
Total	194

QUESTION NO: 1

Which is true for PD Activate and Migrate options? (Choose two.)

- A. Activate is initiated from the source site
- B. Activate is initiated from the destination site
- C. Migrate is initiated from the destination site
- D. Migrate is initiated from the source site

ANSWER: B D

Explanation:

In Nutanix asynchronous replication, activation is a disaster-recovery operation performed at the destination, or remote, site. When the source site is unavailable or its workloads must be recovered remotely, an administrator activates the replicated protection domain on the destination cluster. Activation makes the replicated entities available there so the recovery site can run the protected workloads. This workflow is designed for an unplanned failover scenario and is therefore initiated where the recovered workloads are to be brought online.

Migration is a planned relocation operation initiated from the source site. Before moving workloads, the source cluster can synchronize the latest changes to the destination, helping ensure that the destination has an up-to-date replica. The migration workflow then transfers ownership and makes the protection domain active at the destination in a controlled manner. Initiating migration at the source is important because the current active owner coordinates the final replication and planned handoff. Nutanix documents activation and migration as distinct protection-domain recovery workflows; see the [Prism Web Console Guide: activating a protection domain](#) and the [Prism Web Console Guide: migrating a protection domain](#).

QUESTION NO: 2

An administrator increases the cluster RF to 3. The containers are not modified.

What will the new values in the data resiliency dashboard be for FAILURES TOLERABLE for the Zookeeper and Extent Groups components?

- A. Zookeeper = 1 and Extent Group = 1
- B. Zookeeper = 2 and Extent Group = 2
- C. Zookeeper = 2 and Extent Group = 1
- D. Zookeeper = 1 and Extent Group = 2

ANSWER: C

Explanation:

Zookeeper = 2 and Extent Group = 1 is correct. Changing the cluster replication factor to RF3 increases resilience for cluster-level metadata services, including the ZooKeeper component. With three ZooKeeper replicas distributed across eligible nodes, the service can tolerate the loss of two nodes while retaining a surviving replica and maintaining the configured level of protection shown by the Data Resiliency dashboard.

Extent groups are governed by the replication factor configured for the individual storage container rather than solely by the cluster-wide RF setting. Because the question specifies that the containers are not modified, their existing RF2 configuration remains in effect. An RF2 extent group has two replicas and therefore shows one tolerable failure. Consequently, after the cluster is changed to RF3 but container settings remain unchanged, the dashboard reflects two failures tolerable for ZooKeeper and one failure tolerable for extent groups. Nutanix documents that resiliency reporting evaluates the protection configuration of individual components, including container data replication and cluster services. See the Nutanix community discussion on [data resilience](#) and the [Prism Data Resiliency documentation](#).

QUESTION NO: 3 - (HOTSPOT)

HOTSPOT

Async DR is configured between two sites. A network outage occurs at the primary site.

Which steps must the administrator perform to bring the VMs back into service at the backup site?

Item instructions: For each procedure, indicate the order in which that procedure must take place to meet the item requirements. Not all procedures are valid. Identify any invalid procedures using the drop-down option.

Hot Area:

Procedure	Step
Log into Prism Element at the backup site	▼ Step 1 Step 2 Step 3 Step 4 Invalid Step
Reboot VMs	▼ Step 1 Step 2 Step 3 Step 4 Invalid Step
Go to the Async DR tab	▼ Step 1 Step 2 Step 3 Step 4 Invalid Step
Log into Prism Element at the primary Site	▼ Step 1 Step 2 Step 3 Step 4 Invalid Step
Select the Protection Domain and click Activate	▼ Step 1 Step 2 Step 3 Step 4 Invalid Step
Power on VMs	▼ Step 1 Step 2 Step 3 Step 4 Invalid Step

ANSWER:

Procedure

Log into Prism Element at the backup site

Step

Step 1
Step 2
Step 3
Step 4
Invalid Step

Reboot VMs

Step 1
Step 2
Step 3
Step 4
Invalid Step

Go to the Async DR tab

Step 1
Step 2
Step 3
Step 4
Invalid Step

Log into Prism Element at the primary Site

Step 1
Step 2
Step 3
Step 4
Invalid Step

Select the Protection Domain and click Activate

Step 1
Step 2
Step 3
Step 4
Invalid Step

Power on VMs

Step 1
Step 2
Step 3
Step 4
Invalid Step

Explanation:

When the primary site is unavailable, recovery must be initiated from Prism Element on the backup cluster because that cluster contains the replicated copy of the protected data and is the site that will assume responsibility for running the workloads. The administrator first signs in to Prism Element at the backup site. From there, Async DR provides the recovery workflow and visibility into the protection domains replicated from the unavailable site.

Next, the administrator opens the Async DR tab and selects the relevant protection domain. Activating that protection domain at the backup site is the key disaster-recovery action: it makes the replicated data available for use by the recovery cluster and transitions the protected workload into a state where it can be started at that site. This sequence preserves the intended operational order: access the recovery cluster, locate the replicated protection domain through Async DR, and activate it before attempting to run its virtual machines.

After the protection domain has been activated, the virtual machines are powered on at the backup site so their services can resume. Powering them on only after activation ensures the VMs use the recovered protection-domain data and are running from the designated recovery location. This is the standard operational flow for an unplanned Async DR recovery: Step 1 is logging in to Prism Element at the backup site, Step 2 is navigating to Async DR, Step 3 is selecting the protection domain and clicking Activate, and Step 4 is powering on the VMs. Nutanix documents protection-domain recovery and activation in the [Prism Element Web Console Guide](#).

QUESTION NO: 4

An administrator wants to create a trunked interface on a VM on AOS 5.10.x.

Which two steps should the administrator take first to achieve this? (Choose two.)

- A. SSH to CVM
- B. Update VM dialog
- C. Use acli
- D. Log in over PE web UI

ANSWER: A C

Explanation:

For AOS 5.10.x, configuring VLAN trunking on an AHV virtual machine begins by connecting to a Controller VM through SSH and using the acli command-line interface. The Controller VM provides the supported administrative environment from which the AHV and VM network configuration commands can be run. After establishing the SSH session, the administrator enters acli and uses the VM NIC commands to configure the required trunked VLAN networks on the appropriate virtual NIC. This approach permits the guest operating system to receive tagged traffic for multiple VLANs through one virtual interface, while the physical and virtual networking configuration handles those VLAN tags appropriately. The VM NIC must already be associated with the relevant network configuration, and the administrator identifies the VM and NIC when applying the trunking settings. Nutanix documents acli as the command-line interface for managing AOS entities and VM configuration; the AOS command reference and Prism documentation are available through the [Nutanix AOS 5.10 Command Reference](#) and [Nutanix Prism Web Console Guide](#).

QUESTION NO: 5

An administrator needs to take a single node down for maintenance.

Which two steps should the administrator take? (Choose two.)

- A. Shut down or live migrate all VMs off of the node
- B. Make sure Auto Evacuation is enabled
- C. Verify the data resiliency status of the cluster
- D. Reclaim the licenses to the Nutanix Support Portal

ANSWER: A C

Explanation:

Before placing a node into maintenance, the administrator should shut down or live migrate all VMs off of the node. Live migration is the preferred nondisruptive approach when the workloads and cluster have sufficient resources, because it moves running workloads to other eligible hosts before the target node is serviced. Shutting down workloads is appropriate when migration is not possible or when the maintenance procedure requires it. This preparation ensures guest workloads are not left running on a host that will be powered off or otherwise made unavailable.

The administrator should also verify the data resiliency status of the cluster. Nutanix distributes and protects data across nodes according to the configured replication factor and cluster health. Confirming that the cluster is fully resilient before removing a node verifies that it can tolerate the planned node outage while maintaining the required availability of data and services. If resiliency is degraded, maintenance should be deferred until the condition is resolved or an approved recovery plan is in place. Nutanix documentation and product guides are available through the [Nutanix Documentation portal](#) and the [Nutanix Support Portal documentation library](#).

QUESTION NO: 6

An administrator wants Prism Element to send email notifications when cluster alerts are generated.

What must be configured first?

- A. An SMTP server
- B. A Data Services IP address
- C. An Active Directory trust
- D. A remote site

ANSWER: A

Explanation:

An **SMTP server** must be configured first. Prism Element uses the configured SMTP server to deliver alert notifications to the email recipients defined in the alert configuration. Without an SMTP relay, Prism can generate and display alerts locally but cannot send them through email.

The administrator should configure a reachable SMTP server, appropriate port and security settings, and sender information according to the organization's mail policy. After the SMTP configuration is validated, alert email recipients and notification preferences can be configured. See the [Nutanix Documentation portal](#) for Prism Element alert and email notification configuration guidance.

QUESTION NO: 7

Where can an administrator change a CVM password?

- A. KMS Server Terminal
- B. CVM setting in Prism Element
- C. CVM setting in Prism Central
- D. Prism CVM VM Console

ANSWER: B

Explanation:

CVM setting in Prism Element is correct because Prism Element is the cluster-local management interface used to administer CVM credentials. An administrator authenticated with the required privileges can use the password-management controls in Prism Element to change the CVM password. This action updates the password for the Controller VM account across the Nutanix cluster through the supported management workflow, rather than requiring a manual change from an individual VM console.

Using the Prism Element interface is important because Controller VMs are distributed cluster services and their credentials must remain consistently managed across the environment. Nutanix documents password changes as an administrative operation performed through the Nutanix management interfaces and procedures, which helps preserve the expected cluster configuration and avoids unsupported credential inconsistencies. The administrator should ensure that the correct cluster is selected in Prism Element and follow the documented password-change procedure for the Nutanix software version in use. For further guidance, see Nutanix's [password-modification guidance](#) and the [Prism documentation portal](#).

QUESTION NO: 8

A Nutanix administrator needs to replace a self-signed SSL certificate with a CA generated certificate in Prism.

Which three files are required to install the CA generated certificate in Prism? (Choose three.)

- A. Private Server Key (private_server.key)
- B. Public Certificate (PublicCert.cer)
- C. Certificate Signing Request (server.csr)
- D. Public Server Key (public_server.key)
- E. CA Certificate or Chain (ca_chain_certs.cer)

ANSWER: A B E

Explanation:

Installing a CA-signed SSL certificate in Prism requires the matching private key, the server's public certificate, and the issuing CA certificate chain. **Private Server Key (private_server.key)** is required because it is the private half of the key pair used to create the certificate signing request and must cryptographically match the installed server certificate. **Public Certificate (PublicCert.cer)** is the certificate issued by the CA for the Prism endpoint; it contains the endpoint identity and public key that browsers and API clients validate during the TLS handshake. **CA Certificate or Chain (ca_chain_certs.cer)** supplies the intermediate and, where required, root CA certificates needed for clients to construct a trusted certification path from the Prism server certificate to a trusted authority. Providing the full chain is important because clients may not have the intermediate CA certificate locally. Before installation, the administrator should ensure the certificate subject/SAN entries match the Prism access name and that the private key corresponds to the issued certificate. Nutanix documents the SSL certificate replacement workflow and required certificate components in the [Prism Central Guide](#). General Nutanix product documentation is also available through the [Nutanix Support Portal documentation library](#).

QUESTION NO: 9

Prism Central will be installed manually on an AHV cluster.

Which three disk images must be downloaded from the portal for the Prism Central VM? (Choose three.)

- A. Var
- B. tmp
- C. Boot
- D. Home
- E. Data

ANSWER: C D E

Explanation:

A manual Prism Central deployment on AHV requires the Prism Central VM's three supplied disk-image components: **Boot**, **Home**, and **Data**. The Boot image provides the virtual appliance's operating-system boot disk. The Home image provides the filesystem and application area used by Prism Central software. The Data image provides the persistent storage area required for Prism Central operational data, including its internal services and retained management information. Together, these images form the expected virtual disk layout for the Prism Central appliance when it is created manually rather than deployed through an automated workflow.

Download the images that correspond to the intended Prism Central release from the Nutanix Support Portal, then upload them to the target AHV cluster's image service before creating the VM and attaching the disks according to the manual-deployment procedure. Using all three matched images from the same release is important because the appliance's software and disk layout are delivered as a coordinated set. Nutanix documents Prism Central deployment and VM requirements in the [Prism Central Guide](#); the applicable product documentation can also be accessed through the [Nutanix Support Portal](#).

QUESTION NO: 10

Cluster experiences a CVM failure. However, VMs on the nodes continue to produce I/O.

Which Distributed Storage Fabric can be attributed to this behavior?

- A. Information Lifecycle Management
- B. Data Path Redundancy
- C. VM Snapshotting and Cloning
- D. Scalable Cassandra Database

ANSWER: B

Explanation:

Data Path Redundancy is the Distributed Storage Fabric capability that allows VM I/O to continue when a local Controller VM (CVM) fails. In Nutanix, the CVM normally provides the local storage-controller path for the workloads running on its host. To avoid making that CVM a single point of failure, the Nutanix data path includes alternate access routes to storage services running on other healthy CVMs in the cluster. When the local CVM becomes unavailable, the host and its VMs can use a remote CVM to continue servicing reads and writes, preserving application access while the failed CVM is recovered or restarted. This behavior is integral to Nutanix's distributed architecture: storage services are distributed across cluster nodes rather than being dependent on one centralized controller. The result is continued data availability and VM I/O during an individual CVM outage, assuming the cluster retains sufficient healthy nodes and network connectivity. Nutanix describes its distributed, scale-out architecture and availability model in its [Cloud Infrastructure overview](#) and provides additional product documentation through the [Nutanix Support Portal documentation library](#).

QUESTION NO: 11

Hypervisors communicate with DSF using the industry-standard protocols.

Which industry-standard protocols do hypervisors use to communicate with DSF? (Choose three.)

- A. SAN
- B. NAS
- C. NFS
- D. SMB3
- E. iSCSI

ANSWER: C D E

Explanation:

NFS, SMB3, and iSCSI are the industry-standard storage protocols through which hypervisors can consume storage presented by Nutanix Distributed Storage Fabric (DSF). DSF aggregates the local storage resources across the Nutanix cluster and presents storage to supported hypervisors in the protocol format expected by their storage stack. NFS is used for datastore access in VMware ESXi environments. SMB3 provides the file-based storage connectivity used by Microsoft Hyper-V environments, including Cluster Shared Volumes hosted on SMB shares. iSCSI provides block-storage connectivity for hypervisor configurations that require block devices or volumes. This protocol-based presentation allows the hypervisor to access Nutanix-managed storage without requiring proprietary storage protocols or external SAN infrastructure. The underlying data services—such as replication, resiliency, capacity optimization, and distributed placement—remain managed by the Nutanix platform while the hypervisor uses its native, standards-based interface. Nutanix documentation describes DSF as the distributed storage layer that presents storage services to compute workloads and hypervisors. See the [Nutanix Documentation Portal](#) and the [Nutanix Software-Defined Storage overview](#).

QUESTION NO: 12

An administrator has set up a local web server accessible to the Nutanix cluster.

Which two steps are required to set up LCM for an environment without Internet access? (Choose two.)

- A. Download the `lcm_disconnected_version.tar.gz` LCM bundle from the support portal.
- B. Edit LCM Advanced Settings in Prism Element and enter the address of the web server.
- C. Edit LCM Advanced Settings in Prism Central and enter the address of the web server.
- D. Download the `lcm_dark_site_version.tar.gz` LCM bundle from the support portal.

ANSWER: B D

Explanation:

For a disconnected, or dark-site, environment, Lifecycle Manager must obtain its release metadata and upgrade binaries from an internally reachable web server rather than directly from Nutanix-hosted Internet locations. The required package is the `lcm_dark_site_version.tar.gz` bundle downloaded from the Nutanix Support Portal. This bundle contains the LCM dark-site content that must be made available through the administrator's local web server.

After publishing the bundle's contents on that server, the cluster must be configured with the server location in the LCM Advanced Settings in Prism Element. Prism Element is the cluster-level management interface where LCM on the target cluster is directed to use the dark-site repository URL. Once configured, inventory and upgrade operations can retrieve the required LCM information from the internal repository, allowing lifecycle operations to proceed without outbound Internet access. Nutanix documents the dark-site workflow, including obtaining the dark-site bundle, hosting it internally, and configuring the dark-site URL, in the [Nutanix LCM Guide](#). Additional Nutanix support guidance is available in [Knowledge Base article 7090](#).

QUESTION NO: 13

Which algorithm do snapshots and clones leverage to maximize efficiency and effectiveness?

- A. Redirect-On-Write
- B. Copy-on-Write
- C. Split-mirror
- D. Continuous Data Protection

ANSWER: A

Explanation:

Redirect-On-Write is the algorithm Nutanix AOS uses for snapshots and clones. AOS maintains snapshot metadata at a fine granularity and, when a write targets data protected by a snapshot, redirects the new write to a new location rather than overwriting the existing data block. The pre-write version remains available through the snapshot's metadata, preserving the point-in-time view without requiring a full independent copy of the virtual machine or its disks.

This approach supports storage-efficient snapshots and rapid clones because the initial clone can share unchanged data with its source. It also enables deep snapshot chains while minimizing the performance and capacity impact that would result from copying all data for each snapshot or clone. Nutanix describes its Distributed Storage Fabric as using a fine-grained redirect-on-write algorithm for snapshots, emphasizing efficient cloning and snapshot-chain operation. See the [Nutanix Distributed Storage Fabric overview](#) and Nutanix's [data protection overview](#) for additional context on AOS snapshot-based protection capabilities.

QUESTION NO: 14

Which data savings technique utilizes stripes and parity calculation in a Nutanix cluster?

- A. Compression
- B. Parity strip
- C. Erasure coding
- D. Deduplication

ANSWER: C

Explanation:

Erasure coding is the Nutanix data-savings technique that divides data into multiple stripes and calculates parity information across those stripes. Rather than retaining complete replica copies of data, the cluster distributes data and parity fragments so that it can reconstruct protected data if a disk or node failure occurs. This approach reduces the storage capacity required for resiliency while continuing to provide fault tolerance. In Nutanix environments, erasure coding is particularly valuable for data that is no longer actively being written, because the system can efficiently encode eligible extents and reduce the space consumed compared with replication. Nutanix manages the placement, parity calculation, and reconstruction processes transparently as part of its distributed storage fabric, so administrators receive the capacity benefit without needing to manually create traditional RAID groups. The technique is therefore specifically characterized by stripes plus parity calculations, which distinguishes it as the applicable data-reduction mechanism. Nutanix describes erasure coding as a capacity optimization method that protects data using encoded data and parity strips across the cluster. See the [Nutanix Bible](#) and Nutanix's [Prism Web Console Guide: Data Reduction](#) for further details.

QUESTION NO: 15

A cluster is configured with Replication Factor 3 (RF3).

Which two statements are true for this configuration? (Choose two.)

- A. Three replicas of protected data are maintained
- B. More raw capacity is required than with RF2
- C. Only one replica of each data extent is maintained
- D. Erasure coding is disabled for all containers

ANSWER: A B

Explanation:

RF3 maintains three replicas of protected data across different nodes in the cluster. The additional replica provides greater fault tolerance than RF2 because the cluster can continue to access data when more replica locations become unavailable, subject to the specific failure scenario and cluster health.

RF3 also requires more raw storage capacity than RF2. Because three copies of protected data are maintained rather than two, administrators must account for the additional capacity overhead when sizing the cluster. RF3 is selected when the required resiliency level justifies the increased capacity consumption. See the [Nutanix Bible](#) and the [Nutanix Documentation portal](#) for data resiliency information.

QUESTION NO: 16

Cluster RF3 currently has 5 storage containers with storage pool: 15 TB. You add 2 containers 'A' and 'B' with 1 TB reserved capacity each.

What's now the raw capacity of all seven containers?

- A. original 5 containers: 4 TB each; containers A and B: 3 TB each
- B. original 5 containers: 3 TB each; containers A and B: 4 TB each

- C. original 5 containers: 3 TB each; containers A and B: 3 TB each
- D. original 5 containers: 4 TB each; containers A and B: 4 TB each
- E. 15 TB; adding storage containers and reserving capacity does not change the storage pool's raw capacity.

ANSWER: E

Explanation:

15 TB is the raw capacity, assuming the stated 15 TB is the storage pool's raw-capacity value. In Nutanix, a storage container is a logical partition of a storage pool, not an independently sized physical-capacity unit. Multiple containers can share the same pool, so creating two additional containers does not add disks or increase the pool's raw physical capacity.

The 1 TB reserved-capacity settings protect capacity for the applicable containers by reducing capacity that can be consumed by other workloads in the shared pool. A reservation is therefore an allocation or availability control; it is not new physical storage. Likewise, RF3 determines the amount of physical storage needed to maintain three copies of protected data and affects the pool's effective usable capacity. It does not make container reservations create additional raw capacity. Therefore, there is no valid basis for assigning separate 3 TB or 4 TB raw-capacity values to each container. The cluster continues to have the storage pool's existing 15 TB of raw capacity until physical storage is added to the pool. Nutanix describes storage containers as logical entities within a storage pool and explains capacity usage and resiliency concepts in its [Prism storage-container documentation](#) and [Nutanix storage architecture reference](#).

QUESTION NO: 17

An administrator needs to create a microsegmentation policy that allows only web servers to communicate with application servers.

Which two actions are required in Prism Central? (Choose two.)

- A. Create categories and assign them to the VMs
- B. Create a security policy
- C. Create a Volume Group for each VM category
- D. Configure an IPAM address pool for each category

ANSWER: A B

Explanation:

Creating categories and assigning them to the appropriate VMs is required because Flow uses categories to identify the groups of workloads to which a policy applies. For example, the administrator can assign a Web category to web server VMs and an Application category to application server VMs.

Creating a security policy is also required. The security policy defines the allowed or denied traffic between the categorized workloads, including the applicable protocol and port information. Categories identify the policy endpoints, while the policy enforces the communication rule. See the [Nutanix Flow Network Security overview](#) and the [Nutanix Documentation portal](#) for Flow configuration information.

QUESTION NO: 18

In Nutanix cluster, which feature ensures VMs can be migrated and restarted on another host in case of failure?

- A. High Availability
- B. Protection Domain
- C. Host Affinity Rules

D. Availability zone

ANSWER: A

Explanation:

High Availability is correct because it is the Nutanix cluster capability that detects an unavailable host and automatically recovers the affected user virtual machines on surviving hosts. When a host failure occurs, the cluster identifies the virtual machines that were running there and restarts them elsewhere, provided sufficient resources are available. Nutanix High Availability uses cluster resource awareness and can reserve failover capacity so that the cluster retains enough compute resources to accommodate workloads from a failed host. This minimizes service interruption without requiring an administrator to manually identify, place, and power on the affected workloads. The virtual machines are restarted rather than transparently continuing execution, so applications may experience a brief interruption while the guest operating systems boot and services recover. High Availability is configured and monitored at the cluster level in Prism and is a core resilience feature for maintaining workload availability during an unexpected node or host outage. Nutanix documentation describes High Availability as the mechanism for automatically restarting virtual machines after an AHV host failure and for managing the capacity needed for that recovery. See the [Nutanix Prism Web Console Guide: High Availability](#) and [Nutanix AHV Administration Guide: High Availability](#).

QUESTION NO: 19

What disk formats are supported by the image service?

- A. vhd, raw, vmdk, vdi, iso, qcow2
- B. qcow2, vdi, raw, udf, iso
- C. raw, ova, vmdk, qcow2, vdi
- D. vmdk, qcow2, vdi, ovf, raw

ANSWER: A

Explanation:

The supported disk-format set is **vhd, raw, vmdk, vdi, iso, qcow2**. Nutanix Image Service provides a centralized image repository in Prism, allowing administrators to upload installation media and virtual-machine disk images once and then use them for VM creation and deployment across the cluster. ISO is used for bootable installer or operating-system media, while RAW, VHD, VMDK, VDI, and QCOW2 represent supported virtual disk-image formats originating from common virtualization platforms and tools. When an image is uploaded, Prism identifies its image type and makes it available to supported VM workflows; disk images can be selected as VM disks, and ISO images can be mounted as CD-ROM media. This broad format support is intended to simplify migrations and workload onboarding by accepting source images from multiple environments without requiring every image to be recreated first. For the NCP 5.10 context, the listed formats match the Image Configuration documentation for Prism's Image Service. See Nutanix's [Prism Web Console Guide: Image Configuration](#) and the [Prism Web Console Guide: Image Service](#).