DUMPSBOSS.

Nutanix Certified ProfessionalCloud Native (-6.10)

Nutanix NCP-CN

Version Demo

Total Demo Questions: 13

Total Premium Questions: 135

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Kubernetes Cluster Deployment	63
Topic 2, Kubernetes Cluster Management	50
Topic 3, Kubernetes Application Deployment and Management	14
Topic 4, Kubernetes Networking	1
Topic 5, Kubernetes Storage	3
Topic 6, Kubernetes Troubleshooting	3
Topic 7, Mix Questions	1
Total	135

QUESTION NO: 1

A Platform Engineer is preparing a custom operating-system image for repeated NKP node deployments. The engineer needs to understand the default image-build workflow before publishing the image to the required infrastructure repository.

Which two statements are correct regarding this workflow? (Choose two.)

A. Packer is used to build a reusable operating-system image for NKP nodes.

B. Goss is used to validate that the operating-system image can run NKP.

C. Terraform is used to validate that the operating-system image can run NKP.

D. Ansible is used to publish the image as a Cluster API provider.

ANSWER: A B

Explanation:

Packer is used by the image-building workflow to produce a reusable machine image with the operating-system configuration and NKP node prerequisites required by Cluster API-based provisioning.

Goss validates that the resulting operating-system image meets the expected NKP node requirements. It performs configuration validation rather than creating infrastructure. Terraform and Ansible can be used in broader automation designs, but they are not the default NKP image-build and validation tools described in this workflow.

QUESTION NO: 2

A Platform Engineer is attaching an existing EKS cluster to NKP. The management cluster can reach the EKS API server, but the attachment process fails while attempting to create management resources in the target cluster. The API response indicates that the supplied identity is forbidden from creating cluster-scoped resources.

What should the engineer do?

A. Create a service account with the required cluster-admin permissions and use its credentials for attachment.

B. Create an additional Amazon load balancer for the EKS API server.

C. Replace all EKS worker nodes with an NKP custom image.

D. Configure an NKP infrastructure provider to provision a new EKS cluster.

ANSWER: A

Explanation:

The engineer should create or use a dedicated service account with the cluster-level permissions required for the NKP attachment workflow, then provide credentials for that identity to NKP. Network reachability is already confirmed, so the decisive issue is authorization against the target Kubernetes API.

NKP must be able to create and reconcile the cluster resources needed to register and manage the attached environment. Adding a load balancer or modifying the EKS worker-node image does not resolve an RBAC denial from the Kubernetes API server.

QUESTION NO: 3

A service provider is onboarding a new customer that requires access to two NKP workload clusters. The customer must be isolated from other customers, and its development group must have a defined CPU, memory, and storage budget across the assigned clusters.

Which two actions should the Platform Engineer take? (Choose two.)

A. Create a dedicated NKP workspace and assign the customer's clusters to it.

B. Configure quotas for the customer workspace.

C. Assign the customer group a Global Administrator role.

D. Create a separate infrastructure provider for every customer application.

E. Grant every customer user cluster-admin access on all NKP clusters.

ANSWER: A B

Explanation:

Creating a dedicated NKP workspace and assigning the customer's clusters to it establishes the required tenant boundary. The workspace provides a common scope for the customer's cluster resources, access assignments, and operational configuration.

Configuring workspace quotas establishes the resource budget for the customer's workloads across the workspace scope. The engineer can then assign appropriate workspace RBAC roles to the customer group, but assigning RBAC alone does not limit CPU, memory, or storage consumption. Creating separate infrastructure providers per application is not required to establish customer isolation.

QUESTION NO: 4

Refer to Exhibit:

Using an NKP Ultimate license, a Platform Engineer has created a new workspace and needs to create a new Kubernetes cluster within this workspace. However, the engineer discovers that the Create Cluster option is grayed out, as shown in the exhibit. How should the engineer resolve this issue?

A. Create the cluster only using YAML and not the GUI.

B. Attach existing clusters instead of creating a new cluster.

C. Create an Infrastructure provider for the workspace.

D. Ensure NKP is upgraded to a minimum version of 2.12.

ANSWER: C

Explanation:

Create an Infrastructure Provider for the workspace. is the required resolution because NKP needs an infrastructure provider before it can offer managed-cluster creation in a workspace. An infrastructure provider supplies the target environment and the credentials, configuration, and supported provisioning workflow that NKP uses to deploy the

Kubernetes cluster. Depending on the environment, this provider can represent Nutanix infrastructure or a supported public-cloud or virtualization target.

A newly created workspace does not automatically have a usable provisioning target. Consequently, NKP cannot determine where cluster control-plane and worker resources should be created, so the user interface disables cluster creation. The Platform Engineer should configure the appropriate Infrastructure Provider in NKP and make it available to the intended workspace. After the provider is configured and associated appropriately, NKP can validate that the workspace has an eligible infrastructure target, and the Create Cluster action becomes available.

This behavior is independent of the Ultimate license entitlement: the license enables relevant platform capabilities, while an Infrastructure Provider establishes the concrete location and access configuration required to provision a cluster. Nutanix describes NKP as providing centralized Kubernetes lifecycle management across supported infrastructures in its Nutanix Kubernetes Platform overview. Product documentation and version-specific installation guides are available through the Nutanix Support Portal documentation library.

QUESTION NO: 5

A Platform Engineer manages an NKP v2.12.x environment and is using NKP Image Builder (NIB) to create a custom image. Which two distributions are available for use by the engineer for this task? (Choose two.)

A. Ubuntu

B. Fedora

C. Rocky Linux

D. CentOS

ANSWER: A C

Explanation:

Ubuntu and **Rocky Linux** are the supported distribution families for creating custom NKP node images with NKP Image Builder in the stated NKP release stream. NIB provides a supported, repeatable process for producing machine images that include the operating-system configuration, Kubernetes prerequisites, and Nutanix-specific components required by NKP cluster nodes. Using one of these supported bases ensures that the resulting images align with the operating-system packages, kernel behavior, container runtime requirements, and cloud-provider integration that NKP validates.

Ubuntu is commonly used as an NKP node operating system because its long-term-support releases provide a stable package base and security-maintenance lifecycle. Rocky Linux provides an enterprise Linux-compatible base suitable for organizations that standardize on that ecosystem. Selecting either distribution lets the Platform Engineer create a consistent golden image for control-plane and worker nodes, reducing configuration drift and making cluster deployment and lifecycle operations more predictable. Consult the NKP version-specific documentation for the exact supported release of each distribution and image-building procedure: Nutanix Product Documentation and Nutanix Documentation and Support Resources.

QUESTION NO: 6

A Platform Engineer has deployed NKP and wants to utilize its OOB data storage feature.

What should the engineer enable to support backups within the NKP environment?

A. MinIO

B. Rook Ceph

C. Volumes iSCSI

D. Objects S3

ANSWER: B

Explanation:

Rook Ceph is the appropriate feature to enable because it provides NKP's integrated, cloud-native storage foundation. Rook is a Kubernetes storage orchestrator that deploys and manages Ceph within the cluster. Ceph supplies resilient distributed block, file, and object storage, allowing stateful Kubernetes applications to retain data independently of individual Pods or nodes.

For backup use cases, Ceph's object-storage capabilities can provide an S3-compatible target for backup data, while its persistent storage services support workloads that require durable volumes. This makes Rook Ceph suitable for an NKP environment in which the platform engineer wants an out-of-the-box storage service that can underpin persistent application data and backup operations. Enabling it establishes the storage layer that Kubernetes-native backup tooling and workloads can consume through standard storage and object-storage interfaces.

Rook documents Ceph as a Kubernetes-native way to run and manage Ceph storage, including block, shared filesystem, and object-storage services. Nutanix also documents NKP's data-services capabilities and supported integrations. See the Rook documentation overview and the Nutanix Kubernetes Platform documentation.

QUESTION NO: 7

A Platform Engineer is deploying a new NKP cluster that has internet connectivity. Now, a Cloud Administrator and Security Administrator are discussing the security of communications between the NKP Kubernetes cluster and the container registry. The engineer proposes to have an on-prem private registry.

What is the most significant reason that the engineer should create a private registry instead of configuring a secure connection between the NKP cluster and Github (SaaS)?

A. Private registry license is included with NKP.

B. NKP requires specific registry versions.

C. NKP cannot connect to public clouds.

D. Private registry provides security and privacy.

ANSWER: D

Explanation:

Private registry provides security and privacy. An on-premises private registry gives the organization direct control over where container images are stored, how they are accessed, and which images may be deployed to the NKP cluster. Images can remain inside the organization's security boundary, limiting exposure of proprietary application artifacts, embedded configuration, and image metadata to an external SaaS service. It also enables enterprise controls such as role-based access, private-network access paths, audit logging, vulnerability scanning, image signing or approval workflows, and retention policies that align with internal security and compliance requirements.

For Kubernetes workloads, a registry is a core part of the software supply chain because cluster nodes retrieve images from it during workload deployment. Using a privately managed registry supports a controlled, trusted image source and allows security teams to govern the content that becomes available to workloads. Kubernetes supports authenticated access to private registries through image pull secrets, enabling the cluster to consume images without making them publicly accessible. See the Kubernetes guidance for pulling images from a private registry and Nutanix's overview of the Nutanix Kubernetes Platform.

QUESTION NO: 8

Two application teams share an NKP workload cluster. The platform team must prevent one team from consuming an unlimited amount of CPU and memory, and it must prevent that team's Pods from communicating directly with the other team's database Pods.

Which two Kubernetes controls should the platform team implement? (Choose two.)

A. ResourceQuota

B. NetworkPolicy

C. ClusterRoleBinding

D. LimitRange only

E. IngressClass

ANSWER: A B

Explanation:

A ResourceQuota limits the aggregate resource consumption within a namespace. It can restrict the total CPU and memory requests or limits available to one team, preventing that namespace from consuming capacity needed by other workloads.

A NetworkPolicy controls permitted traffic between Pods. The platform team can use it to deny direct traffic by default and allow only the required application paths, such as communication through an approved service. A LimitRange controls defaults and per-container boundaries but does not define an aggregate namespace budget, while RBAC controls API permissions rather than Pod-to-Pod network traffic.

QUESTION NO: 9

A Platform Engineer needs to do an air-gapped installation of NKP. This environment needs to run without Internet access and be fully operational, including updates. Docker has been installed, and the NKP bundle exists on a bastion host. What is the first command that the engineer must run to begin the process?

A. nkp push bundle --bundle

B. docker load -i konvoy-bootstrap-image-v2.12.0.tar

C. tar -xzvf nkp-air-gapped-bundle_v2.12.0_linux_amd64.tar.gz

D. nkp create cluster nutanix

ANSWER: C

Explanation:

`tar -xzvf nkp-air-gapped-bundle_v2.12.0_linux_amd64.tar.gz` is the required first command because the downloaded NKP air-gapped bundle is delivered as a compressed tar archive. Before the bastion host can provide any NKP artifacts to an isolated environment, the archive must be unpacked so its directory structure and installation resources are available locally. Extraction exposes the NKP command-line tooling, image archives, manifests, and scripts that support the subsequent stages of the air-gapped workflow.

This ordering is important for a fully self-contained deployment and later lifecycle operations. Once the bundle is extracted, the engineer can use the artifacts it contains to load required images into Docker or a private registry, configure local registry access, and invoke the appropriate NKP cluster-installation workflow. The extracted bundle is therefore the local source of truth for software dependencies that would otherwise be retrieved from Internet-hosted locations. The `-x`, `-z`, `-v`, and `-f` flags respectively extract, process gzip compression, display progress, and specify the archive filename. Nutanix documents air-gapped installation procedures in the NKP Air-Gapped Installations documentation; the archive operation follows the standard GNU tar documentation.

QUESTION NO: 10

A Platform Engineer is reviewing provisioning choices for several NKP cluster deployments. One deployment will use virtual machines created on Nutanix infrastructure, while another will use physical servers that were prepared before NKP installation.

Which two statements correctly describe the applicable Cluster API provisioning methods? (Choose two.)

A. CAPX can provision NKP virtual machines on Nutanix infrastructure.

B. CAPPP uses an inventory to identify servers that have already been provisioned.

C. CAPA is required when NKP provisions virtual machines on Nutanix AHV.

D. CAPV is required when NKP provisions physical servers from an inventory file.

ANSWER: A B

Explanation:

CAPX is the Cluster API provider used for Nutanix infrastructure and supports provisioning NKP machines on Nutanix. It is appropriate when the platform creates the virtual machines that will become cluster nodes.

CAPPP is used for pre-provisioned servers and relies on an inventory of the existing machines that will become NKP nodes. It does not dynamically create infrastructure. CAPA is the AWS provider, while CAPV is associated with vSphere rather than Nutanix AHV.

QUESTION NO: 11

At a national defense company, protecting sensitive data is their top priority. With the increase in cyber-attacks, they have decided to implement an air-gapped Kubernetes environment to manage their critical applications, ensuring that no information could leak to the outside. The Kubernetes environment has three clusters deployed for their applications with centralized management. What type of licensing is required to enable this environment?

A. NKP Starter

B. NKP Ultimate

C. NKP Pro

D. NKP UI

ANSWER: B

Explanation:

NKP Ultimate is required because the design combines two enterprise capabilities: operating Kubernetes in a disconnected, air-gapped environment and centrally managing a fleet of multiple clusters. Air-gapped Kubernetes requires that the software artifacts, container images, and required dependencies be made available locally rather than retrieved from public internet registries during deployment or lifecycle operations. This supports the organization's isolation and data-security requirements while allowing its Kubernetes platform to remain deployable and maintainable.

The scenario also describes three application clusters with centralized management. Nutanix Kubernetes Platform fleet-management functionality provides a management-plane approach for attaching, observing, and administering multiple Kubernetes clusters consistently. The Ultimate licensing tier is intended for these advanced deployment and management requirements, including air-gapped environments and multi-cluster fleet operations. Therefore, licensing all relevant NKP components with NKP Ultimate enables the defense company to operate its isolated clusters while retaining centralized operational control. Nutanix describes NKP capabilities and deployment resources in its Nutanix Kubernetes Platform documentation; its Nutanix Kubernetes Platform product page also outlines the platform's enterprise Kubernetes management focus.

QUESTION NO: 12

An NKP Pro environment has a licensing capacity of 100 vCPUs. The current cluster contains three control-plane nodes with 4 vCPUs each and six worker nodes with 8 vCPUs each. The Platform Engineer must expand capacity without exceeding the licensed amount.

Which two expansion plans are compliant? (Choose two.)

A. Expand the existing worker node pool by five nodes with 8 vCPUs each.

B. Add two worker nodes with 8 vCPUs each and a new six-node worker pool with 4 vCPUs per node.

C. Expand the existing worker node pool by six nodes with 8 vCPUs each.

D. Add three control-plane nodes with 4 vCPUs each and four worker nodes with 8 vCPUs each.

ANSWER: A B

Explanation:

The current cluster consumes 60 vCPUs: 12 vCPUs for the three control-plane nodes and 48 vCPUs for the six workers. Adding five 8-vCPU workers adds 40 vCPUs, reaching exactly 100 vCPUs.

Adding two 8-vCPU workers and six 4-vCPU workers also adds 40 vCPUs, reaching 100 vCPUs. Adding six 8-vCPU workers would reach 108 vCPUs. Adding three 4-vCPU control-plane nodes plus four 8-vCPU workers would reach 104 vCPUs.

QUESTION NO: 13

A Platform Engineer needs to create an NKP cluster on vSphere infrastructure, using the vSphere provisioning method. The cluster needs to have 3 worker node pools:

First node pool should consist of 6 worker nodes

Second node pool should consist of 3 worker nodes

Third node pool should consist of 3 worker nodesAdditionally, the worker nodes in the first node pool should be set to 10 CPUs, the second node pool workers should be set to 8 CPUs, and the third node pool workers should be set to 6 CPUs. What is the proper way to create the NKP cluster using the NKP CLI?

A. First, execute the nkp create cluster vsphere command, including the following parameters:--worker-replicas 6--worker-cpus 10Then, execute the nkp create nodepool vsphere command, including the following parameters:--worker-replicas 3--worker-cpus 8Then, execute the nkp create nodepool vsphere command, including the following parameters:--worker-replicas 3--worker-cpus 6

B. First, execute the nkp create cluster vsphere command, including the following parameters:--worker-replicas 6--worker-cpus 10Then, execute the nkp create nodepool vsphere command, including the following parameters:--replicas 3--cpus 6

C. When executing the nkp create cluster vsphere command, include the following parameters:--node-pools 3--worker-replicas 6,3,3--worker-cpus 10,8,6

D. First, execute the nkp create cluster vsphere command, including the following parameters:--worker-replicas 6--worker-cpus 10Then, execute the nkp create nodepools vsphere command, including the following parameters:--replicas 3,3--cpus 8,6

ANSWER: A

Explanation:

First, execute the `nkp create cluster vsphere` command with `--worker-replicas 6` and `--worker-cpus 10`, then create two additional vSphere node pools with their respective replica and CPU values. This is the proper workflow because initial cluster creation defines the cluster's first worker pool, while each additional worker pool is created as a separate node-pool resource. Supplying six replicas and ten vCPUs during cluster creation produces the initial pool with the required capacity. The subsequent `nkp create nodepool vsphere` invocations independently define a pool of three workers with eight CPUs and another pool of three workers with six CPUs. In practice, each command also requires the environment-specific vSphere and cluster parameters, such as the target cluster name and vCenter configuration, but the replica and CPU arguments shown establish the requested per-pool sizing.

This separate-resource approach is important because node pools can have different machine characteristics and scaling values. NKP's vSphere CLI reference documents the cluster-creation command and its worker machine configuration

arguments, and the node-pool CLI reference documents creation of additional vSphere worker pools. See the Nutanix Kubernetes Platform CLI reference and the Nutanix Kubernetes Platform product documentation portal.