

NCP-CN真題材料，NCP-CN考古題更新



從Google Drive中免費下載最新的NewDumps NCP-CN PDF版考試題庫：<https://drive.google.com/open?id=169HSBmMpyR40u84PnVGxTckULHvoKX8K>

在短短幾年內，Nutanix NCP-CN 認證考試已經成為比較有影響力電腦能力認證考試。然而如何簡單順利地通過Nutanix NCP-CN認證考試？我們的NewDumps在任何時間下都可以幫您快速解決這個問題。我們在NewDumps中為您提供了可以成功通過NCP-CN認證考試的培訓工具。NCP-CN認證考試培訓工具的內容是由IT行業專家帶來的最新的考試研究材料組成

Nutanix NCP-CN Exam Syllabus Topics:

Section	Weight	Objectives
Prepare the Environment for an NKP Deployment	22%	- Create a bootstrap cluster - Gather deployment configuration information - Prepare a bastion host - Build machine images or prepare target nodes - Seed a private container registry - Determine license tiers (Community, Standard, Enterprise)
Perform Day 2 Operations	30%	- Configure authentication and authorization - Set up logging and monitoring systems - Monitor cluster performance and health - Perform lifecycle management: upgrades, updates, maintenance - Configure cluster autoscaling - Implement backup and recovery procedures
Manage Building an NKP Cluster	28%	- Apply and manage cluster licenses - Customize and deploy NKP clusters - Customize and deploy Kommander management platform
Conduct NKP Fleet Management	20%	- Configure workspaces for cluster organization - Deploy workload clusters into workspaces - Attach, detach, or remove clusters from workspaces - Configure projects for resource segmentation - Deploy and manage platform applications

>> NCP-CN真題材料 <<

NCP-CN考古題更新 & NCP-CN認證資料

如果你的預算是有限的，但需要完整的價值包，不如嘗試一下我們NewDumps Nutanix的NCP-CN考試培訓資料。我們NewDumps可以為你的IT認證保駕護航，是目前網路上最受歡迎的最可行的培訓資料網站，NCP-CN考試是你職業生涯中的一個里程碑，在這種競爭激烈的世界裏，它比以往任何時候都顯得比較重要，我們保證讓你一次輕鬆的通過考試，也讓你以後的工作及日常工作變得有滋有味。還可以幫你挖掘到許多新的途徑和機會。這實在對著起這個價錢，它所創造的價值遠遠大於這個金錢。

最新的 Nutanix Certified Professional (NCP) NCP-CN 免費考試真題 (Q105-Q110):

問題 #105

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. Ensure NKP is upgraded to a minimum version of 2.12.
- **B. Create an Infrastructure provider for the workspace.**
- C. Attach existing clusters instead of creating a new cluster.
- D. Create the cluster only using YAML and not the GUI.

答案: B

問題 #106

Refer to the exhibit.



NODEPOOL	DESIRED	READY	KUBERNETES VERSION
md-0	4	4	v1.29.9

A current Nutanix cluster is hosting an NKP cluster and a production Kubernetes cluster named production.

Infrastructure administrators recently added three nodes with NVidia GPUs for a new AI initiative, and now a Platform Engineer has been asked to add three workers with the GPUs to the production Kubernetes cluster.

Which first step should the engineer take to achieve this task?

- **A. Create a GPU-compatible OS Image with:**
text
CollapseWrap
Copy
`nkp create image nutanix --gpu \`
`--gpu-name=${GPU_NAME} \`
`--cluster=${NUTANIX_CLUSTER_NAME} \`
`--endpoint=${NUTANIX_PC_ENDPOINT} \`
`--subnet=${NUTANIX_SUBNET} ubuntu-22.04`
- B. Add the GPU Operator to the new workers:
text
CollapseWrap
Copy
`cat <<EOF | kubectl apply -f -`
`kind: Installation`
`apps:`
`nvidia-gpu-operator:`
`enabled: true`
`values: |`
`toolkit:`
`nodepool: md-0`
`EOF`
- C. Configure Multi-Instance GPU (MIG):
text
CollapseWrap
Copy
`kubectl label nodes ${NODE} nvidia.com/mig.config=all-1g.6gb -overwrite`
- D. Create a nodepool of workers with GPU:
text
CollapseWrap
Copy
`nkp create nodepool nutanix -c $CLUSTER_NAME \`

```
--vm-image $NAME_OF_GPU_OS_IMAGE_CREATED_BY_NKP_CLI \
-o yaml --subnets $NUTANIX_SUBNET_NAME \
--prism-element-cluster $NUTANIX_PRISM_ELEMENT_CLUSTER_NAME \
$NODEPOOL_NAME
```

答案： A

解題說明：

Comprehensive and Detailed Explanation From Exact Extract of Nutanix Kubernetes Platform Administration (NKPA) Course:

The NKPA course outlines the process of adding GPU-enabled worker nodes to an existing NKP-managed Kubernetes cluster, such as the production cluster in this scenario. The first step in this process is to ensure that a GPU-compatible OS image is available for the new worker nodes, as GPU support requires specific drivers and configurations (e.g., NVIDIA drivers) that are not included in standard OS images.

The correct first step is to create a GPU-compatible OS image using the command:

text

CollapseWrap

Copy

```
nkp create image nutanix --gpu \
--gpu-name=${GPU_NAME} \
--cluster=${NUTANIX_CLUSTER_NAME} \
--endpoint=${NUTANIX_PC_ENDPOINT} \
--subnet=${NUTANIX_SUBNET} ubuntu-22.04
```

(Option D). This command uses the NKP CLI to create a machine image based on Ubuntu 22.04, tailored for Nutanix AHV infrastructure (nutanix) with GPU support enabled (--gpu). The --gpu-name flag specifies the GPU type (e.g., NVIDIA GPU model), and other parameters define the Nutanix cluster, Prism Central endpoint, and subnet for image creation. The resulting image includes the necessary NVIDIA drivers and dependencies, making it suitable for GPU-enabled worker nodes. The Nutanix Cloud Native (NCP-CN) 6.10 Study Guide states: "Before adding GPU-enabled workers to an NKP cluster on Nutanix, the first step is to create a GPU-compatible OS image using nkp create image nutanix --gpu, ensuring the image includes the required GPU drivers for the target infrastructure." This image is then used in subsequent steps (like Option A) to create a node pool with GPU-enabled workers.

Without this image, the node pool creation in Option A would fail due to the lack of a suitable --vm-image.

Incorrect Options:

- * A. Create a nodepool of workers with GPU: This step requires a pre-existing GPU-compatible OS image (specified via --vm-image). Since the question does not indicate that such an image already exists, creating the image (Option D) must happen first.
- * B. Add the GPU Operator to the new workers: The NVIDIA GPU Operator can be installed to manage GPU resources, but this step occurs after the GPU-enabled workers are added to the cluster, not as the first step.
- * C. Configure Multi-Instance GPU (MIG): MIG configuration is an advanced GPU setup step that occurs after the workers are added and the GPU Operator is installed, not as the first step.

:

Nutanix Kubernetes Platform Administration (NKPA) Course, Section on GPU Support in NKP Clusters.

Nutanix Cloud Native (NCP-CN) 6.10 Study Guide, Chapter on Managing NKP Clusters.

Nutanix Cloud Bible, NutanixKubernetesPlatform Section: <https://www.nutanixbible.com> NVIDIA GPU Operator Documentation:

<https://docs.nvidia.com/datacenter/cloud-native/gpu-operator> Before adding GPU-enabled worker nodes to the cluster, a GPU-compatible OS image must be created. This ensures that the required drivers and kernel modules are present for GPU workloads.

This step is critical as it ensures the OS image is properly prepared for GPU operations before the nodes are integrated into the cluster.

References: Nutanix Kubernetes Platform Administration Guide - GPU Node Preparation and Image Management

問題 #107

To keep an NKP cluster and applications healthy and drive productivity forward, a Platform Engineer needs to stay informed of all events occurring within the cluster. What component of kube-prometheus-stack will help the engineer to stay informed of these events in NKP?

- A. node-exporter
- B. prometheus-operator
- C. service monitors
- D. alertmanager

答案： D

解題說明：

The kube-prometheus-stack is a key component of NKP's monitoring stack, providing tools for metrics collection, visualization, and alerting. The NKPA course explains that Alertmanager, a component of the kube-prometheus-stack, is responsible for handling alerts generated from Prometheus metrics. It aggregates, deduplicates, and routes notifications to the appropriate channels (e.g., email, Slack, PagerDuty), ensuring that the Platform Engineer stays informed of critical events and anomalies in the NKP cluster, such as node failures, resource exhaustion, or application errors.

The Nutanix Cloud Native (NCP-CN) 6.10 Study Guide states: "Alertmanager in the kube-prometheus-stack processes alerts from Prometheus, enabling administrators to stay informed of cluster events through configured notification channels." By configuring Alertmanager with appropriate routing rules and receivers, the engineer can receive real-time notifications about cluster events, driving proactive management and productivity.

Incorrect Options:

- * A. prometheus-operator: The operator manages Prometheus and related resources but does not directly handle event notifications.
- * B. service monitors: Service monitors define how Prometheus scrapes metrics, not how events are communicated.
- * D. node-exporter: Node-exporter collects node-level metrics, not event notifications.

:

Nutanix Kubernetes Platform Administration (NKPA) Course, Section on Monitoring and Alerting.

Nutanix Cloud Native (NCP-CN) 6.10 Study Guide, Chapter on Day 2 Operations.

Nutanix Cloud Bible, NutanixKubernetesPlatform Section: <https://www.nutanixbible.com> Prometheus Documentation: <https://prometheus.io/docs/alerting/alertmanager>

問題 #108

A technology company has decided to migrate its infrastructure to NKP to improve the scalability and management of its applications. After a successful initial implementation, the operations team faces a new challenge of validating the HelmReleases to ensure that all applications are running correctly and avoid problems in production. Which command should the company execute to know the right status of their HelmReleases?

- A. `kubectl get helmreleases -n ${PROJECT_NAMESPACE}`
- B. `kubectl apply -f fluent-bit-overrides.yaml`
- C. `kubectl get namespaces`
- D. `kubectl edit helmreleases -n ${PROJECT_NAMESPACE}`

答案： A

解題說明：

NKP uses Helm and Flux for application deployment, where applications are managed as HelmReleases, a custom resource that defines Helm chart deployments. The NKPA course explains that to validate the status of HelmReleases, administrators can use the kubectl command to query these resources. The correct command is `kubectl get helmreleases -n ${PROJECT_NAMESPACE}`, which lists all HelmReleases in the specified project namespace, showing their status (e.g., Ready, Failed) and other details like chart version and reconciliation status.

The Nutanix Cloud Native (NCP-CN) 6.10 Study Guide states: "To check the status of HelmReleases in NKP, use `kubectl get helmreleases -n <namespace>` to view the current state of application deployments managed by Flux." This command helps the operations team verify that applications are running correctly and identify any issues in production.

Incorrect Options:

- * A. `kubectl get namespaces`: This lists namespaces, not HelmReleases or their status.
- * C. `kubectl edit helmreleases -n ${PROJECT_NAMESPACE}`: This edits HelmRelease resources, not displays their status.
- * D. `kubectl apply -f fluent-bit-overrides.yaml`: This applies a configuration for Fluent Bit, unrelated to HelmRelease status.

:

Nutanix Kubernetes Platform Administration (NKPA) Course, Section on Application Management.

Nutanix Cloud Native (NCP-CN) 6.10 Study Guide, Chapter on Day 2 Operations.

Nutanix Cloud Bible, NutanixKubernetesPlatform Section: <https://www.nutanixbible.com> Flux Documentation: <https://fluxcd.io/docs/components/helm>

問題 #109

A Platform Engineer is a member of an IT team that provides Kubernetes clusters for three groups within a company named Fin Group, Inc.:

Fin VD

Fin Insurance

Fin TravelThe engineer created workspaces for each group. Fin Group Inc. has its own Active Directory implementation, while each

- A. Create a role named admin-tenant-X, where X is the name of the group, and assign that role to manage the corresponding workspace.
- B. Create a role binding and assign it to manage the corresponding workspace.
- C. Configure the global Active Directory and assign a workspace admin user to each group.
- **D. Configure a dedicated identity provider for each group to access their own workspace.**

• • • • •

[illegible]

BONUS!!! 免費下載NewDumps NCP-CN考試題庫的完整版: <https://drive.google.com/open?id=169HSBmMpyR40u84PnVGxTckULHvoKX8K>