



- Expert Verified, Online, **Free**.



CERTIFICATION TEST

- CertificationTest.net - Cheap & Quality Resources With Best Support

An administrator receives reports about a Nutanix environment. The investigation finds the following;

- ⇒ VMs are experiencing very high latency
- ⇒ Each node is equipped with a single SSD, utilized at 95%
- ⇒ Each node is equipped with three HDDs, utilized at 40%

Why are the guest VMs experiencing high latency?

- A. CVMs are overwhelmed by disk balancing operations.
- B. All VM write operations are going to HDD.
- C. All VM read operations are coming from HDD.
- D. VMs are unable to perform write operations

Suggested Answer: A

Community vote distribution

A (50%)

B (50%)

🗳️ 👤 **ed838d7** 10 months, 2 weeks ago

This exam is obsolete. It is not the NCM-MCI 6.5 break fix exam with 15 questions. It is an outdated exam. I would guess either 5.15 or 6.0, not 6.5.
upvoted 2 times

🗳️ 👤 **e4af987** 10 months, 2 weeks ago

I think A - which comes first? the Reprovisioning of where the data rests or the because now the Reads are coming from HDD? I think A.
upvoted 1 times

🗳️ 👤 **pharaoh2000** 1 year, 11 months ago

Selected Answer: B

B. On the Nutanix Cluster once a disk (SSD-PCIe, SSD-SATA or DAS-SATA storage tier) reaches 95% full Stargate will not write to that disk
<https://portal.nutanix.com/page/documents/kbs/details?targetId=kA0600000008hfcCAA>

D. The question does not indicate the the data is cold, only cold data is read from HDD.
upvoted 1 times

🗳️ 👤 **EhCloudGrrrU** 2 years, 1 month ago

Selected Answer: B

Answer is 'B' - if the SSDs are 95% full, they can long take on new writes, therefore writes have to go to HDD, which is much slower.
upvoted 2 times

🗳️ 👤 **nabillose** 2 years, 1 month ago

Correct answer is C
upvoted 1 times

🗳️ 👤 **AlexNut** 3 years, 2 months ago

Selected Answer: A

After using of 75% ssd starts overwhelmed tearing operation, correct is A
upvoted 3 times

🗳️ 👤 **BillBaits** 3 years, 5 months ago

Correct
upvoted 1 times

In a Nutanix cluster, a Protection Domain contains 50 entities that are replicated to a remote Single Node Replication Target. The current schedule configuration is as follows:

- ☞ Repeat every 6 hours
- ☞ Local Retention Policy 1
- ☞ Remote Retention Policy 8
- ☞ Starting time 12 am

At 8 am on Monday, the administrator discovers that a protected VM is corrupted. The latest good state was Sunday 2 pm. The administrator needs to maintain current protection.

Which strategy should the administrator use to meet these requirements?

- A. From the Remote site, activate the Protection Domain, then re-protect the entity.
- B. From the Remote site, restore the VM from the local snapshot by selecting the correct snapshot.
- C. From the local site, retrieve the correct remote snapshot, then restore the VM locally.
- D. From the local site, restore the VM from the local snapshot by selecting the correct snapshot.

Suggested Answer: C

🗨️ **nabillose** 1 year, 1 month ago

correct answer is C, because we have retention policy 1 in the local site.
upvoted 1 times

🗨️ **FOG** 1 year, 4 months ago

I think is C because you have Local Retention Policy 1 and it's corrupted
upvoted 1 times

An administrator has a custom backup application that requires a 2TB disk and runs in Windows. Throughput is considerably lower than expected.

The application was installed on a VM with the following configuration:

- * Four vCPUs with one core/vCPU
- * 4GB of Memory
- * One 50GB vDisk for the Windows installation
- * One 2TB vDisk for the application

What is the recommended configuration change to improve throughput?

- A. Increase the number of cores per vCPU
- B. Increase the vCPUs assigned to the VM
- C. Span the 2TB disk across four vDisks
- D. Add 4GB of memory to the VM

Suggested Answer: C

Currently there are no comments in this discussion, be the first to comment!

An administrator is responsible for the following Nutanix Enterprise Cloud environment:

A central datacenter with a 20-node cluster with 1.5PB of storage

▪

⇒ Five remote sites each with a 4-node cluster with 200TB storage

The remote sites are connected to the datacenter via 1GB links with an average latency of 6 ms RTT.

What is the minimum RPO the administrator can achieve for this environment?

- A. 0 minutes
- B. 15 minutes
- C. 1 hour
- D. 6 hours

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ **mogmog** 10 months, 1 week ago

0 min RPO as sync rep is not possible due to 6ms RTT. Requires <5ms RTT

Central DC: 75TB per node

Remote DC: 50TB per node.

Based on 6.5 document. Near sync and Async 1h is support. Lowest is 15mins

https://portal.nutanix.com/page/documents/details?targetId=Leap-Xi-Leap-Admin-Guide-v2022_6:ecd-ecdr-resource-requirements-r.html
upvoted 1 times

🗳️ **nabillose** 1 year, 1 month ago

response D

upvoted 1 times

🗳️ **EhCloudGrrrU** 1 year, 9 months ago

Selected Answer: D

On nodes that have storage greater than 64 TB (80 TB for all-flash nodes), the minimum supported snapshot frequency is six (6) hours. For such nodes, this limitation has the following consequences:

The minimum supported RPO for asynchronous replication is also 6 hours.

NearSync is not supported.

Any cluster that includes such a node is also subject to this limitation.

https://portal.nutanix.com/page/documents/details?targetId=Web-Console-Guide-Prism-v5_15:wc-dr-near-sync-requirements-limitations-r.html
upvoted 2 times

🗳️ **SG1ATLANTIS** 1 year, 12 months ago

Selected Answer: D

Definitely response D

upvoted 3 times

🗳️ **Azimuth** 2 years, 2 months ago

Selected Answer: D

On nodes that have storage greater than 40 TB (48 TB for all-flash nodes), the minimum supported snapshot frequency is 6 hours. For such nodes, this limitation has the following consequences:

The minimum supported RPO for asynchronous replication is also 6 hours.

NearSync is not supported.

Any cluster that includes such a node is also subject to this limitation.

This limitation does not apply to hypervisor snapshots

upvoted 4 times

An administrator of a Nutanix Enterprise Cloud has microsegmentation enabled and a firewall VM installed. The security team notifies the administrator that one of the VMs hosted in the environment has been exhibiting suspect network activity. The administrator wants to isolate the VM from the production network, but must still be able to access it to perform diagnostics. What should the administrator do to meet the requirement?

- A. Disable the vNIC on the affected VM
- B. Quarantine the VM using the Forensic Method
- C. Create a firewall rule that blocks VM traffic but permits diagnostic access
- D. Create a security policy with a service chain directing that VMs traffic to the firewall

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **mogmog** 1 year, 4 months ago

Selected Answer: B

Forensic. Isolates the VM from all traffic except traffic from categories specified in the built-in quarantine policy. The allowed categories contain forensic tools that enable you to perform forensics on the VM.

https://portal.nutanix.com/page/documents/details?targetId=Nutanix-Flow-Guide-v5_17:mul-quarantine-rule-configuration-c.html
upvoted 2 times

🗨️ 👤 **Cromo19** 10 months, 3 weeks ago

Agreed, answer is B.
upvoted 1 times

🗨️ 👤 **Treasure** 1 year, 7 months ago

The "Forensic" option allows the limitation of inbound/outbound traffic to that which is specified from the "Add Forensic Tools" tab.
upvoted 2 times

🗨️ 👤 **nabillose** 1 year, 7 months ago

Quarantining the VM using the Forensic Method would also isolate the VM, but it would not allow for diagnostic access.
Option: C
upvoted 1 times

A customer recently set up Async Replication between Site A and Site B. The customer wants to conduct a planned failover and clicks Activate on Site B. The customer then runs the following command on Site A: `ncli pd deactivate_and__destroy_vms name=<protection_domain_name>`. What does this do to the customer environment?

- A. VMs get deleted from Site B and the protection domain is now Active.
- B. VMs are powered off on Site A and must be manually powered on at Site B.
- C. VMs get deleted from Site A and the protection domain is no longer active.
- D. Customer must then manually power off VMs at Site A and power them on at Site B.

Suggested Answer: C

```
<ncli> pd deactivate_and_destroy_vms name=DataProtection
Use this command only when you are preparing to do a failback from an unplanned
failover. Executing this command will delete the VMs from this site and might ca
use data loss when used incorrectly. Are you sure (y/N)? Error: Mode change for
protection domain DataProtection failed with error: Canno
t deactivate protection domain DataProtection because 1 VM(s) (POC_DEMO) are Pow
ered ON
```

Community vote distribution



prashant13485 4 months ago

Selected Answer: C

Why i feel that the question says what will happen if we run the command and not what to do prior to run the command? In that case C would be the right answer because the image clearly states if you run VMs will get deleted.

upvoted 1 times

Cromo19 1 year, 10 months ago

Selected Answer: D

With the screenshot from the command output, it is clear that if the VM is not powered off before running it, it fails.

upvoted 1 times

pharoah2000 2 years, 5 months ago

D is the best answer here

1. power off VMs on Site A, without powering off VM, CLI command won't work
2. run the CLI on Site A > Site A become Secondary
3. power on VMs at Site B

upvoted 2 times

rmaio 10 months, 1 week ago

"3. power on VMs at Site B" is actually wrong as the question states they've been already powered on (" The customer wants to conduct a planned failover and clicks Activate on Site B"). C is the right one.

upvoted 1 times

njmadeja 2 years, 7 months ago

D. is correct

upvoted 1 times

nabillose 2 years, 7 months ago

Correct :C

upvoted 2 times

An administrator needs to create and start five new VMs for a Data Analytics Project (OLAP). The VM Profile is as follows:

☞ vCPU:4

☞ vRAM: 64 GB

☞ vDisk: 1.5 TB

Each of the four nodes of the Nutanix cluster has the following:

☞ 24 vCPU. 20% overall usage

☞ 192 GB RAM, 60% overall usage

☞ 2 x 1920 GB SSD

☞ 4 x 2 TB HDD

At cluster level, there is a single RF2 container that is 30% utilized and has an extent store capacity of 13.5 TB.

Which component requires administrator attention?

- A. Physical RAM, because it is not enough to power on all of the new VMs.
- B. Physical Cores, because they are not enough to power on all of the new VMs.
- C. Storage, because the capacity is not enough to create VMs.
- D. Flash Tier because it is not enough to accommodate the workloads.

Suggested Answer: A

Community vote distribution

A (100%)

🗉 👤 **Cromo19** 10 months, 3 weeks ago

Selected Answer: A

Answer is A according to elijahxc's explanation below.

upvoted 1 times

🗉 👤 **elijahxc** 1 year ago

Option A. As stated in the question, 60% of the total RAM in the Nutanix cluster is in use (based on '192 GB RAM, 60% overall usage'), which amounts to 768 GB (192 GB per node * 4 nodes).

This means that only 40% of the remaining RAM in the cluster, equivalent to 307.2 GB (768 GB * 40%), is available for the creation of the five VMs.

Since the total amount of vRAM needed for the five VMs is 320 GB (64 GB per VM * 5 VMs), it exceeds the available RAM (320 GB > 307.2 GB), making Option A the correct answer.

upvoted 2 times

🗉 👤 **nabillose** 1 year, 7 months ago

Correct : A

upvoted 2 times

An administrator manages the following two Nutanix AOS 5.15 cluster environment:

☞ Corp-cluster-01

☞ Corp-cluster-02

The VM images must be available only on Corp-cluster-01, but cannot be checked out to cluster Corp-cluster-02. The images also cannot be checked out to any other clusters that are registered with Prism Central in the future.

Which two configuration settings must the administrator choose when creating the image placement policy that satisfies the stated requirements? (Choose two.)

- A. Create an image placement policy that identifies cluster Corp-cluster-01 as the target cluster
- B. Set the policy enforcement to Soft.
- C. Set the policy enforcement to Hard.
- D. Create an image placement policy that identifies cluster Corp-cluster-02 as the target cluster.

Suggested Answer: AC

🗉  **nabillose** 1 year, 1 month ago

The two configuration setting that the administrator must be choose are : A and C
upvoted 2 times

An administrator is supporting a business critical environment and deploys metro availability to achieve a zero data loss configuration. The two clusters are connected by a 1GbE connection. A new workload is going to be deployed to this cluster. This workload requires a sustained 150MB/s of write throughput and 20MB/s of read throughput.

Which change must be made to deploy the workload successfully on this cluster?

- A. The bandwidth must be increased to support this workload.
- B. The workload must be configured to read at greater than 12.5MB/s.
- C. The replication frequency must be less than 60 minutes.
- D. Zero data loss nearsync must be used to support this workload.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

Refer to the exhibit.

User	Quota	Enforcement
All Users	No quota limit	
fileslab\user1	5 GiB	Soft
All Users in "AAPM"	10 GiB	Hard

User 1 and User 2 belong to the AAPM group.

What are two descriptions of how Files perform quota management? (Choose two.)

- A. User 2 can continue to add another 3GB file but will receive a daily notification email.
- B. User 1 can continue to add another 5GB file but will receive a daily notification email.
- C. User 1 can add an 8GB file to the share without receiving notification.
- D. User 2 can add an 8GB file without receiving a notification email.

Suggested Answer: BD

Community vote distribution

BD (100%)

🗉 👤 **adzlan** 1 year, 2 months ago

Selected Answer: BD

Hard Limit: Prevents further writes after reaching quota limit.

Soft Limit: Does not prevent writes. Sends email notifications to email recipients.

Note: You cannot set both a soft and hard limit for the same user or group.

upvoted 2 times

🗉 👤 **Yasser_Mohamed** 2 years, 1 month ago

BD

<https://portal.nutanix.com/page/documents/kbs/details?targetId=kA00e000000LMQtCAO>

upvoted 2 times

An organization is planning an upgrade to AOS 5.15 and wants to understand which cluster products and/or services are supported for physical traffic isolation.

Which Nutanix component supports its network traffic being isolated onto its own virtual network?

- A. Volumes
- B. Objects
- C. Containers
- D. Files

Suggested Answer: A

Community vote distribution

A (100%)

  **begoun** 9 months, 3 weeks ago

Selected Answer: A

https://portal.nutanix.com/page/documents/details?targetId=Nutanix-Security-Guide-v6_10:wc-network-segmentation-service-isolation-support-r.html

upvoted 1 times

  **Cromo19** 10 months, 3 weeks ago

Selected Answer: A

Answer is A

upvoted 1 times

An administrator is planning to deploy a Nutanix cluster to support a few high-performance VMs. This deployment will have the following considerations:

- ⇒ Individual VMs will likely generate network throughputs in the range of 90-100Mbps
- ⇒ Due to the configuration of the VMs, it is unlikely that a node will host more than one or two of them at a time
- ⇒ Individual VMs will communicate with only a few (one or two) remote hosts at a time
- ⇒ Multicast will not be used in the environment
- ⇒ The hosts are connected via two 1GbE network connections

How should the administrator configure the network bonds to meet this requirement while maintaining the simplest switch configuration?

- A. Configure bond0 as balance-sib
- B. Configure bond0 as active-backup
- C. Configure bond0 as balance-tcp
- D. Configure bond0 as balance-lacp

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **double_nope** Highly Voted 3 years, 6 months ago

In the answers, balance-sib should be balance-slb.
upvoted 5 times

🗳️ 👤 **Cromo19** Most Recent 1 year, 4 months ago

Selected Answer: A

Answer is A
upvoted 1 times

🗳️ 👤 **pharoah2000** 1 year, 11 months ago

Selected Answer: A

A: balance-slb is the best answer here
both physical interfaces are in use
balance-slb does not use LACP, switch configuration is simple.
balance-slb do not support multicast

B. only one physical interface is in use at any time

C. It requires LACP to be configured on the switch, not the simplest configuration

D. no such bond mode
upvoted 3 times

🗳️ 👤 **Treasure** 2 years, 1 month ago

Because:
"Individual VMs will likely generate network throughputs in the range of 90-100Mbps"
"Multicast will not be used in the environment"
"simplest switch configuration"
So the answer should be "balance-slb".
upvoted 2 times

🗳️ 👤 **Daner** 2 years, 5 months ago

Answer must be B, Multicast is against the Nutanix Best Practices with balance SLB. Active-Backup configuration is enough for this Use Case
upvoted 3 times

🗳️ 👤 **Zippy12** 1 year, 3 months ago

Multicast isn't an issue, as it's explicitly not being used per the description.

Also, it gives a likely max of 2 VMs x 100MBps at a time. 200 MBps is 1600Mbps, which is more than 1GbE adapter can provide at a time. Active-Backup is definitely not sufficient for this, so everyone else saying balance-slb is correct.

upvoted 1 times

An administrator configures authentication in Prism Central (PC) for the NTNX-Admins group. Users report that they are unable to log in. The administrator confirms that the Active Directory (AD) settings in PC are correct. Which condition could result in this behavior?

- A. A role mapping needs to be associated with the AD group.
- B. The users must be added to the Domain Admins AD group.
- C. The Local authentication type must also be selected in PC.
- D. A client chap certificate must also be added in PC.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ **Cromo19** 10 months, 3 weeks ago

Selected Answer: A

Answer is A

upvoted 1 times

🗳️ **mogmog** 1 year, 4 months ago

Selected Answer: A

just need to map ad group to a role.

upvoted 2 times

🗳️ **pharoah2000** 1 year, 5 months ago

A is correct

upvoted 1 times

🗳️ **nabillose** 1 year, 7 months ago

A is correct answer

upvoted 1 times

🗳️ **maxgiver** 2 years ago

Selected Answer: A

A is correct. It's not necessary to have domain admin ACL

upvoted 1 times

🗳️ **EhCloudGrrrU** 2 years, 3 months ago

Selected Answer: A

Role mapping must be created.

upvoted 1 times

🗳️ **Sarunpk6** 2 years, 4 months ago

Aanswer is "A"

upvoted 1 times

🗳️ **Azimuth** 2 years, 9 months ago

Selected Answer: A

Roles in Prism Central must be granted to AD groups. Best practices are to grant the least permissions possible. With this in mind, we would not add the Domain Admins group in AD to any of the roles in Prism Central. Instead, a new group would be created in AD, linked in Prism Central, and only the necessary personnel who needed access would be added.

upvoted 3 times

🗳️ **jssanche2019** 3 years, 1 month ago

A is correct. No permissions are granted to the directory users by default. To grant permissions to the directory users, you must specify roles for the users in that directory (see Configuring Role Mapping).

upvoted 4 times

An administrator has an existing Nutanix seven-node cluster running at RF2 that must be changed to have the following capabilities:

- ☞ RF3 set for the cluster
- ☞ Performance tier deduplication

What is the minimum memory needed on the CVMs to allow for these capabilities?

- A. 24GB
- B. 28GB
- C. 32GB
- D. 20GB

Suggested Answer: B

Community vote distribution

B (100%)

🗉 👤 **Cromo19** 10 months, 3 weeks ago

Selected Answer: B

Agreed with Akane_ and anistuta, the answer is B.

upvoted 1 times

🗉 👤 **Akane_** 2 years, 4 months ago

Selected Answer: B

Redundancy factor 3 has the following requirements:

- Redundancy factor 3 can be enabled at the time of cluster creation or after creation too.
- A cluster must have at least five nodes for redundancy factor 3 to be enabled.
- For guest VMs to tolerate the simultaneous failure of two nodes or drives in different blocks, the data must be stored on containers with replication factor 3.
- Controller VMs must be configured with a minimum of 28 GB(20 GB default+8 GB for the feature) of memory.

upvoted 2 times

🗉 👤 **anistuta** 2 years, 7 months ago

Selected Answer: B

https://portal.nutanix.com/page/documents/details?targetId=Web-Console-Guide-Prism-v6_0:wc-cvm-memory-configuration-c.html

upvoted 2 times


An administrator needs to deploy an application with a large amount of data connected via Nutanix Volumes. Which two actions should the administrator take when designing the Volume Group? (Choose two.)

- A. Use multiple subnets for iSCSI traffic
- B. Enable RSS (Receive Side Scaling)
- C. Enable thick provisioning on the Volume Group(s)
- D. Distribute workload across multiple virtual disks

Suggested Answer: BD

Community vote distribution

BD (100%)

 **Cromo19** 10 months, 3 weeks ago

Selected Answer: BD

The answers are B and D.

Ref: <https://portal.nutanix.com/page/documents/solutions/details?targetId=BP-2049-Nutanix-Volumes:BP-2049-Nutanix-Volumes>
upvoted 1 times

An administrator deploys Exchange on a Nutanix cluster. The administrator creates two containers with RF2 configuration:

- ⇒ One container for Exchange
- ⇒ One container for the other VMs

The Exchange VM is experiencing I/O performance issues.

Which changes to Storage Optimization should be used on the Exchange container to enhance I/O performance?

- A. No data reduction
- B. Post process Deduplication only
- C. Inline compression and EC-X
- D. Inline compression and Deduplication

Suggested Answer: C

Community vote distribution

C (100%)

🗉 👤 **Cromo19** 10 months, 3 weeks ago

Selected Answer: C

The answer is C

upvoted 1 times

🗉 👤 **pharoah2000** 1 year, 5 months ago

Selected Answer: C

Enable compression at the container level.

Enable erasure coding (EC-X).

<https://portal.nutanix.com/page/documents/solutions/details?targetId=BP-2036-Microsoft-Exchange-Server%3ABP-2036-Microsoft-Exchange-Server&a=b256b8d046f9c2b1c9abaf7c9656f93f26126c302c8e5d37396989d3645d74b22ac2b2dfba5eed98>

upvoted 3 times

🗉 👤 **jisjv** 1 year, 8 months ago

Nutanix Recommendations

Configure the Nutanix container for replication factor 2.

Enable compression at the container level.

Don't use container-level deduplication.

Enable erasure coding (EC-X).

upvoted 1 times

🗉 👤 **CazSand** 1 year, 10 months ago

Selected Answer: C

I think this should be C

<https://portal.nutanix.com/page/documents/solutions/details?targetId=BP-2036-Microsoft-Exchange-Server:BP-2036-Microsoft-Exchange-Server>

upvoted 3 times

An administrator learns that some security settings in a cluster have been changed. The environment is out of compliance with required policies. Which function should be used to revert any settings that have changed from the defaults and prevent further violations?

- A. Cluster Lockdown with strong password enforcement
- B. Advanced Intrusion Detection Environment (AIDE)
- C. Security Configuration Management Automation (SCMA)
- D. Defense Information Systems Agency STIG enforcement


Suggested Answer: C

Reference:

https://portal.nutanix.com/page/documents/details?targetId=Nutanix-Security-Guide-v5_17:secsecurity-scma-implementation-wc-c.html

Community vote distribution

C (100%)

 **Cromo19** 10 months, 3 weeks ago

Selected Answer: C

Answer is C: https://portal.nutanix.com/page/documents/details?targetId=Nutanix-Security-Guide-v5_17:secsecurity-scma-implementation-wc-c.html
upvoted 1 times


A customer with a four-node RF2 cluster is adding application VMs to their system. After adding these Vms, the Prism dashboard shows 81% storage utilization.

What is the consequence of running the cluster at 81% storage utilization?

- A. The customer has the ability to add more VMs up to the 100% storage utilization.
- B. There is available capacity in the storage fabric and the cluster is resilient.
- C. Node failure is imminent due to storage utilization.
- D. The cluster is not resilient in the storage fabric

Suggested Answer: D

Community vote distribution



e4af987 10 months, 1 week ago

Do the math: $M = 0.9 \times (T - N1)$. Assume total storage is 100G, that would put the max at 67.5GB or 67.5%. See this article: <https://next.nutanix.com/how-it-works-22/recommended-guidelines-for-maximum-storage-utilization-on-a-cluster-38531>
upvoted 1 times

Cromo19 1 year, 4 months ago

Selected Answer: D

Answer is D. Ref <https://portal.nutanix.com/page/documents/kbs/details?targetId=kA0600000008ducCAA>
upvoted 3 times

huskerjeff 2 years, 2 months ago

D IS THE answer because each node has 25% of storage. So if your at 81% at lose 1 node that would result in 81% with only 75% of capacity
upvoted 3 times

maxgiver 2 years, 6 months ago

Selected Answer: B

False it's 90%
Good answer is B
<https://next.nutanix.com/how-it-works-22/recommended-maximum-storage-utilization-37234>
upvoted 2 times

maxgiver 2 years, 6 months ago

False it's 90%
Good answer is B
<https://next.nutanix.com/how-it-works-22/recommended-maximum-storage-utilization-37234>
upvoted 1 times

An administrator has two clusters: Site A and Site B.

Async DR is configured between the clusters as follows:

- ☞ Site A is the Primary Site
- ☞ Site B is the Secondary inactive Site, no VMs running
- ☞ Site A is fully in Production
- ☞ Site B used as Backup/Disaster Recovery

The customer wants to perform a planned failover between the sites. How should the administrator accomplish this?

- A. Click Activate on the Secondary site within the Data Protection Dashboard
- B. Click Activate on the Secondary site and run deactivate NCLI command on Primary site
- C. Choose Migrate on the Secondary site within the Data Protection Dashboard
- D. Choose Migrate on the Primary site within the Data Protection Dashboard

Suggested Answer: D

Reference:

<https://portal.nutanix.com/page/documents/details?targetId=Web-Console-Guide-Prism-v510:wcprotection-domain-failover-wc-t.html>

Community vote distribution

D (100%)

🗨️ 👤 **Cromo19** 10 months, 3 weeks ago

Selected Answer: D

Answer is D.

upvoted 1 times

A systems administrator needs to add more VMs to their Nutanix cluster. Which two actions should the administrator perform to determine if the current cluster could accommodate the new VMs? (Choose two.)

- A. Perform an Inventory with Life Cycle Management.
- B. Determine utilization with Cluster Runway.
- C. Enable Deduplication and Erasure Coding.
- D. Utilize Optimize Resources for VM efficiency.

Suggested Answer: *BD*

Currently there are no comments in this discussion, be the first to comment!

An administrator responsible for a VDI environment needs to investigate reports of slow logins. The administrator finds that increasing the number of vCPUs from

2 to 4 will reduce the login times. Production workloads are consuming 75% of the host CPU on the cluster. The administrator increases the vCPU count on all of the VDI VMs.

What are two impacts on the cluster? (Choose two.)

- A. Increase CPU utilization %
- B. Increase CPU ready %
- C. Increase memory utilization %
- D. Increasing CPU counts will decrease memory utilization

Suggested Answer: AB

Currently there are no comments in this discussion, be the first to comment!

An administrator is deploying an application using Nutanix Volumes. After configuring the volume group and connecting the guest's iSCSI initiator to the volume group, they begin performance testing. The administrator learns that performance on the virtual disk attached to the iSCSI initiator is far less than that of a virtual disk connected directly to the VM.

Which situation is causing this condition?

- A. The VM's iSCSI traffic is being routed to the data services IP subnet
- B. The performance test is generating small block random 10
- C. The cluster data services IP address is a bottleneck, shared with other iSCSI traffic
- D. Data Locality ensures that VM attached disks are always faster than iSCSI volumes

Suggested Answer: C

Community vote distribution

A (100%)

🗲️ 👤 **pharoah2000** 11 months, 4 weeks ago

A. is correct

C. iSCSI Data Services IP enables path optimization in the cluster, preventing bottlenecks
upvoted 2 times

🗲️ 👤 **Treasure** 1 year, 1 month ago

"far less", So the answer C is more suitable
upvoted 1 times

🗲️ 👤 **JHNSMTH** 1 year, 7 months ago

Selected Answer: A

iSCSI would be routed due to different vlans, resulting in longer times.
upvoted 1 times

🗲️ 👤 **Azimuth** 2 years, 2 months ago

If the guest VM doesn't have an interface on the iSCSI network, the traffic will have to pass from the guest VM network, through a router, and then to the iSCSI network. I believe the answer is A.
upvoted 2 times

A guest VM is reported to have poor storage performance. It has an I/O profile of 80% read, 20% write, and the HDDs form more than 50% of the read source.

What should an administrator do to resolve this issue?

- A. Increase SSD capacity
- B. Increase the OPLOG
- C. Increase the write cache
- D. Use HDDs with higher RPM

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

An administrator is deploying a three-tier application on a Nutanix Enterprise Cloud consisting of web, application, and database VMs. Traffic between the servers must be limited to specific services/ports.

There is an application layer firewall VM deployed on one host in the cluster.

In which way should the administrator meet this requirement?

- A. Create a network for each type of VM, and use policy-based routing to control the traffic flow
- B. Configure a security policy and service chain redirecting the server traffic through the firewall
- C. Separate the web application and database VMs on to dedicated VLANs
- D. Categorize the VMs in Prism Central and utilize Flow to set a security policy

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

An administrator needs to make sure an RF2 cluster can survive a complete rack failure without negatively effecting workload performance. The current cluster configuration consists of the following:

- ⇒ 30 nodes: distributed 10 nodes per rack across three 42U rack
- ⇒ Each nodes is configured with 40TB usable storage all flash (Cluster Total 1.2 PB Usable)
- ⇒ Current cluster utilization is 900TB storage

Which configuration changes should be made to make sure that the cluster meets the requirements?

- A. Expanse the cluster by adding 3 additional nodes per rack, 9 nodes total.
- B. Expanse the cluster by adding 2 additional nodes per rack, 6 nodes total.
- C. Expanse the cluster by adding 9 nodes to a new 42U rack
- D. Expanse the cluster by adding 8 nodes to a new 42U rack

Suggested Answer: A

🗨️ 👤 **Chacal001** 1 year, 3 months ago

you need 50TB additional per rack (in order to get 450TB per rack), each node has 40TB so you need 2 more per rack, but 1 additional node (per rack) in order to get N+1.

upvoted 2 times

🗨️ 👤 **Josantru** 1 year, 7 months ago

What is the reason? 39 nodes?

upvoted 1 times

A company needs to replicate their VMs to a remote site and must have an RPO of less than 1 hour. When creating a protection domain, what should an administrator do to meet this requirement?

- A. Enable Async Replication with an RPO of 60 minutes
- B. Enable Near-Sync Replication with an RPO of 15 minutes
- C. Enable Entity Sync on the designated VMs
- D. Enable Availability Zones and send the VM to multiple targets

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **pharoah2000** 11 months, 3 weeks ago

Selected Answer: B

- A. Async Replication can not do less than 1 hour
 - B. Nar-Sync can do less than 1 hour
 - C. No such option
 - D. Availability Zones has nothing to do with replication policy
- upvoted 2 times

🗳️ 👤 **ttyso** 1 year, 5 months ago

Selected Answer: B

near sync have 15 RPO min
upvoted 2 times

🗳️ 👤 **EhCloudGrrrU** 1 year, 9 months ago

Selected Answer: B

It's near-sync with a 15 min RPO.
upvoted 2 times

🗳️ 👤 **Josantru** 2 years ago

The question indicates less than one hour. Shouldn't it be option b? With near-sync, 15 minutes.
upvoted 4 times

An administrator is monitoring the Nutanix v5.15-based AOS cluster performance logs and notices that a SQL server VM is greatly exceeding its intended maximum IOPS. The administrator has confirmed that a QoS policy was previously created for the group of VMs this server is a member of.

What are two reasons that this VM would exceed its maximum configured IOPS from the QoS Policy? (Choose two.)

- A. The VM is missing the required snapshot needed to implement the QoS Policy.
- B. The VM was not created as an AFS virtual machine.
- C. The SQL Server VM has volume groups attached.
- D. It was cloned from another SQL Server VM that was a member of the QoS Policy.

Suggested Answer: CD

Community vote distribution

CD (100%)

 **jssanche2019** Highly Voted 2 years, 7 months ago

C and D are corrects. If a VM has volume groups attached, then QoS is not applicable and While creating a clone of a VM, any QoS attributes throttle limit set on the original VM might not be applied to the new cloned VM


upvoted 8 times

 **hi68340** Most Recent 7 months, 4 weeks ago

Selected Answer: CD

C and D are corrects

upvoted 1 times

 **ziggy675** 10 months, 3 weeks ago

- A. The VM is missing the required snapshot needed to implement the QoS Policy.
- D. It was cloned from another SQL Server VM that was a member of the QoS Policy.

Explanation:

The absence of the required snapshot: When a QoS policy is created for a group of VMs, it requires a snapshot to be taken in order to implement the policy. If the VM in question is missing this required snapshot, the QoS policy may not be enforced, allowing the VM to exceed its intended maximum IOPS.

Cloning from another VM: If the SQL Server VM was cloned from another SQL Server VM that was already a member of the QoS policy, the new VM may inherit the QoS policy settings, including the maximum configured IOPS. However, if the new VM has different workload characteristics or if the QoS policy was not properly adjusted for the new VM, it may exceed the intended maximum IOPS.

upvoted 1 times

 **nabillose** 1 year, 1 month ago

C and D are corrects

upvoted 2 times

 **EhCloudGrrrU** 1 year, 9 months ago

Selected Answer: CD

'C' and 'D' are correct.

upvoted 2 times

A customer has a Nutanix cluster with 10Gb connectivity via switch fabric extenders. The administrator receives NCC health check errors of latency greater than 200ms.

Which action should the administrator take to resolve the NCC errors?

- A. Replace the switch fabric extenders with 10G line rate switches
- B. Upgrade NCC and increase the CVM memory by 4Gb
- C. Add 2 additional 10G uplinks from the switch fabric extenders per node
- D. Upgrade NCC and increase the vCPU of the CVM

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!


An administrator knows that the disks in a cluster are out of date and a new firmware is available. For security reasons, these clusters do not have access to the Internet.

Which two steps must be completed to update the firmware? (Choose two.)

- A. Configure LCM with a URL to access the firmware bundle.
- B. Download a dark site bundle and stage it on an internal webserver.
- C. Download the updated firmware .tgz file from the Nutanix portal.
- D. Select Upgrade Software, then select disks and upload the updated firmware.

Suggested Answer: AB

Community vote distribution



🗳️ 👤 **Xoxoo** 1 year, 3 months ago

Selected Answer: AB

D is wrong as under Upgrade Software in Prism Element, you can only upgrade AOS, File Server, Hypervisor, NCC and Foundation. No option for you to upgrade firmware as this is done in the LCM in prism element.

upvoted 1 times

🗳️ 👤 **pharoah2000** 1 year, 5 months ago

B and C are correct.

A. Configure LCM with the URL of an internal webserver

D. No such option

upvoted 1 times

🗳️ 👤 **Poddo** 1 year, 6 months ago

correct answer is A and B

upvoted 1 times

🗳️ 👤 **nabillose** 1 year, 7 months ago

correct answer is A and B

upvoted 1 times

🗳️ 👤 **ttyso** 1 year, 11 months ago

Selected Answer: AB

<https://next.nutanix.com/installation-configuration-23/how-to-configure-a-web-server-for-lcm-dark-site-on-a-linux-machine-38045>

upvoted 1 times

🗳️ 👤 **maxgiver** 2 years ago

Selected Answer: CD

No no no

Correct Answer is C and D

upvoted 1 times

A customer needs a DR solution designed around the following requirements:


- ☞ There are 8 ms of latency between Sites A and B
- ☞ The administrator needs the shortest RPO/RTO possible
- ☞ Application consistent snapshots are required
- ☞ Use of Nutanix VSS

Which Data Protection Method should be used?

- A. Async
- B. Metro Availability
- C. Near Sync
- D. Cloud Connect

Suggested Answer: A

Community vote distribution



🗳️ **Xoxoo** 1 year, 3 months ago

Selected Answer: C

C can be correct now as NearSync can support application-consistent snapshots from AOS 5.20 onwards. However, minutely snapshots are not application-consistent but hourly snapshots are.

Source: https://portal.nutanix.com/page/documents/details?targetId=Prism-Element-Data-Protection-Guide-v5_20:wc-dr-near-sync-c.html

upvoted 2 times

🗳️ **pharoah2000** 1 year, 5 months ago

A is correct

C. is incorrect

+ NearSync does not support application-consistent snapshots

+ NearSync replication schedules of less than 15 minutes do not support self-service restore.

https://portal.nutanix.com/page/documents/details?targetId=Prism-Element-Data-Protection-Guide-v6_6:wc-cluster-dr-nearsync-limitations-r.html

upvoted 2 times

🗳️ **ttyso** 1 year, 11 months ago

Selected Answer: A

can not be NearSync.because of app consistency

upvoted 2 times

🗳️ **Chacal001** 2 years, 3 months ago

But you can add more snapshot schedules each 1 hr, it supports application consistent snapshots.

upvoted 4 times

🗳️ **EhCloudGrrrU** 2 years, 3 months ago

Selected Answer: A

Due to app consistency it cannot be NearSync.

upvoted 2 times

🗳️ **SG1ATLANTIS** 2 years, 5 months ago

Response A. NearSync doesn't support application consistent

https://portal.nutanix.com/page/documents/details?targetId=Prism-Element-Data-Protection-Guide-v6_0:wc-cluster-dr-nearsync-limitations-r.html

upvoted 3 times

A VM is exhibiting one or more of the following baseline values based on the past 30 days:

⇒ CPU usage < 20%

⇒ CPU ready time < 5%

Memory usage < 50% (moderately) or < 20% (severely)

•
⇒ Memory swap rate = 0 Kbps

Which type of VM is being described?

- A. Constrained VM
- B. Inactive VM
- C. Bully VM
- D. Over-Provisioned VM

Suggested Answer: D

Over-provisioned VM

Over-provisioned VM is the opposite of a constrained VM, meaning it is a VM that is over-sized and wasting resources which are not needed. A VM is considered over-provisioned when it exhibits one or more of the following baseline values, based on the past 30 days:

- CPU usage < 20% and CPU ready time < 5%
- Memory usage < 50% (moderately) or < 20% (severely) and memory swap rate = 0 Kbps

To prevent host resource wastage, resize (decrease) the over-provisioned VMs.

Constrained VM

Constrained VM is one that does not have enough resources for the demand and can lead to performance bottlenecks. A VM is considered constrained when it exhibits one or more of the following baseline values, based on the past 30 days:

- CPU usage > 90% (moderate), 95% (high)
- CPU ready time > 5% , 10%
- Memory usage > 90%, 95%
- Memory swap rate > 0 Kbps (no moderate value)

One or more constrained VMs might cause a performance bottleneck. To provide adequate host resources, resize (increase) the constrained VMs.

Bully VM

Bully VM is one that consumes too many resources and causes other VMs to starve. A VM is considered a bully when it exhibits one or more of the following conditions for over an hour:

- CPU ready time > 5%
- Memory swap rate > 0 Kbps
- Host I/O Stargate CPU usage > 85%

One or more bully VMs might cause cluster performance to degrade. Identifying bully VMs can help in analyzing whether one or more of those VMs are misbehaving or need additional resources.

Inactive VM

Inactive VM is in either of the following states:

- **Dead VM** : A VM is considered dead when it has been powered off for at least 30 days.
- **Zombie VM** : A VM is considered a zombie when it is powered on but does fewer than 30 read or write I/Os (total) and receives or transfers fewer than 1000 bytes per day for the past 30 days.

Currently there are no comments in this discussion, be the first to comment!

An administrator needs to forecast infrastructure requirements for a new program and its associated applications. Prior to the projected start of the new program, all existing applications will be decommissioned.

How should the administrator perform this task?

- A. Check the Disregard Existing Workloads radio button in the Runway scenario.
- B. Check the Disregard Existing Nodes radio button in the Runway scenario.
- C. Add up the recovered workloads and manually remove from the Runway configuration.
- D. Power down the workloads during a maintenance window and run the Capacity Runway.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

An administrator is deploying several new application VMs to a Nutanix cluster. The application is very transactional with a 28/72 random, read/write ratio.

What should the administrator do?

- A. Add a minimum of four SSDs to each node to ensure all writes go to the SSDs.
- B. Add a minimum of four virtual disks to each VM, and aggregate them on guest OS level.
- C. Add one large SSD to each node and enable Flash Mode for the VMs.
- D. Add one large virtual disk to each VM and use it for all application data.

Suggested Answer: B

>Multiple vdisks should be attached (they can be unified with LVM for example) to get more performance from VM's, as Nutanix OS limiting oplog size per vdisk

(to avoid a "noisy neighbor" problem).

Reference:

<https://next.nutanix.com/server-virtualization-27/iops-latency-issue-1168>

Currently there are no comments in this discussion, be the first to comment!

An alert about RX errors on eth2 on a node is reported in the cluster. The administrator logs in to the CVM in question, checks the ping_* files in the data/logs/ syststats folder, and notices intermittent ping loss. The node in question has the following network configuration:

```
nutanix@NTNX-16xxxxxxxxx5-A-CVM:xx.xx.xx.76:~$ manage_OVS show_uplinks
```

```
Bridge: br0
```

```
Bond: br0-up
```

```
bond_mode: active-backup
```

```
interfaces: eth3 eth2
```

```
lacp: off
```

```
lacp-fallback: True
```

```
lacp_speed: slow
```

```
Bridge: br1
```

```
Bond: eth0
```

```
bond_mode: active-backup
```

```
[interfaces: eth0
```

```
lacp: off
```

```
lacp-fallback: True
```

```
lacp_speed: slow
```

```
Bridge: br1
```

```
Bond: eth0
```

```
bond_mode: active-backup
```

```
interfaces: eth0
```

```
lacp: off
```

```
lacp-fallback: false
```

```
lacp_speed: slow
```

```
Bridge: br2
```

```
Bridge: br5
```

```
nutanix@NTNX-16xxxxxxxxx5-A-CVM:xx.xx.xx.76:~$ manage_ovs show_interfaces
```

```
name mode link speed
```

```
eth0 1000 True 1000
```

```
eth1 1000 True 1000
```

```
eth2 10000 True 10000
```

```
eth3 10000 True 10000
```

```
nutanix@NTNX-16xxxxxxxxx5-A-CVM:xx.xx.xx.76:~$
```

Which action should be used to troubleshoot without disrupting the VMs running on this node?

- A. Replace the cable from eth2 to the switch
- B. Remove eth3 from br0-up and monitor for new alerts
- C. Check the port on the switch side for any errors
- D. Enable balance-s1b on br0-up and monitor if a problem persists

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **EhCloudGrrrU** 1 year, 3 months ago

Selected Answer: C

Answer is 'C'.

upvoted 2 times

🗳️ 👤 **zmu** 1 year, 9 months ago

Only D is without disrupting the VMs running

upvoted 2 times

🗳️ 👤 **jerem1183** 1 year ago

To enable balance-s1b you have to put the node and CVM in the maintenance mode so it cannot be D

upvoted 1 times

A three-node cluster has a Nutanix Files instance with three FSVMs.

What happens after a node failure?

- A. The affected FSVM will get live-migrated to another node.
- B. The affected FSVM will not be restarted on another node due to its agent VM setting.
- C. The affected FSVM will not be restarted on another node due to its VM-to-VM anti-affinity settings.
- D. The affected FSVM will be restarted on another node.

Suggested Answer: *D*

Reference:

https://portal.nutanix.com/page/documents/details?targetId=Files-v3_6:fil-files-high-availability-c.html

Currently there are no comments in this discussion, be the first to comment!

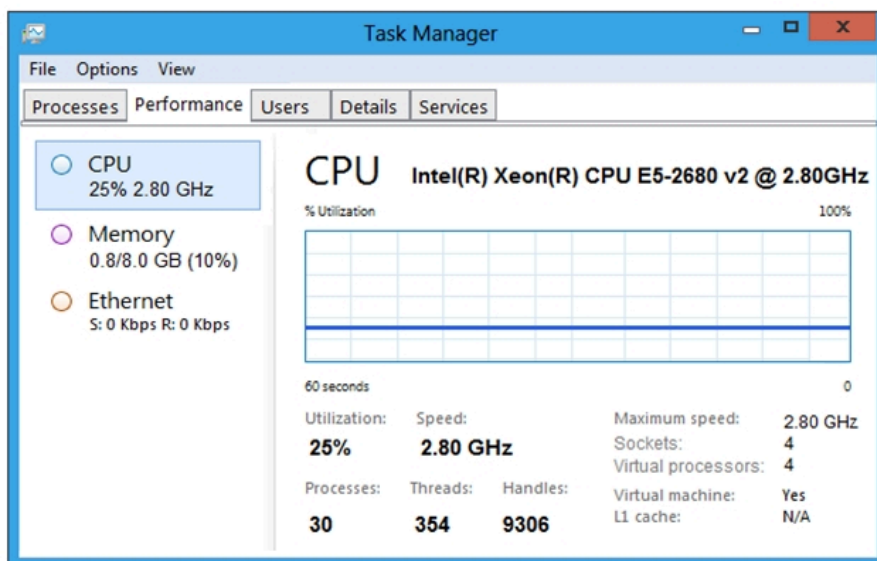

Refer to the exhibits.

An administrator has reported that a new VM is not performing well.

Prism Element, VM Table Metrics

VM Name: NCNCAPVM
 Cores: 4
 Memory Capacity: 8GiB
 Storage: 9.86 GiB / 40 GiB
 CPU Usage: 25.85%
 Memory Usage: 0%
 Controller Read IOPS: 0
 Controller Write IOPS: 0
 Controller I/O Bandwidth: 13 KBps
 Controller Avg IO Latency: 2.33ms
 Backup ...: Yes
 Flash Mode: No

Prism Element, Analysis Chart





After analyzing the information presented in the exhibits, which option shows the best conclusion regarding this workload?

- A. VM is CPU constrained since the hypervisor cannot provide the required resources.
- B. VM is CPU constrained because more CPU needs to be added to the VM.
- C. VM is Storage constrained, because the workload is waiting for storage access.
- D. VM is not CPU constrained, because the VM is not running, a multithreaded application.

Suggested Answer: D

Community vote distribution

D (100%)

  **EhCloudGrrrU** 1 year, 3 months ago

Selected Answer: D

Answer is 'D' - no clue why storage is even chosen here, it's not part of the diagram.

upvoted 4 times

An administrator is performing a set of routine checks and notes that Prism Central is running with 4 vCPUs and 21 GB RAM. The organizations environment is based on two Nutanix clusters with 170 VMs plus a Legacy vSphere infrastructure. No manual changes have been done on Prism Central VM.

Which configuration justifies the current Prism Central configuration?

- A. Nutanix Move is running on Prism Central to import VMs from the Legacy Cluster.
- B. Prism Central was installed for Large Deployment environments.
- C. Prism Central is running on the legacy cluster and requires more RAM.
- D. Nutanix Leap and Nutanix Flow have been enabled on Prism Central

Suggested Answer: D

Community vote distribution

C (100%)

  **pharoah2000** 11 months, 4 weeks ago

Selected Answer: C

<https://portal.nutanix.com/page/documents/kbs/details?targetId=kA00e000000CvAnCAK>

upvoted 2 times

  **pharoah2000** 11 months, 4 weeks ago

A is incorrect - Nutanix Move does not import VM, it does VM migration

B and D are incorrect - refer to <https://portal.nutanix.com/kb/8932>

C is correct, 1-VM PC requires 6 vCPUs, 26 GB of memory, and 500 GiB of storage.

upvoted 1 times

A customer wants to run SAP HANA on a Nutanix cluster with the following characteristics:

- ☞ AHV nodes
- ☞ SAP HANA version 2
- ☞ Production database files

Which configuration should be recommended to the customer?

- A. Start with a cluster containing three or more nodes
- B. Enable only compression for the container that contains the Production database
- C. Do not enable compression, deduplication, or erasure coding on the storage container
- D. Place the SAP HANA Database on the same socket as CVM

Suggested Answer: C

Do not configure storage-saving functionalities such as compression, deduplication, or erasure coding (EC-X) on a storage container that holds production database files. Because of the way the SAP HANA Persistence Engine stores data, these features deliver no benefits. Nutanix SAP Engineering has tested compression with SAP HANA workloads; the tests indicate no noticeable performance impact but also show no reduction in the space the SAP HANA workload consumes.

Reference:

<https://www.nutanix.com/go/sap-hana-on-nutanix-bpg>

Community vote distribution

C (100%)

🗳️ 👤 **pharoah2000** 11 months, 4 weeks ago

C is correct

Don't configure storage-saving functionalities such as compression, deduplication, or erasure coding (EC-X) on a storage container that holds production database files.

<https://portal.nutanix.com/page/documents/solutions/details?targetId=BP-2097-SAP-HANA-on-AHV:BP-2097-SAP-HANA-on-AHV>

upvoted 1 times

🗳️ 👤 **EhCloudGrrrU** 1 year, 9 months ago

Selected Answer: C

C is correct - do not turn on any storage savings.

upvoted 2 times

🗳️ 👤 **Josantru** 2 years ago

I think, the correct is B.

upvoted 1 times

🗳️ 👤 **examgroup123456** 1 year, 6 months ago

<https://portal.nutanix.com/page/documents/solutions/details?targetId=BP-2097-SAP-HANA-on-AHV:BP-2097-SAP-HANA-on-AHV>

Correct Answer is C

upvoted 1 times

An administrator receives reports that VDI desktop performance in an 8-node Nutanix VDI environment is poor. Opening applications takes between 1 and 2 minutes. When investigating the issue, the following conditions are found:

- ⇒ Cluster memory utilization: 80%
- ⇒ Cluster SSD utilization: 70%
- ⇒ Average VM CPU Wait Time: 11%
- ⇒ CVM CPU utilization: 75%

Which action should be taken to improve VDI performance?

- A. Increase the amount of SSD storage in the cluster.
- B. Add memory to the nodes in the cluster.
- C. Add CPU resources to the cluster.
- D. Increase the number of vCPU cores allocated to the CVM.

Suggested Answer: C

🗉 👤 **jerem1183** 1 year ago

Key element: CPU Wait Time at 11%

upvoted 2 times

An administrator receives an alert in Prism indicating that interface eth2, on an AHV host is receiving many CRC errors. After toggging into the problematic host, the following command is run to show the indicated output:

```
[root@AHV-Host ~]# ethtool -S eth2 | grep error
rx_errors: 0
tx_errors: 0
rx_over_errors: 0
rx_crc_errors: 478593
rx_frame_errors: 0
rx_fifo_errors: 0
rx_missed_errors: 0
tx_aborted_errors: 0
tx_carrier_errors: 0
tx_fifo_errors: 0
tx_heartbeat_errors: 0
rx_long_length_errors: 0
rx_short_length_errors: 0
rx_csum_offload_errors: 0
```

What is causing this issue?

- A. Incorrect link speeds on the switch
- B. The interface is incorrectly configured with Jumbo Frames
- C. A misconfigured bond
- D. A physical layer network problem

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **pharoah2000** 11 months, 4 weeks ago

D is correct

<https://portal.nutanix.com/kb/1381>

upvoted 2 times

🗳️ 👤 **18041987** 1 year, 2 months ago

Selected Answer: D

Has to be D

upvoted 2 times

🗳️ 👤 **EhCloudGrrrU** 1 year, 9 months ago

Selected Answer: D

It is 'D' a physical layer problem.

upvoted 3 times

🗳️ 👤 **quevi15** 1 year, 10 months ago

Correct answer should be "physical layer network problem"

upvoted 4 times




An administrator needs to migrate workloads from a Nutanix cluster running VMware ESXi to Nutanix AHV. The migration process needs to be completed in multiple VM groups and support roll back in case any problems are found during user acceptance testing. Which action should the administrator perform?

- A. Use Cross Hypervisor DR to replicate VMs from ESXi to AHV
- B. Use VMware Converter
- C. Use Nutanix one-click cluster conversion from ESXi to AHV
- D. Use storage live migration

Suggested Answer: A




Community vote distribution

A (100%)

  **jssanche2019**  3 years, 1 month ago



A is correct. Cross hypervisor disaster recovery (CHDR) provides an ability to migrate the VMs from one hypervisor to another (ESXi to AHV or AHV to ESXi) by using the Protection Domain semantics of protecting VMs, taking snapshots, replicating the snapshots, and then recovering the VMs from the snapshots

upvoted 10 times

  **alibubba**  1 year, 1 month ago

This statement "support roll back in case any problems" actually does point to in-place upgrade option. But yes DR replication from ESXi to AHV would be safest option.

upvoted 1 times

  **Sprwl** 1 year, 7 months ago

 **Selected Answer: A**

Also Voting for „A“

With covert cluster you cannot rollback.

upvoted 3 times

  **EhCloudGrrrU** 2 years, 3 months ago

 **Selected Answer: A**

It's 'A'. Comment is correct.

upvoted 2 times

Refer to the exhibit.

Overview

Table

Async DR		Remote Site
Name		Async DR Alerts
ADR_local		• Unable to locate VM(s) Lin_Ntnx_Uvm_1 protected by protection domain ADR_local_no_cg
ADR_local_no_cg		
ADR_remote		• Unable to locate VM(s) Lin_Ntnx_Uvm_1 protected by protection domain ADR_local_no_cg
ADR_remote_no_cg		
ADR_remote_palo02		• Unable to locate VM(s) Lin_Ntnx_Uvm_1 protected by protection domain ADR_local_no_cg

After performing ESXi to AHV conversion on a cluster, the alerts shown in the exhibit display on Prism. What should the administrator do to resolve this issue efficiently?

- A. Remove the VMs, create a new Protection Domain, and add VMs to it.
- B. Remove VMs from the Protection Domain and add them back in.
- C. Delete the Protection Domain, create a new Protection Domain, and add the VMs to it.
- D. Create a new Protection Domain and add the VMs to it.

Suggested Answer: B

Currently there are no comments in this discussion, be the first to comment!

For several days, an administrator notices the following alerts:

- ☞ CVM NIC Speed Low Warning Alerts
- ☞ Warning Alerts of CVM NIC not performing at optimal speed
- ☞ CVM is disconnected from the network Critical Alert
- ☞ Network Visualization page shows excessive dropped packets on CVM/Host.

Which steps should be taken to determine which problem should be addressed first?

A.

- ☞ Access the Hardware Page to verify resources are available
- ☞ Analyze all CVM Speed Alerts in the Alerts/Events page
- ☞ Analyze output for the network and interface properties and connectivity issues

B.

- ☞ Verify Host/CVM connectivity on the Network Visualization page
- ☞ Use to verify the bridge and bond configuration
- ☞ Review alerts/events page for the CVM disconnected error

C.

- ☞ Review Alerts page for NIC speed alerts and alert timing
- ☞ Analyze the genesis.out log file for process failures
- ☞ Assess the NIC properties in the Network Visualization > Host Properties page

D.

- ☞ Restart networking services on the CVM
- ☞ Determine the current configuration of the affected CVM via output
- ☞ Access the Alerts/Events page for the CVM network connection failures

Suggested Answer: B

Currently there are no comments in this discussion, be the first to comment!

An administrator needs to relocate an AHV cluster to a new datacenter during a maintenance window. The cluster will use the same IPs in the new datacenter.

Which two steps should be taken to prepare for this task? (Choose two.)

- A. Stop all Nutanix Files clusters
- B. Relocate the linked LDAP servers
- C. Shut down all user VMs in the cluster
- D. Reconfigure IPMI for the new datacenter

Suggested Answer: AC

Reference:

<https://next.nutanix.com/installation-configuration-23/physical-relocation-of-nutanix-clusters-38403>

Currently there are no comments in this discussion, be the first to comment!

An administrator needs to evaluate if the existing configuration of Protection Domains meets the specified corporate recovery policy for SQL Data. The corporate policy requires that all SQL Servers be recoverable at the DR site within 30 minutes, and must result in a maximum of 60 minutes of data loss. SQL data must be required.

The protection domain is configured as follows:

- ⇒ Application consistent snapshots are enabled
- ⇒ The SQL server VMs are assigned to the Protection Domain as their own Consistency Group
- ⇒ The Protection Domain schedule is configured as snapshots taken every 90 minutes from 7 am until 8 pm

What must be changed to meet the required policy?

- A. Revise the schedule to be 15 minutes
- B. Revise the schedule to be 60 minutes
- C. Revise the schedule to be 30 minutes
- D. Revise the schedule to be 45 minutes

Suggested Answer: B

🗨️ 👤 **klwood2000** 10 months ago

C. You want to replicate half the time of the RPO. If you have a 60 min RPO, and you lose a replication point, you have to roll back 120 minutes. If you replicate every 30 minutes and lose a cycle, at most you have to go back is 60 minutes which meets the requirement.

upvoted 1 times

An administrator needs to make sure an RF3 cluster can survive a failure of two complete racks without negatively affecting performance. The current cluster configuration consists of the following:

- ⇒ 30 All Flash Nodes: distributed 10 nodes per rack across three 42U racks
- ⇒ Each node is configured with 20TB usable storage all flash (Cluster Total 600TB Usable)
- ⇒ Current cluster utilization is 500TB storage
- ⇒ Storage containers have Erasure Coding enabled

Which configuration changes should be made to make sure the cluster meets the requirements?

- A. Expand the cluster to 50 nodes distributed evenly across 5 racks
- B. Expand the cluster to 48 nodes distributed evenly across 6 racks
- C. Expand the cluster to 60 nodes distributed evenly across 6 racks
- D. Expand the cluster to 40 nodes distributed evenly across 5 racks

Suggested Answer: C

🗨️ 👤 **sa12** 1 year ago

"A" could be correct too, with 5, 10 node rack we will have 1000TB total cluster capacity. If we deduct the 500TB in use, we can still lose two racks (total of 400TB) and still maintain two copies of each block of data on the cluster.

upvoted 2 times

🗨️ 👤 **Iski** 11 months, 4 weeks ago

Erasure Coding requires an additional rack.

upvoted 1 times

An administrator receives an error indicating that the CVMs in the cluster are not syncing to any NTP servers. An investigation of the issue finds:

- ☞ The NTP servers are configured in Prism
- ☞ The time on all CVMs is the same
- ☞ Both the CVMs and AHV hosts are configured for the UTC time zone

Which two steps can be taken to troubleshoot this issue? (Choose two.)

- A. Confirm that the NTP servers are reachable from the CVMs.
- B. Restart genesis on all CVMs.
- C. On a CVM. run the command `a Hash ntpq -pn`.
- D. Restart the chronos service on all CVMs.

Suggested Answer: AC

🗨️ 👤 **double_nope** 1 year ago

Answer "c." should read:

"C. On a CVM. run the command `allssh ntpq -pn`"

upvoted 3 times

HOTSPOT -

An administrator is configuring an All Flash cluster to support the following workloads:

- ⇒ Full clone virtual desktops
- ⇒ Hadoop
- ⇒ Write heavy databases

Which correct configuration options should the administrator use to meet the item requirements? For each procedure, indicate whether it is a correct or incorrect procedure that must be configured to meet the item requirements using the drop-down option.

Hot Area:

Answer Area

Procedure	Correct/Incorrect
One container with compression and deduplication enabled for full clone virtual desktops	<div><div></div><div>Correct</div><div>Incorrect</div></div>
One container with compression and Erasure Coding enabled for Hadoop	<div><div></div><div>Correct</div><div>Incorrect</div></div>
One container with compression enabled for heavy databases	<div><div></div><div>Correct</div><div>Incorrect</div></div>
One container with deduplication enabled for full clone virtual desktops	<div><div></div><div>Correct</div><div>Incorrect</div></div>
One container with Erasure Coding enabled for Hadoop	<div><div></div><div>Correct</div><div>Incorrect</div></div>
One container with deduplication enabled for Hadoop	<div><div></div><div>Correct</div><div>Incorrect</div></div>

Suggested Answer:

Answer Area

Procedure

One container with compression and deduplication enabled for full clone virtual desktops

Correct/Incorrect

▼

Correct

Incorrect

One container with compression and Erasure Coding enabled for Hadoop

▼

Correct

Incorrect

One container with compression enabled for heavy databases

▼

Correct

Incorrect

One container with deduplication enabled for full clone virtual desktops

▼

Correct

Incorrect

One container with Erasure Coding enabled for Hadoop

▼

Correct

Incorrect

One container with deduplication enabled for Hadoop

▼

Correct

Incorrect

Currently there are no comments in this discussion, be the first to comment!