



Reference Solution for 2V0-621.exam.147q 2V0-621

2V0-621

VMware Certified Professional 6 – Data Center Virtualization

Score: 800/1000

Version:

Time Limit: 120 Minutes

Exam A (147 questions)

Question 1

When upgrading vCenter Server, an administrator notices that the upgrade fails at the vCenter Single Sign-On installation.



What must be done to allow the upgrade to complete?

- ☒ Verify that the VMware Directory service can stop by manually restarting it.
- ☐ Verify that the vCenter Single Sign-On service can stop by manually restarting it.
- ☐ Uninstall vCenter Single Sign-On service.
- ☐ Uninstall the VMware Directory service.

Explanation:

Explanation:


Prerequisites

- ? Your current vCenter Single Sign-On must have been installed on a separate virtual machine (VM) or physical server from your vCenter Server instance.
- ? Verify your configuration meets the upgrade requirements, see vCenter Server for Windows Requirements.
- ? Complete the preparation to upgrade tasks. See Before Upgrading vCenter Server
- ? Verify that you have made a backup of your vCenter Server configuration and database.
- ? To verify that the VMware Directory Service is in a stable state and can stop, manually restart it. The VMware Directory service must be stopped for the vCenter Server upgrade software to uninstall vCenter Single Sign-On during the upgrade process.
- ? Download the vCenter Server Installer. See Download the vCenter Server for Windows Installer

Reference: <https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.upgrade.doc%2FGUID-A94C1617-5F15-402A-B462-1AC6A041C73E.html>

Question 2

Refer to the Exhibit.



The Prod-DB virtual machine has a VM Override as shown in the Exhibit.

What step, if taken, would require all virtual machines in the cluster to migrate automatically?

- ☐ Deselect the virtual machine from VM Overrides.
- ☐ Add all virtual machines to the VM Overrides.
- ☐ Change Response for Host Isolation to Use Cluster Settings.
- ☒ Change the Automation level to Use Cluster Settings.

Explanation:

Explanation:

To migrate virtual machine in the cluster automatically, you need to change the automation level to use cluster settings. The automation level will then use cluster settings to automatically migrate virtual machines in the cluster.

Question 3

Which two statements are true regarding Virtual SAN Fault Domains? (Choose two.)

- ☒ They enable Virtual SAN to tolerate the failure of an entire physical rack.
- ☒ Virtual SAN ensures that no two replicas are provisioned on the same domain.
- ☐ Virtual SAN ensures that all replicas are provisioned on the same domain.
- ☐ They require VMware High Availability (HA) to ensure component distribution across domains.

Explanation:

Explanation:

Managing Fault Domains in Virtual SAN Clusters If your Virtual SAN cluster spans across multiple racks or blade server chassis in a data center and you want to make sure that your hosts are protected against rack or chassis failure, you can create fault domains and add one or more hosts to each fault

domain. A fault domain consists of one or more Virtual SAN hosts grouped together according to their physical location in the data center. When configured, fault domains enable Virtual SAN to tolerate failures of entire physical racks as well as failures of a single host, capacity device, network link or a network switch dedicated to a fault domain.

Reference: <https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.virtualsan.doc%2FGUID-8491C4B0-6F94-4023-8C7A-FD7B40D0368D.html>

Question 4

A failed upgrade from vCenter Server version 5.x to version 6.0 produces the following error:

[00800 error 'Default'] Database version id '600' is incompatible with this release of VirtualCenter.

What is the cause of the upgrade failure?

- ☒ There was a database schema upgrade failure during the installation.
- ☐ The VMWAREVCMSSDS service was upgraded before the vCenter Server service.
- ☐ The VMware Directory Service database failed during the installation.
- ☐ There was an incompatible ODBC driver version for the database.

Explanation:

Explanation: Attempts to upgrade the VMware vCenter Server from 5.x to 6.0 might fail when validating the database

Upgrading VMware vCenter Server with an Oracle Database from 5.x to 6.0 might fail. This error occurs when you install a vCenter Server against an external Microsoft SQL database. You need to create the database schema manually by referencing to the information in DB_and_schema_creation_scripts_mssql.txt in the DB scripts folder. An error message similar to the following is displayed:

The user associated with the DSN has insufficient privileges.

This issue is resolved in this release. VMware vCenter Server 6.0 Update 1

check release notes whitepaper on vmware.com

Upgrading from vSphere 5.x to vSphere 6.0 Best Practices (2130664)

https://kb.vmware.com/selfservice/search.do?cmd=displayKC&docType=kc&docTypeID=DT_KB_1_1&externalId=2130664

Question 5

An administrator tries to connect the vSphere 5.5 Client to an ESXi 6.x host.

◆◆

What will happen when this takes place?

- ☐ The operation will fail, since the vSphere Client is deprecated in vSphere 6.x.
- ☐ The operation will fail and the administrator will need to delete the client and install the 6.x version.
- ☒ The operation will prompt the administrator to run a script to upgrade the vSphere Client.
- ☐ The operation will update the vSphere Client silently in the background, then connect.

Explanation:

Explanation:

Whenever a user tries to connect to a higher version of the ESXi host, the user will be prompted to upgrade to the higher version to escape the compatibility issues.

Question 6

Which three connection types are supported between a remote site and vCloud Air? (Choose three.)

- ☒ Secure Internet Connectivity
- ☐ Private Connect
- ☒ Direct Connect
- ☐ Internet Connectivity
- ☒ Secure VPN

Explanation:

Explanation:

VPN and Remote Networks

Another aspect of network security is the connectivity you establish between your on premise data center and the vCloud Air cloud. Based on the workload, a virtual machine can have various connectivity needs.

vCloud Air supports Internet, secure VPN, and Direct Connect secure connections between your remote site and the vCloud Air cloud.

Each type of connection has different security features:


- Secure Internet connectivity with firewall rules (a gateway service)
See “About Firewall Rules,” on page 31 in this guide for information.
- Secure VPN:
 - IPsec VPN—secure site-to-site VPN
See “About IPsec VPN,” on page 33 and “Set up an IPsec VPN Connection to a Remote Site,” on page 35 in this guide for information.
 - SSL VPN (Data Center Extension)—extension of your existing IP address range from your on-premise data center into the cloud with Layer 2 extension
See “SSL VPN for Data Center Extension,” on page 36 in this guide for information.
- Direct Connect—private connectivity providing a dedicated, connection (ideal for regulated applications)

Reference:

http://pubs.vmware.com/vca/topic/com.vmware.ICbase/PDF/vcloud_air_networking_guide.pdf

Question 7

Refer to the Exhibit.



```
ds-lab-vsan12.dsl.vmware.com - PuTTY
[root@ds-lab-vsan12:~] esxcli vsan policy getdefault
Policy Class  Policy Value
-----
cluster      (("hostFailuresToTolerate" i1))
vdisk        (("hostFailuresToTolerate" i1))
vmnamespace  (("hostFailuresToTolerate" i1))
vmswap       (("hostFailuresToTolerate" i1) ("forceProvisioning" i1))
vmem         (("hostFailuresToTolerate" i1) ("forceProvisioning" i1))
[root@ds-lab-vsan12:~]
```

An administrator recently created a Virtual SAN but no Storage Policies were defined. A few virtual machines were deployed to this cluster. The administrator analyzes the default Virtual SAN policy as shown in the Exhibit.



Based on the exhibit, which two statements are true? (Choose two.)

- ☒ Losing one cluster node will not affect data availability.
- ☒ Losing one Hard Disk in a cluster node will not affect data availability.
- ☐ Creating a virtual machine Swap file will fail if it violates default storage policy
- ☐ Creating a virtual machine will succeed even if it violates default storage policy

Explanation:

Explanation:

Policies Rules: Defines the number of host, disk, or network failures a virtual machine object can tolerate. For n failures tolerated, $n+1$ copies of the virtual machine object are created and $2n+1$ hosts with storage are required.

Default value is 1. Maximum value is 3.

To understand more check the link:

<https://pubs.vmware.com/vsphere-55/index.jsp?topic=%2Fcom.vmware.vsphere.storage.doc%2FGUID-C8E919D0-9D80-4AE1-826B-D180632775F3.html>

Question 8

Refer to the Exhibit.



Adapter	Type	Status	Identifier	Targets	Devices	Paths
vmhba34	Block SCSI	Unknown		0	0	0
QLogic NetXtreme II iSCSI Adapter						
vmhba32	iSCSI	Unbound	bnx2i-001018ed2ab0(iqn.1998...	0	0	0
vmhba33	iSCSI	Unbound	bnx2i-001018ed2ab2(iqn.1998...	0	0	0
iSCSI Software Adapter						
vmhba39	iSCSI	Online	iqn.1995-05.com.broadcom.vs...	1	1	1

The Exhibit shows the status of vmhba32 and vmhba33 as Unbound.

What is a likely reason for this status?

- ☒ The host bus adapter is not associated with a vmknfc.
- ☐ The Dynamic Target Discovery was not configured.
- ☐ The Static Target Discovery was not configured.
- ☐ The host bus adapter is not associated with a vmnic.

Explanation:

Explanation:

For sharing a volume across hosts, a VMFS volume is bound to its underlying block device storage. When a low level block copy is performed to copy or move the VMFS volume, the copied volume will be unbound.

When using Broadcom network interfaces with Hardware iSCSI Offload capabilities such as BCM5709 and BCM57711, each of the network interfaces shows up as individual storage adapters in the Configuration tab of an ESX host and each adapter also has an individual IQN. However, when you configure the adapter for iSCSI, you observe these symptoms:

You are unable to add any discovery addresses.

When observing the Dynamic or Static Discovery tabs in the iSCSI Initiator Properties window, you see the message:

The host bus adapter is not associated with a vmknfc. To configure targets the adapter should be associated with a vmknfc. Refer to the VMware documentation to associate the adapter with a vmknfc.

vSphere 6.0 Documentation Center - VMware

www.vmware.com/a/info/?id=1320

Question 9

Which two tasks does the Pluggable Storage Architecture (PSA) perform? (Choose two.)

- ☒ Handles I/O queueing to the logical devices.
- ☐ Handles physical path discovery, but is not involved in the removal.
- ☒ Handles physical path discovery and removal.
- ☐ Handles I/O queueing to FC storage HBAs.

Explanation:

Explanation :

the PSA performs the following tasks:

- ? Loads and unloads multipathing plug-ins.
- ? Hides virtual machine specifics from a particular plug-in.
- ? Routes I/O requests for a specific logical device to the MPP managing that device.
- ? Handles I/O queueing to the logical devices.
- ? Implements logical device bandwidth sharing between virtual machines.
- ? Handles I/O queueing to the physical storage HBAs.
- ? Handles physical path discovery and removal.
- ? Provides logical device and physical path I/O statistics.

Reference: <https://pubs.vmware.com/vsphere-55/index.jsp?topic=%2Fcom.vmware.vsphere.storage.doc%2FGUID-C1C4A725-8BE4-4875-919E-693812961366.html>

Question 10

Refer to the Exhibit.

```
6:16:06pm up 74 days 19:37, 470 worlds, 3 VMs, 4 vCPUs; CPU load average: 0.13, 0.10, 0.09
```

ADAPTR PATH	NPTH	CMDS/s	READS/s	WRITES/s	MBREAD/s	MBWRTN/s	DAVG/cmd	KAVG/cmd	GAVG/cmd	QAVG/cmd
vmhba0 -	1	6.68	0.00	0.00	0.00	0.00	0.25	0.03	0.28	0.01
vmhba1 -	1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
vmhba32 -	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
vmhba33 -	2	184.06	85.35	97.27	5.31	4.88	160.91	0.01	160.92	0.00

An administrator is using the `esxtop` command to troubleshoot storage performance issues on a virtual machine. The `esxtop` capture is shown in the Exhibit.



Based on the exhibit, which two statements are true? (Choose two.)

- ☒ The iSCSI device is experiencing high latency.
- ☐ The ESXi kernel is experiencing high latency.
- ☒ The Guest OS is experiencing high latency and response time.
- ☐ The NFS device is experiencing high latency.

Explanation:

Explanation:

Using `esxtop` to identify storage performance issues for ESX / ESXi (multiple versions) (1008205)

Note: In ESXi 5.x and later, you may see messages indicating that performance has deteriorated. For more information, see Storage device performance deteriorated (2007236).

Question 11

An administrator wants to allow users to login to the vSphere Web Client using the Use Windows session authentication check box for faster authentication.



Which three requirements must be met for this feature to be available and functional? (Choose three.)

- ☐ Install the vSphere Web Client Integration browser plug-in on the vCenter Server and Platform Services Controller machines.
- ☒ Install the vSphere Web Client Integration browser plug-in on each workstation from where a user will sign in.
- ☒ The users must be signed into Windows using Active Directory user accounts.
- ☒ The administrator must create a valid Identity Source in Single Sign-On for the users domain.
- ☐ The administrator must create a valid Single Sign-On Identity Source using Integrated Windows Authentication.

Explanation:

Explanation:

Explanation:

Procedure

- 1 Navigate to the vSphere Web Client login page.
- 2 If the Use Windows session authentication check box is not available, click Download the Client Integration Plug-in at the bottom of the login page.
- 3 If the browser blocks the installation by issuing certificate errors or by running a pop-up blocker, follow the Help instructions for your browser to resolve the problem.
- 4 Close other browsers if you are prompted to do so.

After installation, the plug-in is available for all browsers.

- 5 Exit and restart your browser.


After the restart, you can select the Use Windows session authentication check box.

Reference:

<https://pubs.vmware.com/vsphere-55/index.jsp?topic=%2Fcom.vmware.vsphere.security.doc%2FGUID-C7BF3BD5-76E6-406F-8CF1-309D31F1DEBA.html>

Question 12

Refer to the Exhibit.



A vSphere 6.x DRS cluster is configured as shown in the Exhibit.

Based on the exhibit, which statement is true?

- ☐ A virtual machine can be powered on in the Web Resource Pool with a 3 GHz CPU Reservation.
- ☐ A virtual machine can be powered on in the Web Resource Pool with a 4 GHz CPU Reservation.
- ☒ A virtual machine can be powered on in the DB Resource Pool with a 3 GHz CPU Reservation.
- ☐ A virtual machine can be powered on in the DB Resource Pool with a 4 GHz CPU Reservation.

Explanation:

Explanation:

Scaling out (not up) your VMs ie: Don't have one 8 vCPU SQL DB server, use 4 smaller 2vCPU VMs

If there is resource contention, RP-QA receives 4GHz and 2GB of memory, and RP-Marketing 2GHz and 1GB.

Otherwise, they can receive more than this allotment. Those resources are then available to the virtual machines in the respective resource pools.

So now you know better than to use reservations to solve CPU contention.

<https://pubs.vmware.com/vsphere-6.0>

Question 13

An administrator is upgrading vCenter Server and sees this error:

◆◆

The DB User entered does not have the required permissions needed to install and configure vCenter Server with the selected DB. Please correct the following error(s): %s

◆◆

Which two statements explain this error? (Choose two.)

- ☒ The database is set to an unsupported compatibility mode.
- ☒ The permissions for the database are incorrect.
- ☐ The permissions for vCenter Server are incorrect.
- ☐ The database server service has stopped.

Explanation:

Explanation:

Cause

This issue occurs when the vCenter Server instance being upgraded is connected to a SQL server which was upgraded and the vCenter Server database compatibility level is not changed to match the SQL server version.

Resolution

To resolve this issue, change the database compatibility level to match the SQL Servers level

◆

To change the database compatibility level:

Right-click the vCenter Server database and click Properties.

Click Options.

Change the compatibility level that matches the SQL server version.

Using the vSphere Command Line Interface (CLI), you can upgrade ESXi 5.x hosts to version 6.0.

For more information about upgrading ESXi 5.x hosts using esxcli commands, see the Upgrading Hosts by Using esxcli Commands

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2006904

Question 14

Which two scenarios could cause Storage DRS to be disabled on a Virtual Disk (VMDK)? (Choose two.)

- ☒ The VMDK is an independent disk.
- ☒ The virtual machine has vSphere Fault Tolerance enabled.
- ☐ The VMDK is hosted on NFS storage.
- ☐ The virtual machine has a CD-ROM/ISO image connected.

Explanation:

Explanation:

Storage DRS is Disabled on a Virtual Disk

Even when Storage DRS is enabled for a datastore cluster, it might be disabled on some virtual disks in the datastore cluster.

Problem

You have enabled Storage DRS for a datastore cluster, but Storage DRS is disabled on one or more virtual machine disks in the datastore cluster.

Cause

The following scenarios can cause Storage DRS to be disabled on a virtual disk.

- ? A virtual machine's swap file is host-local (the swap file is stored in a specified datastore that is on the host). The swap file cannot be relocated and Storage DRS is disabled for the swap file disk.
- ? A certain location is specified for a virtual machine's .vmx swap file. The swap file cannot be relocated and Storage DRS is disabled on the .vmx swap file disk.
- ? The relocate or Storage vMotion operation is currently disabled for the virtual machine in vCenter Server (for example, because other vCenter Server operations are in progress on the virtual machine). Storage DRS is disabled until the relocate or Storage vMotion operation is re-enabled in vCenter Server.
- ? The home disk of a virtual machine is protected by vSphere HA and relocating it will cause loss of vSphere HA protection.
- ? The disk is a CD-ROM/ISO file.
- ? If the disk is an independent disk, Storage DRS is disabled, except in the case of relocation or clone placement.
- ? If the virtual machine has system files on a separate datastore from the home datastore (legacy), Storage DRS is disabled on the home disk. If you use Storage vMotion to manually migrate the home disk, the system files on different datastores will be all be located on the target datastore and Storage DRS will be enabled on the home disk.
- ? If the virtual machine has a disk whose base/redo files are spread across separate datastores (legacy), Storage DRS for the disk is disabled. If you use Storage vMotion to manually migrate the

disk, the files on different datastores will be all be located on the target datastore and Storage DRS will be enabled on the disk.

- ? The virtual machine has hidden disks (such as disks in previous snapshots, not in the current snapshot). This situation causes Storage DRS to be disabled on the virtual machine.
- ? The virtual machine is a template.
- ? The virtual machine is vSphere Fault Tolerance-enabled.
- ? The virtual machine is sharing files between its disks.
- ? The virtual machine is being Storage DRS-placed with manually specified datastores

Reference:

<https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.troubleshooting.doc%2FGUID-B749AAA0-7B67-4A79-BEDC-395DFEC9FC60.html>

Question 15

An administrator is upgrading an ESXi host from 5.5 to 6.0 and runs the following command:

◆◆

```
esxcli software vib list --rebooting-image
```

◆◆

What◆does this command show?

- ☒ VIBs active after a reboot.
- ☐ VIBs that require a reboot.
- ☐ VIBs that are in the boot image.
- ☐ ◆VIBs that are third-party.

Explanation:

Explanation:

Install vCLI or deploy the vSphere Management Assistant (vMA) virtual machine. See◆Getting Started with vSphere Command-Line Interfaces. For troubleshooting , run◆esxcli◆commands in the ESXi Shell.

Procedure

- 1 Enter one of the following commands.

Option	Description
For VIBs	<code>esxcli --server=server_namesoftware vib list --rebooting-image</code>
For Profiles	<code>esxcli --server=server_namesoftware profile get --rebooting-image</code>

- 2 Review the output that is returned.
The output displays information for the ESXi image that will become active after the next reboot. If the pending-reboot image has not been created, the output returns nothing.

Reference: <https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.upgrade.doc%2FGUID-F0FE86AB-70B1-4E7E-A209-E5C5E0A4E57A.html>

Question 16

An administrator runs the command `esxcli storage core device list` and sees the following output:

```
mpx.vmhba1:C0:T0:L0 Display Name: RAID 5 (mpx.vmhba1:C0:T0:L0) Has Settable Display Name: false Size: 40960 Device Type: Direct-Access Multipath Plugin: NMP Devfs Path: /vmfs/devices/disks/mpx.vmhba1:C0:T0:L0 Status: off Is Local: true
```



What can be determined by this output?

- ☐ The device is being used for vFlash Read Cache.
- ☒ The device is in a Permanent Device Loss (PDL) state.
- ☐ The device is a local Solid State Device (SSD).
- ☐ The device is in an All Paths Down (APD) state.

Explanation:

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2004684

Question 17

Refer to the Exhibit.

Rule-Set 1

Select rules to create your VM storage policy.

- The VM storage policy will match datastores that satisfy any of the rule sets.
- A rule set will match datastores that satisfy all of the selected rules.

Rules based on common capabilities

<Add capability>

Rules based on vendor-specific capabilities

VSAN

Number of failures to tolerate

Number of disk stripes per object

<Add capability>

Rules based on tags

Add tag-based rule...

Storage Consumption Model

A virtual disk with size 100 GB would consume:

Storage space
300.00 GB
Initially reserved storage space
0.00 B
Reserved flash space
0.00 B

A Storage Policy for a Virtual SAN is set to the default policy, as shown in the Exhibit.

Which change would reduce the storage consumption by one third?

- ☒ Number of failures to tolerate = 1
- ☐ Number of disk stripes per object = 2
- ☐ Number of failures to tolerate = 3
- ☐ Number of disk stripes per object = 1

Explanation:

Explanation:

The new policy is added to the list.

What to do next:

Apply this policy to a virtual machine and its virtual disks. Virtual SAN will place the virtual machine objects in accordance with the requirements specified in the policy. For information about using and editing storage policies, see About Virtual Machine Storage Policies.

Also Check:

Policies Rules: Defines the number of host, disk, or network failures a virtual machine object can tolerate. For n failures tolerated, n+1 copies of the virtual machine object are created and 2n+1 hosts with storage are required.

Default value is 1. Maximum value is 3.

To understand more check the link:

<https://pubs.vmware.com/vsphere-55/index.jsp?topic=%2Fcom.vmware.vsphere.storage.doc%2FGUID-C8E919D0-9D80-4AE1-826B-D180632775F3.html>

Question 18

What are two limitations of Link Aggregation Control Protocol (LACP) on a vSphere Distributed Switch? (Choose two.)

- ☐ IP Hash load balancing is not a supported Teaming Policy.
- ☒ Software iSCSI multipathing is not compatible.
- ☐ Link Status Network failover detection must be disabled.
- ☒ It does not support configuration through Host Profiles.

Explanation:

Explanation:

LACP Limitations on a vSphere Distributed Switch Link Aggregation Control Protocol (LACP) on a vSphere distributed switch allows network devices to negotiate automatic bundling of links by sending LACP packets to a peer. However, there are some limitations when using LACP with a vSphere distributed switch.

LACP only works with IP Hash load balancing and Link Status Network failover detection.

LACP is not compatible with iSCSI software multipathing.

vSphere only supports one LACP group per distributed switch, and only one LACP group per host.

LACP settings do not exist in host profiles.

LACP between two nested ESXi hosts is not possible.

LACP does not work with port mirroring.

Reference:

<https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.networking.doc%2FGUID-3FDE1E96-9217-4FE6-8B76-6E3A64766828.html>

Question 19

Which secondary Private VLAN (PVLAN) type can communicate and send packets to an Isolated PVLAN?

- ☐ Community
- ☐ Isolated
- ☒ Promiscuous
- ☐ Primary

Explanation:

Explanation:

Private VLANs Private VLANs are used to solve VLAN ID limitations and waste of IP addresses for certain network setups.

A private VLAN is identified by its primary VLAN ID. A primary VLAN ID can have multiple secondary VLAN IDs associated with it. Primary VLANs are Promiscuous, so that ports on a private VLAN can communicate with ports configured as the primary VLAN. Ports on a secondary VLAN can be either Isolated, communicating only with promiscuous ports, or Community, communicating with both promiscuous ports and other ports on the same secondary VLAN.

Reference:

<https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.networking.doc%2FGUID-A9287D46-FDE0-4D64-9348-3905FEAC7FAE.html>

Question 20

An administrator is assigning a user the Content Library administrator role. The user will only be creating the library for a single vCenter Server.

◆◆

What is the lowest level of the permission heirarchy that this role can be granted to the user and still allow them to create a Content Library?

- ☒ Global
- ☐ Datacenter Folder
- ☐ Virtual Center
- ☐ Datacenter

Question 21

Which two options are available in the Virtual Machine Component Protection (VMCP) setting for Datastore with All Paths Down (APD)? (Choose two.)

- ☒ Issue Events
- ☒ Power off and restart virtual machines
- ☐ Reset virtual machines
- ☐ Leave Powered On

Explanation:

Explanation:

PDL failures

A virtual machine is automatically failed over to a new host unless you have configured VMCP only to Issue events.

APD events

The response to APD events is more complex and accordingly the configuration is more fine-grained.

After the user-configured Delay for VM failover for APD period has elapsed, the action taken depends on the policy you selected. An event will be issued and the virtual machine is restarted conservatively or aggressively. The conservative approach does not terminate the virtual machine if the success of the failover is unknown, for example in a network partition. The aggressive approach does terminate the virtual machine under these conditions. Neither approach terminates the virtual machine if there are insufficient resources in the cluster for the failover to succeed.

If APD recovers before the user-configured Delay for VM failover for APD period has elapsed, you can choose to reset the affected virtual machines, which recovers the guest applications that were impacted by the IO failures.

Reference: <https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.avail.doc%2FGUID-F01F7EB8-FF9D-45E2-A093-5F56A788D027.html>

Question 22

Which two settings are required for Virtual Machine Component Protection (VMCP) to protect from All Paths Down (APD) and Permanent Device Loss (PDL)? (Choose two.)

- ☒ Host Monitoring
- ☒ VM Restart Priority
- ☐ Virtual machine Monitoring
- ☐ Response for Host Isolation

Explanation:

Explanation:

If either the Host Monitoring or VM Restart Priority settings are disabled, VMCP cannot perform virtual machine restarts. Storage health can still be monitored and events can be issued,

Reference:

<https://pubs.vmware.com/vsphere60/index.jsp?topic=%2Fcom.vmware.vsphere.avail.doc%2FGUID-F01F7EB8-FF9D-45E2-A093-5F56A788D027.html>

Question 23

An administrator needs to migrate a legacy physical application server to a virtual machine within a vSphere 6 cluster. As part of the conversion, the administrator must reduce the size of the virtual disks.

What action should the administrator take to create a virtual machine with smaller virtual disks than the original physical server?

- ☐ Shut down the physical server and use VMware Converter cold cloning with volume-based cloning at the disk level.
- ☐ Use VMware Converter hot cloning with volume-based cloning at the block level.
- ☐ Shut down the physical server and use VMware Converter cold cloning with volume-based cloning at the file level.
- ☒ Use VMware Converter hot cloning with volume-based cloning at the file level.

Explanation:

Explanation: When reducing the disk volume size, the cloning operation switches from a Block level clone to a File level clone operation, meaning each individual file on the filesystem is copied over to the virtual machines hard drive rather than a block by block image of the drive.

◆

This can cause issues if the filesystem on the source machine is not in good health (for example, if it is severely fragmented, the physical drive has unrecorded bad sectors, or system files are damaged or missing in the index tables of the NTFS file system).

◆

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1014872

Question 24

After configuring a Virtual SAN cluster, an administrator using the vSphere Web Client notices that the Virtual SAN datastore is smaller than expected. The cluster contains:

Three ESXi hosts

Each host has one 10GB SDD

Each host has one 100GB HDD



Why would the Virtual SAN datastore show as 100GB instead of 300GB?

- ☒ There is a network problem with the Virtual SAN vmkernel ports.
- ☐ The Virtual SAN VASA provider is disabled.
- ☐ vSphere High Availability is enabled on the Virtual SAN cluster.
- ☐ The Virtual SAN cluster must be managed using the vSphere Web Client.

Explanation:

Explanation: Here we have three ESXi hosts (each with 100GB HDD) which makes the complete storage to 300 GB. It generally shows 100GB instead of 300GB because the three hosts are not able to connect properly due to network issues.

Question 25

An administrator is configuring virtual machines to use Worldwide Port Names (WWPNs) to access the storage.



Which two conditions are required? (Choose two.)

- ☒ The switches in the fabric must be N-Port ID Virtualization aware.
- ☒ The virtual machines must be using passthrough Raw Disk Mapping (RDMp).
- ☐ The virtual machines must be using Virtual Machine Disk (VMDK).
- ☐ The switches in the fabric must be Storage I/O Control aware.

Explanation:

Explanation: N-Port ID Virtualization

N-Port ID Virtualization (NPIV) is an ANSI T11 standard that describes how a single Fibre Channel HBA port can register with the fabric using several worldwide port names (WWPNs). This allows a fabric-attached N-port to claim multiple fabric addresses. Each address appears as a unique entity on the Fibre Channel fabric

How NPIV-Based LUN Access Works

NPIV enables a single FC HBA port to register several unique WWNs with the fabric, each of which can be assigned to an individual virtual machine. SAN objects, such as switches, HBAs, storage devices, or virtual machines can be assigned World Wide Name (WWN) identifiers. WWNs uniquely identify such objects in the Fibre Channel fabric. When virtual machines have WWN assignments, they use them for all RDM traffic, so the LUNs pointed to by any of the RDMs on the virtual machine must not be masked against its WWNs. When virtual machines do not have WWN assignments, they access storage LUNs with the WWNs of their host's physical HBAs. By using NPIV, however, a SAN

administrator can monitor and route storage access on a per virtual machine basis. The following section describes how this works.

Reference: https://pubs.vmware.com/vsphere-4-esx-vcenter/index.jsp#com.vmware.vsphere.config_fc.doc_40/esx_san_config/managing_san_systems/_n-port_id_virtualization.html

Question 26

An administrator decides to change the root password for an ESXi 6.x host to comply with the company's security policies.

What are two ways that this can be accomplished? (Choose two.)

- ☒ Use the Direct Console User Interface to change the password.
- ☒ Use the passwd command in the ESXi Shell.
- ☐ Use the password command in the ESXi Shell.
- ☐ Use the vSphere client to update local users.

Explanation:

Explanation: Limit ESXi Access

By default, the ESXi Shell and SSH services are not running and only the root user can log in to the Direct Console User Interface (DCUI). If you decide to enable ESXi or SSH access, you can set timeouts to limit the risk of unauthorized access.

ESXi enforces password requirements for direct access from the Direct Console User Interface, the ESXi Shell, SSH, or the vSphere Client.

Reference: <https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-security-guide.pdf>

Question 27

An administrator wants to power on a virtual machine (VM) while connected to an ESXi host using SSH. The VM has the following Name and ID:

VM Name = SQL001

VMID = 12345

◆

Which command would successfully power on the virtual machine?

- ☒ vim-cmd vmtoolsd/power.on 12345
- ☐ vim-cmd vmtoolsd/power.on SQL001
- ☐ vmware-vim-cmd vmtoolsd/power.on 12345
- ☐ vmware-vim-cmd vmtoolsd/power.on SQL001

Explanation:

Explanation:

ESXi 4.x, 5.x and 6.0

To power on a virtual machine from the command line:

1. List the inventory ID of the virtual machine with the command:

```
vim-cmd vmsvc/getallvms |grep <vm name>
```

Note: The first column of the output shows the **vmid**.

2. Check the power state of the virtual machine with the command:

```
vim-cmd vmsvc/power.getstate <vmid>
```

3. Power-on the virtual machine with the command:

```
vim-cmd vmsvc/power.on <vmid>
```

Reference:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1038043

Question 28

What is the minimum Virtual Hardware version required for vFlash Read Cache?

- ☐ Version 8
- ☐ Version 9
- ☒ Version 10
- ☐ Version 11

Explanation:

Explanation:

Virtual Machines vSphere Flash Read Cache is available only to virtual machines with hardware version 10 (VMX-10). All legacy virtual machine hardware must be upgraded to hardware version 10 to use Virtual Flash Read Cache.

Reference:

https://www.vmware.com/files/pdf/vSphere_55_Flash_Read_Cache_Whats_New_WP.pdf

Question 29

The command `esxcli network vm list` displays four virtual machines connected to the Production vSwitch. Within the vSphere Web Client, five virtual machines are seen.

◆◆

What explains this behavior?

- ☒ The fifth virtual machine is currently powered off.
- ☐ The fifth virtual machine has two vnics.
- ☐ The fifth virtual machine has an invalid IP address.
- ☐ The fifth virtual machine has an invalid MAC address.

Explanation:

Explanation: You can use vCLI network commands to view network statistics and troubleshoot your networking setup. The nested hierarchy of commands allows you to drill down to potential trouble spots.

1. List all virtual machine networks on a host.


`esxcli network vm list` The command returns for each virtual machine the World ID, name, number of ports, and networks, as in the following example.

Reference:

https://pubs.vmware.com/vsphere51/index.jsp?topic=%2Fcom.vmware.vcli.examples.doc%2Fcli_manage_networks.11.4.html

Question 30

Refer to the Exhibit.



An administrator reviews the Health of a virtual machine, as shown in the Exhibit.



Based on the exhibit, which three metrics can be used to determine the virtual machine's Workload characteristics? (Choose three.)

- ☒ CPU
- ☒ Memory
- ☒ Network IO
- ☐ Threads
- ☐ vNUMA Stats

Explanation:

Explanation:

Virtual Machine Alert Definitions

The vCenter adapter provides alert definitions that are generated on the virtual machines in your environment.

Health/Symptom-based

These alerts have the following impact and criticality information.

Check the Link for detailed info: <http://pubs.vmware.com/vrealizeoperationsmanager-61/index.jsp?topic=%2Fcom.vmware.vcom.core.doc%2FGUID-746FD64E-3380-44A6-A154-0BC63B4624F0.html>

Question 31

An administrator is concerned about possible vCPU over-commitment for an ESXi 6.x host.

Which two Performance Counters should be reviewed in the vSphere Web Client Performance Charts to confirm if there is contention on the host? (Choose two.)

- ☐ Wait
- ☒ Ready
- ☐ Core Utilization
- ☒ Co-Stop

Explanation:

Explanation:

To check vCPU over commitment for an Esxi 6.0 host >>>>performance Tab>>>>Graph etc.

Ready - Amount of time the virtual machine was ready to run, waiting in a queue to be scheduled.

Co-Stop - Amount of time a SMP virtual machine was ready to run, but incurred delay due to co-vCPU scheduling contention.

These performance metrics can be reviewed using the Performance tab in the vSphere Client or using the esxtop or resxtop command-line utilities. Choose the most appropriate method for your environment. Make note of the four metrics displayed. Each is measured in milliseconds

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1017926

Question 32

An administrator has created a resource pool named Marketing HTTP with a Memory Limit of 24 GB and a CPU Limit of 10,000 MHz.

◆

The Marketing HTTP resource pool contains three virtual machines:◆

Mktg-SQL has a memory reservation of 16 GB.

Mktg-App has a memory reservation of 6 GB.

Mktg-Web has a memory reservation of 4 GB.

◆

What would happen if all three virtual machines are powered on?

- ☐ All three virtual machines can power on, but will have memory contention.
- ☐ All three virtual machines can power on without memory contention.
- ☒ Only two of the three virtual machines can power on.
- ☐ Only one of the virtual machines can power on.

Explanation:

Explanation:

Applying a limit on a vCPU and memory will slow your VM down no matter what because of distribution of resources. Even if there are no other VMs running on that 4 socket quad core host.

Powering on all the virtual machine will enable 100% resources to be used by VM hence, only few required VM's will be powered "On" which requires the maximum utilization from the available resources.

Example: When a single (or two) VM has a limit of 5000MHz and 16GB, 6 GB Memory = 22gb utilization will be high and is the only VM running on a host then it will run at full speed as it will be constantly rescheduled for 5000MHz and maximum memory (22-24 GB)

Check:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2032906

Or Resource Pool Metrics:

<https://pubs.vmware.com/.../GUID-6F436275-169F-4DE5-B393-3B771623C0F8.htm>.

Question 33

What three supported methods can be used to upgrade a host from ESXi 5.x to ESXi 6.x? (Choose three.)

- ☒ vSphere Update Manager
- ☐ vihostupdate
- ☒ esxcli
- ☒ vSphere Auto Deploy
- ☐ esxupdate

Explanation:

Explanation:

Upgrade Options for ESXi 6.0 VMware provides several ways to upgrade ESXi 5.x hosts to ESXi 6.0 hosts.

The details and level of support for an upgrade to ESXi 6.0 depend on the host to be upgraded and the upgrade method that you use. Verify support for the upgrade path from your current version of ESXi to the version to which you are upgrading. See VMware Product Interoperability Matrixes at http://www.vmware.com/resources/compatibility/sim/interop_matrix.php.

You can upgrade a ESXi 5.x host, asynchronously released driver or other third-party customizations, interactive upgrade from CD or DVD, scripted upgrade, or upgrade with vSphere Update Manager. When you upgrade a ESXi 5.x host that has custom VIBs to version 6.0, the custom VIBs are migrated. See Upgrading Hosts That Have Third-Party Custom VIBs.

Methods supported for direct upgrade to ESXi 6.0 are:

- ? vSphere Update Manager.
- ? Interactive upgrade from CD, DVD, or USB drive.
- ? Scripted upgrade.

? vSphere Auto Deploy. If the ESXi 5.x host was deployed by using vSphere Auto Deploy, you can use vSphere Auto Deploy to reprovision the host with an ESXi 6.0 image.

? The esxcli command.

Reference: <https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.upgrade.doc%2FGUID-FE668788-1F32-4CB2-845C-5547DD59EB48.html>

Question 34

An administrator notices that the time on an ESXi 6.x host is incorrect.

◆◆

Which two actions should the administrator take to correct this issue? (Choose two.)

- ☒ Modify the time for the host using the vSphere client.
- ☒ Correct the NTP settings in the /etc/ntp.conf file.
- ☐ Configure NTP from the Direct Console User Interface.
- ☐ Use the vicfg-ntp command from the vSphere Management Appliance.

Explanation:

Explanation:

Resolution

To configure NTP on ESX/ESXi 4.1 and ESXi 5.x hosts using the vSphere Client:

1. Connect to the ESX/ESXi host using the vSphere Client.
2. Select a host in the inventory.
3. Click the Configuration tab.
4. Click Time Configuration.
5. Click Properties.
6. Click Options.
7. Click NTP Settings.
8. Click Add.
9. Enter the NTP Server name. For example, pool.ntp.org.

Note: When entering the multiple NTP Server names, use a comma (,) followed by a space () between the entries.

10. Click OK.
11. Click the General tab.

12. Click Start automatically under Startup Policy.

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2012069

Question 35

When attempting to remove a host from a vSphere Distributed Switch (vDS), an administrator observes the error message:

The resource '16' is in use

What are two reasons why this error would be displayed? (Choose two.)

- ☒ VMkernel network adapters on the vDS are in use.
- ☒ Virtual machine network adapters are connected to the vDS.
- ☐ Network I/O Control has been configured on the vDS.
- ☐ There is active network traffic on the vDS.

Explanation:

Explanation: Cause

You cannot remove the host from the distributed switch or delete the host proxy switch because of the following reasons.

There are VMkernel adapters on the switch that are in use.

There are virtual machine network adapters connected to the switch.

Reference: <https://pubs.vmware.com/vsphere-55/index.jsp?topic=%2Fcom.vmware.vsphere.troubleshooting.doc%2FGUID-038AC93F-D710-48ED-8E3B-258A23FB2930.html>

Question 36

How does vSphere High Availability calculate the memory slot size of a virtual machine?

- ☒ Virtual machine memory reservation + overhead of largest virtual machine
- ☐ Virtual machine memory reservation - overhead of largest virtual machine
- ☐ Virtual machine memory reservation + overhead of smallest virtual machine
- ☐ Virtual machine memory reservation - overhead of smallest virtual machine

Explanation:

Explanation:

Virtual machine memory reservation + overhead of largest virtual machine

Metrics Example: Suppose if :

Memory (2048MB)

Memory Overhead (110.63MB)

The first one, ♦Memory, is an easy one. This is the amount of memory you provisioned your VM with, in this case 2048MB. The second field is ♦Memory Overhead. Memory Overhead is the amount of memory the VMkernel thinks it will need to run the virtualized workload, in this case 110.63MB. This typically would include things like page tables, frame buffers etc.

<https://www.vmware.com/files/pdf/.../VMware-PerfBest-Practices-vSphere6-0.pdf>

Question 37

Which command shows the Physical Uplink status for a vmnic?

- ☐ esxcli network ip get ♦
- ☒ esxcli network nic list
- ☐ esxcli network vmnic list
- ☐ esxcli network ifconfig get

Explanation:

Explanation:

Command:

network nic list

Description:

This command will list the Physical NICs currently installed and loaded on the system.

Reference: https://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vcli.ref.doc_50%2Fesxcli_network.html

Question 38

Which two vCenter ♦Server ♦services are migrated automatically as part of an upgrade from a Distributed vCenter Server running 5.x? (Choose two.)

- ☐ vCenter Single Sign-on Service
- ☒ vSphere Web Client
- ☒ vSphere Inventory Service
- ☐ Storage Policy Based Management

Explanation:

Explanation:

vCenter Server 5.x for Windows services that are migrated to become part of the vCenter Server group of services during the upgrade process include:

Inventory Services

vSphere Web Client

vSphere Auto Deploy

vSphere Syslog Collector

vSphere ESXi Dump Collector

vCenter Server and vCenter Single Sign-On are the only services that are not migrated. vCenter Single Sign-On instances are upgraded in place to become part of an external Platform Services Controller if they are deployed on a system other than the system where the vCenter Server resides.

Reference:

<https://pubs.vmware.com/vsphere60/index.jsp?topic=%2Fcom.vmware.vsphere.upgrade.doc%2FGUID-4BFB12D8-9FCA-4AB1-A44F-2986966F0AD5.html>

Question 39

An administrator has migrated a vCenter Server Appliance from version 5.5 to version 6.x. During the migration, the administrator selected DHCP for the appliance and obtained a hostname from the DHCP server. The administrator adjusts the hostname after the migration and uses a static IP and hostname.

◆◆

What should the administrator do immediately after this change to prevent service failures?

- ☒ Regenerate the SSL certificates.
- ☐ Re-register components to Single Sign-On.
- ☐ Update the /etc/hosts file.
- ☐ Execute the command `services.restart vmware-vpxd`.

Explanation:

Explanation:

This is a known issue affecting vCenter Server 6.0.

Currently, there is no resolution.

To work around this issue, regenerate the SSL certificates on the vCenter Server 5.x before upgrading to 6.0.

To regenerate the SSL certificates on the vCenter Server 5.x before upgrading to 6.0,

LinkKB: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2110943

Question 40

What are two ways to view the DNS settings for an ESXi 6.x host? (Choose two.)

- ☒ Use the vicfg-dns command from the vSphere Management Appliance.
- ☒ View the /etc/resolv.conf file on the ESXi host.
- ☐ Use vicfg-dns command on the ESXi host.
- ☐ View the /etc/dns.conf file on the ESXi host.

Explanation:

Explanation: To verify that the DNS information in /etc/resolv.conf is correct for your environment run the command:

```
[root@server root]# cat /etc/resolv.conf
```

Run vicfg-dns without command-specific options to display DNS properties for the specified server.

vicfg-dns <conn_options>The information includes the host name, domain name, DHCP setting (true or false), and DNS servers on the ESXi host.

Reference: https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vcli.examples.doc%2Fcli_manage_networks.11.8.html

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1003796

Question 41

Which three statements are true when restoring a Resource Pool Tree? (Choose three.)

- ☐ Distributed Resource Scheduler must be set to manual.
- ☒ Restoring a snapshot can only be done on the same cluster from which it was taken.
- ☒ No other resource pools can be present in the cluster.
- ☒ Restoring a resource pool tree must be done in the vSphere Web Client.
- ☐ Enabling Enhanced vMotion Compatibility on the cluster is required.

Explanation:

Explanation:

Restore a Resource Pool Tree in the vSphere Web Client You can restore a previously saved resource pool tree snapshot.

Prerequisites

vSphere DRS must be turned ON.

You can restore a snapshot only on the same cluster that it was taken.

No other resource pools are present in the cluster.

Reference:

<https://pubs.vmware.com/vsphere51/index.jsp?topic=%2Fcom.vmware.vsphere.resmgmt.doc%2FGUID-43B3A1EF-B7FF-421C-96FA-33FA230688BB.html>

Question 42

An administrator creates a custom ESXi firewall rule using an XML file, however the rules do not appear in the vSphere Web Client.



Which action should the administrator take to correct the problem?

- ☐ Load the new rules using esxcli network firewall reload.
- ☒ Load the new rules using esxcli network firewall refresh.
- ☐ Verify the entries in the XML file and then reboot the ESXi host.
- ☐ Remove the ESXi host from the inventory and add it back.

Explanation:

Explanation:

Revert the access permissions of the service.xml file to the read-only default by running the command:

```
# chmod 444 /etc/vmware/firewall/service.xml
```

Refresh the firewall rules for the changes to take effect by running the command:

```
# esxcli network firewall refresh
```

or

```
# localcli network firewall refresh
```

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2008226

Question 43

What will occur if the .nvram file is deleted from a powered off virtual machine?

- ☒ The .nvram file will get created the next time the virtual machine is powered on.
- ☐ Restoring the file from backup is needed to allow the virtual machine to power on.
- ☐ The virtual machine will fail to power on and enter an Orphaned state.
- ☐ The virtual machine will fail to power on and enter an Inaccessible state.

Explanation:

Explanation:

The .nvram file. This small file contains the BIOS that is used when the VM boots. It is similar to a physical server that has a BIOS chip that lets you set hardware configuration options. A VM also has a virtual BIOS that is contained in the NVRAM file. The BIOS can be accessed when a VM first starts up by pressing the F2 key. Whatever changes are made to the hardware configuration of the VM are then saved in the NVRAM file. This file is in binary format and if deleted it will be automatically re-created when a VM is powered on.

Reference:

<http://searchvmware.techtarget.com/tip/Understanding-the-files-that-make-up-a-VMware-virtual-machine>

Question 44

An administrator needs to recover disk space on a previously-used thin provisioned virtual disk. The volumes where the administrator needs to recover the disk blocks are on VAAI-compliant storage arrays.

◆◆

Which two actions should the administrator take accomplish this task? (Choose two.)

- ☐ Perform a Storage vMotion to another volume in order to force free space recovery to occur. This recreates the volume in a new location and recovers all unused space.
- ☒ Use VMware Converter to migrate the virtual machine to a new datastore. This will recreate the volumes and recover all unused space.
- ☐ Issue the vmkfstools -vmfs unmap command within the VMFS volume directory on the ESXi host console.
- ☒ Execute the esxcli storage vmfs unmap command.

Explanation:

Explanation:

To reclaim unused storage blocks on a VMFS datastore for a thin-provisioned device, run this command:

```
esxcli storage vmfs unmap --volume-label=volume_label|--volume-uuid=volume_uuid--reclaim-unit=number
```

The command takes these options:

`-l|--volume-label=volume_label`

The label of the VMFS volume to UNMAP. This is a mandatory argument. If you specify this argument, do not use `-u|--volume-uuid=volume_uuid`.

`-u|--volume-uuid=volume_uuid`

The UUID of the VMFS volume to UNMAP. This is a mandatory argument. If you specify this argument, do not use `-l|--volume-label=volume_label`.

`-n|--reclaim-unit=number`

The number of VMFS blocks to UNMAP per iteration. This is an optional argument. If it is not specified, the command uses a default value of 200.

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2057513

Question 45

Refer to the Exhibit.

-- Exhibit --



Review the Exhibit. An administrator has configured permissions for a group called VMGroup and a user named VMUser. A new Role has been created called PowerVM. ♦The group and role have these characteristics: ♦♦

PowerVM role can power on VMs

VMGroup granted PowerVM role on VMFolder

VMUser is a member of VMGroup

VMUser granted No Access on VMFolder

♦

Based on the exhibit, which statement best explains why VMUser is denied access to the VMFolder?

- ☒ The VMUser permission overrides the VMGroup permission.
- ☐ The No Access role overrides the PowerVM role.
- ☐ The VMGroup permission overrides the VMUser permission.
- ☐ The PowerVM role overrides the No Access role.

Explanation:

Explanation:

User Permissions Overriding Group Permissions:

This example illustrates how permissions assigned directly to an individual user override permissions assigned to a group that the user is a member of.

In this example, permissions are assigned to a user and to a group on the same object.

Role 1 can power on virtual machines.

Group A is granted Role 1 on VM Folder.

User 1 is granted No Access role on VM Folder.

User 1, who belongs to group A, logs on. The No Access role granted to User 1 on VM Folder overrides the group permission. User 1 has no access to VM Folder or VMs A and B.

Example 3: User Permissions Overriding Group Permissions

<https://pubs.vmware.com/vsphere-60>

The setting in vSphere is a property of the ESX(i) host, found within the configuration tab in the Virtual Machine Startup/Shutdown options section.

A VM startup sequence can be set with timing for each host.

The properties section will allow the startup options to be configured. It is a good idea to shorten their startup time if the hosts power up and are ready to go. This may not be the case with SAN or NAS systems in use, which may take more time to start up.

To power on a virtual machine from the command line:

List the inventory ID of the virtual machine with the command:

```
vim-cmd vmsvc/getallvms |grep <vm name>
```

Note: The first column of the output shows the vmid.

Check the power state of the virtual machine with the command:

```
vim-cmd vmsvc/power.getstate <vmid>
```

Power-on the virtual machine with the command:

```
vim-cmd vmsvc/power.on <vmid>
```

Question 46

An administrator is upgrading an ESXi 5.5 host to ESXi 6.x and gets the following error:

```
❖❖
```

```
MEMORY_SIZE❖
```

```
❖❖
```

What does this indicate?

- ☒ Insufficient memory on the ESXi host to complete the upgrade.
- ☐ Insufficient memory for Auto Deploy to complete the upgrade.
- ☐ Insufficient memory in vCenter Server to complete the upgrade.
- ☐ Insufficient memory for Update Manager to complete the upgrade.

Explanation:

Explanation:

A-) ESXi 6.0 requires the NX/XD bit to be enabled for the CPU in the BIOS.

ESXi requires a minimum of 4GB of physical RAM. It is recommended to provide at least 8 GB of RAM to run virtual machines in typical production environments.

To support 64-bit virtual machines, support for hardware virtualization (Intel VT-x or AMD RVI) must be enabled on x64 CPUs.

<http://pubs.vmware.com/vsphere60/index.jsp?topic=%2Fcom.vmware.vsphere.upgrade.doc%2FGUID-DEB8086A-306B-4239-BF76-E354679202FC.html>

Question 47

When troubleshooting a performance related issue, an administrator sees the following warning message on an ESXi Console:

Significant imbalance between NUMA nodes detected. Performance may be impacted.

Which action should the administrator take to correct this problem?

- ☒ Ensure that the RAM modules are evenly balanced between processor sockets in the physical server.
- ☐ Ensure that all memory banks associated with CPU Socket 0 in the physical server are completely filled.
- ☐ Ensure that the physical server has the maximum amount of RAM modules that it can support.
- ☐ Ensure that all RAM modules in the host are Error-correcting code (ECC) modules.

Explanation:

Explanation:

This message displays when memory is not evenly distributed between processor sockets.

Each CPU socket has its own RAM slots. To resolve this issue, evenly balance RAM modules between processor sockets. That is, ensure that each bank of RAM slots has the same amount of RAM.



For example, if CPU socket 1 has a total of 8 GB of RAM, ensure that CPU socket 2 also has 8 GB of RAM.

Ensure that memory on the nodes are balanced. For more information and assistance, contact your vendor.



Note: Damaged or faulty physical RAM modules can also cause this issue.

Reference:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1018754

Question 48

Users are reporting CPU related performance problems on the virtual machine Email-Prod throughout the day.



The resource settings◆for the VM and the ESXi host the VM is running on are shown◆below:



VM:

Email-Prod

vCPUs: 4

Memory: 96GB



ESXi host:◆

CPU: 2 x 8 Core Processors

Memory: 128GB

NUMA architecture. 2 Nodes



Which two options would alleviate the observed performance problem for Email-Prod? (Choose two.)

- ☐ Enable◆the advanced parameter Numa.PageMigEnable.
- ☒ Enable the advanced parameter Numa.AutoMemAffinity.
- ☒ Enable CPU affinity to separate 2 vCPUs for each NUMA node.
- ☐ Enable CPU affinity to bind all vCPUs to one NUMA node.

Explanation:

Explanation:

B-) Memory You can change the amount of RAM that a VM will use like this

memsize = "128"

If you don't assign a value for memory VMware will use the default-setting which is 32MB

memsize = "32"

Whenever you assign RAM manually make sure the value is a multiple of 4 - otherwise the VM will not start:

memsize = "255"

causes this error-message:

Memory size 255 not a multiple of 4◆

<https://pubs.vmware.com/vsphere-60/index.jsp#com.vmware.vsphere.hostclient.doc/GUID-99E66B59-53A8-42D0-A08F-17DBE7D4A864.html?resultof=%2522>

C-) This is expected behavior based on the current architecture of the scheduler. However, these points help to limit the impact of the issue :Lower the number of vCPUs.

From the example above, you would lower the number of vCPUs from 8 to 6 or less. If the virtual machines are sized such that they are a whole multiple or divisor of the NUMA node size, this helps with the number of virtual machines that you can power on. If 6 vCPU virtual machines are used, you can run up to at least 8 of those virtual machines (with 100% CPU utilization) without incurring substantial ready times.

<https://www.vmware.com/files/pdf/.../VMware-PerfBest-Practices-vSphere6-0.pdf>

Question 49

Which log file would you examine to identify an issue which occurred during the pre-upgrade phase of a vCenter Server upgrade process?

- ☒ vcdb_req.out
- ☐ vcdb_export.out
- ☐ vcdb_import.out
- ☐ vcdb_inplace.out

Explanation:

Explanation:

Example: Database Upgrade Locations

For pre-upgrade checks, review the %TEMP%\..\vcsUpgrade\vcdb_req.out file.


The vcdb_req.err file tracks any errors that were identified during the pre-upgrade phase.

Reference:

<https://pubs.vmware.com/vsphere60/index.jsp?topic=%2Fcom.vmware.vsphere.upgrade.doc%2FGUID-5EAC8B84-0A95-41EC-AAF4-6CBBB3A5152A.html>

Question 50

Refer to the Exhibit.



An administrator has configured a vSphere 6.x DRS cluster as shown in the Exhibit.

Based on the exhibit, which statement is true?

- ☒ A virtual machine can be powered on in the Test Resource Pool with a 6 GB Memory Reservation.
- ☐ A virtual machine can be powered on in the Dev Resource Pool with a 8 GB Memory Reservation.
- ☐ A virtual machine from both the Test Resource Pool and the Dev Resource Pool can be powered on with a 4 GB Memory Reservation.
- ☐ No more virtual machines can be powered on due to insufficient resources.

Explanation:

Explanation:

A virtual machine can be powered on in the Test Resource Pool with a 6 GB Memory Reservation because: Total is 8GB

1Gb used by test resource pool VM which is powered on and 2gb by Test dev pool VM powered on hence distribution is less and expandable memory quota is more/left.

To understand limits and theory check the information given below:

Memory Maximums

The ESXi host maximums represents the limits for ESXi host memory.

Table 3-2. ESXi Host Memory Maximums

Item Maximum

RAM per host 6 TB

12 TB is supported on specific OEM certified platform. See VMware Hardware Compatibility Limits for guidance on the platforms that support vSphere 6.0 with 12 TB of physical memory.

Number of swap files 1 per virtual machine

And,

Resource Pool

Resource pools per host 1600

Children per resource pool 1100

Resource pool tree depth 8

Additional 4 resource pools are used by system internals.


Resource pools per cluster 1600

Configuration Maximums - vSphere 6.0 - VMware

<https://www.vmware.com/pdf/vsphere6/r60/vsphere-60-configuration-maximums.pdf>

Question 51

Refer to the Exhibit.



An administrator has configured Distributed Resource Scheduler (DRS) groups and Affinity Rules as shown in the Exhibit.

◆Based on the exhibit, which two statements are true? (Choose two.)

- ☐ If ESXi-A and ESXi-D failed, VM-A and VM-B would not failover.
- ☒ A new conflicting affinity rule will be disabled by default.
- ☐ VM-B and VM-D can run on the same hosts.
- ☒ The administrator must disable Rule3 in order to enable Rule4.

Explanation:

Explanation:

B-) NOTE If you use the Specify Failover Hosts admission control policy and designate multiple failover hosts, DRS does not attempt to enforce VM-VM affinity rules for virtual machines that are running on failover hosts.

D-) vSphere HA and DRS Affinity Rules If you create a DRS affinity rule for your cluster, you can specify how vSphere HA applies that rule during a virtual machine failover. The two types of rules for which you can specify vSphere HA failover behavior are the following: n VM anti-affinity rules force specified virtual machines to remain apart during failover actions. n VM-Host affinity rules place specified virtual machines on a particular host or a member of a defined group of hosts during

failover actions. When you edit a DRS affinity rule, select the checkbox or checkboxes that enforce the desired failover behavior for vSphere HA. n HA must respect VM anti-affinity rules during failover -- if VMs with this rule would be placed together, the failover is aborted.

<https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-availability-guide.pdf>

Question 52

An administrator is troubleshooting a virtual machine performance issue using vRealize Operations.



Which two badges would help to identify possible resource contention concerns? (Choose two.)

- ☒ Health > Workload
- ☐ Health > Faults
- ☐ Risk > Time Remaining
- ☒ Risk > Stress

Explanation:

Explanation:

The **Workload** badge shows how hard an object is working. A higher workload score indicates that an object is doing more work. Obviously, you don't want objects out there doing zero work, as that is waste but, as the same time, you also don't want objects completely maxed out with a workload score of 100 either. **Workload** is an absolute measurement that calculates the demand for a resource divided by the capacity of an object. Resources might include CPU, memory, disk I/O, or network I/O. **vC Ops** will help you to balance workload across your resource objects effectively

Stress badge reports the stress that an object is under. Just as your stress level is related to your workload, so is the stress score in vC Ops. The stress score is based on long-term high workload (where the workload score is instantaneous workload). Stress in the virtual infrastructure might come from VMs that are undersized, hosts that are overloaded, or datastores with consistently high I/O latency. Stress is reported between 0 and 100 with 100 being very high stress and 0 being no stress.

Reference:

<http://blogs.vmware.com/management/2014/04/david-davis-on-vcenter-operations-post-8-understanding-vcenter-operations-badges.html>

Question 53

An administrator creates a Private VLAN with a Primary VLAN ID of 2. The administrator then creates three Private VLANs as follows:

Marketing
PVLAN ID. 4
PVLAN Type. Isolated
Accounting
PVLAN ID. 5
PVLAN Type. Community
Secretary
PVLAN ID. 17
PVLAN Type. Isolated

◆

◆

Users in the Accounting PVLAN are reporting problems communicating with servers in the Marketing PVLAN.

◆◆

Which two actions could the administrator take to resolve this problem? (Choose two.)

- ☒ Change the PVLAN type for the Accounting network to Promiscuous.
- ☒ Change the PVLAN ID for the Accounting network to 2.
- ☐ Change the PVLAN type for Marketing network to Promiscuous.
- ☐ Change the PVLAN ID for Accounting network to 4.

Explanation:

Explanation: Please refer to this link for understanding network connectivity issues in detail: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1004109

Question 54

Which two scenarios would cause a Fault Tolerance-enabled virtual machine to fail to power the Secondary virtual machine? (Choose two.)

- ☒ The host has entered a Network Partitioned state.
- ☒ vSphere High Availability (HA) is disabled on the host cluster.
- ☐ Enhanced vMotion Compatibility (EVC) is enabled on the host cluster.
- ☐ vSphere Distributed Power Management (DPM) is enabled on the host cluster.

Explanation:

Explanation:

A-) This issue occurs when the SSL certificate thumbprint presented to the master host is not what the master host is expecting. This is indicated by the thumbprint mismatch error in the `fdm.log` file of the master host https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2012649

In general terms, a second virtual machine is created to work in tandem with the virtual machine on which you have enabled Fault Tolerance. ♦ This virtual machine resides on a different host in the cluster and runs in virtual lockstep with the primary virtual machine. ♦ When a failure is detected, the second virtual machine takes the place of the first one with the least possible interruption of service. ♦ More specific information about how this is achieved can be found in the ♦ Protecting Mission-Critical Workloads with VMware Fault Tolerance ♦ whitepaper.

B-) You may need to complete this process for multiple clusters. VMware HA can be disabled only if there are no virtual machines with VMware Fault Tolerance (FT) enabled. If there are virtual machines with VMware FT enabled in the cluster you are disabling, turn off VMware FT before disabling VMware HA. The process of turning off VMware FT is described in ♦ Disabling or Turning Off VMware FT (1008026) ♦.

Note: Virtual Machine Monitoring will also be disabled when HA is disabled.

Question 55

A vSphere Web Client task fails while creating a VMFS datastore on a disk with these characteristics: ♦

The disk was ♦ formatted with a Master Boot Record (MBR) ♦ partition table

The disk was ♦ not ♦ erased

The disk is ♦ visible in ♦ the ♦ vSphere Web Client

What action needs to ♦ be ♦ performed ♦ to ♦ resolve ♦ the ♦ issue?

- ☒ Delete the partitions manually with partedUtil. ♦
- ☐ Create a VMFS3 file system first, then upgrade it.
- ☐ Create a VMFS5 file system with the command ♦ esxcli storage filesystem add
- ☐ Delete the data with the vmkfstools command.

Explanation:

Explanation: Deleting a partition

A single partition can be deleted from a partition table on a block disk device using the ♦ partedUtil ♦ command line utility.

To delete a partition, run this command:

```
partedUtil delete "/vmfs/devices/disks/DeviceName" ♦ PartitionNumber
```

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1036609# Deleting a partition

Question 56

When a Content Library is deleted, what happens to the files contained on the backing storage?

- ☐ The files will be unchanged.
- ☒ The files will be deleted.
- ☐ The files will remain and be marked as orphaned.
- ☐ The containing folder will be marked as orphaned.

Explanation:

Explanation:

Delete a Content Library

When you no longer need a content library, you can invoke the delete method on either the LocalLibrary or the SubscribedLibraryservice depending on the library type.

Procedure

1Access the SubscribedLibrary or the LocalLibrary service by using the vCloud Suite Endpoint.

2Retrieve the library ID you want to delete.

3Call the delete function on the library service and pass the library ID as argument.

All library items cached on the storage backing are removed asynchronously. If this operation fails, you must manually remove the content of the library.

<https://pubs.vmware.com/vsphere-60/index.jsp#com.vmware.vapi.progguide.doc/GUID-C3A2DD49-BF41-43CD-9AEC-D6F28A9F3AF4.html?resultof=%2522%>

Question 57

Which two supported tools can be used to upgrade virtual machine hardware? (Choose two.)

- ☒ vSphere Web Client
- ☒ vSphere Update Manager
- ☐ vmware-vmupgrade.exe
- ☐ esxcli vm hardware upgrade

Explanation:

Explanation:

Upgrade Virtual HardwareYou can upgrade the hardware version of virtual machines to the latest version of ESXi. For virtual machines that are running on ESXi 5.x, VMware recommends that you upgrade the virtual hardware to version 8.

Consider the following points:

? When you upgrade from virtual hardware version 4 to version 8 the upgrade is reversible if you take a virtual machine backup or snapshot before performing the upgrade.

? Upgraded virtual machines cannot be powered on by an ESX 2.x host, even if relocated to a VMFS2 datastore.

? To automate this process, consider using Update Manager for virtual machine upgrades. See the Installing and Administering VMware vSphere Update Manager documentation. Update Manager takes automatic snapshots before performing virtual machine upgrades. See Perform an Orchestrated Upgrade of Virtual Machines with vSphere Update Manager.

? When you upgrade virtual hardware, no downtime is required for vCenter Server or ESX/ESXi hosts. For virtual machines, the only significant downtime is the time to reboot the guest operating systems.

Reference: https://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vsphere.upgrade.doc_50%2FGUID-A45CBEE5-C4D2-484E-A699-A5A577244DE0.html

Question 58

What are two reasons why a local flash device would be unavailable for use with Virtual SAN?
(Choose two.)

- ☒ It has a VMFS datastore present.
- ☒ It is in use by the vFlash Read Cache feature.
- ☐ It is smaller than the minimum capacity required for Virtual SAN usage.
- ☐ It does not have any partitions created.

Explanation:

Explanation:

Local Flash Devices Are Unavailable for Use with Virtual SAN or Virtual Flash

A local flash device becomes unavailable for virtual flash resource or Virtual SAN configuration when it is formatted with VMFS or any other file system.

Problem

When you attempt to configure either Virtual SAN or virtual flash resource, a local flash disk does not appear on the list of disks to be used.

Cause

This problem might occur when a local flash intended for use with either feature has been already formatted with VMFS. Neither Virtual SAN nor virtual flash can share the flash disk with VMFS or any other file system.

Also, because virtual flash and Virtual SAN are mutually exclusive consumers of flash disks, both features cannot share the same flash disk. If the flash disk is already claimed by one feature, for example Virtual SAN, you are not be able to use it for another, such as virtual flash, unless you release the disk.

Solution

Use only unformatted flash disks for virtual flash resource and Virtual SAN configuration.

Avoid formatting the flash disks with VMFS during ESXi installation or Auto Deploy.

If the flash disk is already formatted with VMFS, remove the VMFS datastore. For information see the vSphere Storage documentation.

To use the flash disk as a virtual flash resource, do not claim this disk for Virtual SAN. If the disk is claimed by Virtual SAN, remove the disk from Virtual SAN. The flash disk is released from Virtual SAN and becomes available on the list of disks to use with virtual flash. For information about removing disks from Virtual SAN, see the Administering VMware Virtual SAN documentation.

If you intend to use the flash disk with Virtual SAN, do not use the disk for a virtual flash resource. If the flash disk is used as the virtual flash resource, remove the virtual flash configuration. The disk becomes available for Virtual SAN. See the vSphere Storage documentation.

Reference:

<https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.troubleshooting.doc%2FGUID-A17F5D39-C1F3-47B1-9645-5C247412F370.html>

Question 59

Which three logs can be reviewed to troubleshoot a vCenter Server upgrade failure? (Choose three.)

- ☒ vminst.log
- ☒ vim-vcs-msi.log
- ☒ pkgmgr.log
- ☐ vc-upgrade.log
- ☐ firstboot.log

Explanation:

Explanation:

Retrieve Installation Logs Manually You can retrieve the installation log files manually for examination.

Procedure

1 Navigate to the installation log file locations.

? %PROGRAMDATA%\VMware\vCenterServer\logs directory, usually

C:\ProgramData\VMware\vCenterServer\logs

? %TEMP% directory, usually C:\Users\username\AppData\Local\Temp

The files in the %TEMP% directory include vminst.log, pkgmgr.log, pkgmgr-comp-msi.log, and vim-vcs-msi.log.

2 Open the installation log files in a text editor for examination.

Reference:

<https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.install.doc%2FGUID-CF80C555-BFB1-4FF1-B4E6-6935D424838E.html>

Question 60

What is the name of the command line utility that checks for VMFS5 metadata corruption?

- ☐ vmkfstools --check
- ☒ voma
- ☐ vmfsanalyzer
- ☐ esxcli vmfs check

Explanation:

Explanation:

In ESXi 5.1, VMware introduced the vSphere On-disk Metadata Analyzer (VOMA) for performing VMFS file system metadata checks. This utility scans the VMFS volume metadata and highlights any inconsistencies to which you may be required to open a support request.

Reference:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2036767

Question 61

An administrator is unable to patch an ESXi 6.x host using VMware Update Manager



What is an alternative option for patching a host?

- ☒ Upload the offline bundle to a datastore and execute the command `esxcli software vib install -d` to apply it manually.
- ☐ Upload the vib to a datastore and execute the command `esxcli software vib install -d` to apply it manually.
- ☐ Upload the offline bundle to a datastore and execute the command `esxupdate install -v` to apply it manually.
- ☐ Upload the vib to a datastore and execute the command `esxupdate install -v` to apply it manually.

Explanation:

Explanation:

Update a Host with Individual VIBs You can update a host with VIBs stored in a software depot that is accessible through a URL or in an offline ZIP depot.

Important

If you are updating ESXi from a zip bundle in a VMware-supplied depot, either online from the VMware Web site or downloaded locally, VMware supports only the update method specified for VMware-supplied depots in the topic Upgrade or Update a Host with Image Profiles. The `esxcli software vib update` and `esxcli software vib install` commands are not supported for upgrade operations. See Differences Between vSphere Upgrades and Updates and Upgrade or Update a Host with Image Profiles

Reference:

<https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.upgrade.doc%2FGUID-1E773360-CB1C-4BC2-B2A4-B73AB5706FAF.html>

Question 62

An administrator has decided to create 10 ESXi 6.x hosts via Auto Deploy for a new Test/Dev cluster. The hosts are configured to obtain their networking configuration via DHCP.



Which Direct Console User Interface option should the administrator use to renew the DHCP lease for the hosts?

- ☐ Restore Network Settings
- ☐ Test Management Network
- ☒ Restart Management Network
- ☐ Network Restore Options

Explanation:

Explanation: Restart the Management Network Restarting the management network interface might be required to restore networking or to renew a DHCP lease.

Restarting the management network will result in a brief network outage that might temporarily affect running virtual machines.

If a renewed DHCP lease results in a new network identity (IP address or host name), remote management software will be disconnected.

Procedure

1. From the direct console, select Restart Management Network and press Enter.
2. Press F11 to confirm the restart.

Reference: https://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vsphere.install.doc_50%2FGUID-E46F9F62-D1BB-4485-815D-08DBAE320CC3.html

Question 63

An administrator observes the following symptoms for a virtual machine: ♦

CPU usage is consistently above 90%

CPU ready value is consistently above 20%.

Application performance is impacted.



Which two actions should the administrator take to improve the performance of this virtual machine? (Choose two.)

- ☐ Increase the number of vCPUs assigned to this virtual machine.
- ☐ Decrease the number of vCPUs assigned to this virtual machine.
- ☒ Verify that VMware Tools is installed on every ♦ virtual machine on the host
- ☒ Increase the CPU shares assigned to the virtual machine.

Explanation:

Explanation:

Solutions for Consistently High CPU Usage

Temporary spikes in CPU usage indicate that you are making the best use of CPU resources. Consistently high CPU usage might indicate a problem. You can use the vSphere Client CPU performance charts to monitor CPU usage for hosts, clusters, resource pools, virtual machines, and vApps.

Problem

? Host CPU usage constantly is high. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the host.

? Virtual machine CPU usage is above 90% and the CPU ready value is above 20%. Application performance is impacted.

Cause

The host probably is lacking the CPU resources required to meet the demand.

Solution

? Verify that VMware Tools is installed on every virtual machine on the host.

? Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked bar chart on the host's **Virtual Machine** view shows the CPU usage for all virtual machines on the host.

? Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine.

? Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles.

? Increase the amount of memory allocated to the virtual machine. This action decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.

? Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.

? If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.

? Upgrade the physical CPUs or cores on the host if necessary.

? Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

Reference:

<https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.monitoring.doc%2FGUID-5F8147A1-6416-4D29-BA3D-E4CED3966016.html>

Question 64

An administrator is troubleshooting a performance problem with a virtual machine (VM). The VM and ESXi host configuration is as follows:

The application which runs within the virtual machine is highly sensitive to memory latency, but has low processor utilization.

The virtual machine has 6 vCPUs.

The ESXi host CPU topology has 2 Physical CPU Sockets, each with 8 Logical CPUs.

The ESXi host CPU has Hyper Threading enabled.

◆◆

The administrator would like to improve CPU performance for this VM. Other VMs on the host have no performance issues. Which action should the administrator take to resolve the issue?

- ☒ Set the Advanced Parameter numa.vcpu.preferHT = TRUE in the virtual machine configuration file.
- ☐ Set the Advanced Parameter numa.vcpu.preferHT = FALSE in the virtual machine configuration file.
- ☐ Set the Advanced Parameter numa.PreferHT = 1 in the ESXi host configuration file.
- ☐ Set the Advanced Parameter numa.PreferHT = 0 in the ESXi host configuration file.

Explanation:

Explanation:

On systems with hyper-threading enabled, a wide-VM that makes use of full processor cores across NUMA nodes but has less vCPUs configured than the number of logical processors (hardware threads) in each physical NUMA node might benefit from using logical processors with local memory rather than using full cores with remote memory. This can be configured by setting the numa.vcpu.preferHT option to TRUE in the specific VM's advanced configuration.

Reference:

https://books.google.co.in/books?id=GnVvAwAAQBAJ&pg=PA78&lpg=PA78&dq=Advanced+Parameter+numa.vcpu.preferHT&source=bl&ots=o9dp9BqWj&sig=XT1BXwxJ6qeTg__zHfO8uLnZNT0&hl=en&sa=X&ved=0ahUKEwjUy-eejcrNAhWGqY8KHQCKBzwQ6AEIVjAH#v=onepage&q&f=false

Question 65

An administrator is able to manage an ESXi 6.x host connected to vCenter Server using the vSphere Web Client but is unable to connect to the host directly

◆◆

Which action should the administrator take to correct this behavior?

- ☐ Restart management agents on the ESXi host.
- ☒ Disable Lockdown Mode on the ESXi host through vCenter Server.
- ☐ Disable the ESXi firewall with the command `esxcli network firewall unload`.
- ☐ Reboot the ESXi host.

Explanation:

Explanation:

ESXi 6.0:

Starting with vSphere 6.0, you can select normal Lockdown mode or strict Lockdown mode, which offer different degrees of lockdown.◆

Normal Lockdown mode:

In normal lockdown mode the DCUI service is not stopped. If the connection to the vCenter Server is lost and access through the vSphere Web Client is no longer available, privileged accounts can log in to the ESXi host's Direct Console Interface and exit lockdown mode. Only these accounts can access the Direct Console User Interface.◆

Accounts in the Exception User list for lockdown mode who have administrative privileges on the host. The Exception Users list is meant for service accounts that perform very specific tasks. Adding ESXi administrators to this list defeats the purpose of lockdown mode.◆

Users defined in the DCUI.Access advanced option for the host. This option is for emergency access to the Direct Console Interface in case the connection to vCenter Server is lost. These users do not require administrative privileges on the host.◆

Strict Lockdown mode:

In strict lockdown mode the DCUI service is stopped. If the connection to vCenter Server is lost and the vSphere Web Client is no longer available, the ESXi host becomes unavailable unless the ESXi Shell and SSH services are enabled and Exception Users are defined. If you cannot restore the connection to the vCenter Server system, you have to reinstall the host.

For more information on Lockdown mode in vSphere 6.0, see the Lockdown Mode section◆in the vSphere Security Guide.

To enable or disable Lockdown mode from the DCUI:

Log directly in to the ESXi host.

Open the DCUI on the host.

Press◆F2◆for◆Initial Setup.

Press◆Enter◆to toggle the◆Configure Lockdown Mode◆setting.

To enable or disable Lockdown mode from the vSphere Web Client:

Browse to the host in the vSphere Web Client inventory.

Click the Manage tab and click Settings.

Under System, select Security Profile.

In the Lockdown Mode panel, click Edit.

Click Lockdown Mode and select one of the lockdown mode options.

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1008077

Question 66

An administrator is unable to upgrade a vCenter Server Appliance from version 5.1 Update 2 to version 6.0.

◆◆

What is a likely reason for this?

- ☒ vCenter Server Appliance 6.0 does not support upgrades from version 5.1 Update 2.
- ☐ vCenter Server Appliance must be joined to an Active Directory domain before upgrading to version 6.0.
- ☐ vCenter Server Appliance 5.1 Update 2 uses an incompatible database for upgrading to version 6.0.
- ☐ vCenter Server Appliance 6.0 must be upgraded using the vSphere Web Client.

Explanation:

Explanation:

If your current vCenter Server Appliance version is earlier than 5.1 Update 3, you must upgrade to 5.1 Update 3 or later before upgrading to vCenter Server Appliance 6.0.

Reference:

<https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.upgrade.doc%2FGUID-66836F60-A095-4749-86C9-1DAFB5D21070.html>

Question 67

Refer to the Exhibit.

```
[root@ds-lab-vsan12:~] esxcli storage nmp psp roundrobin deviceconfig set --useano=0 -d naa.600000e00d11000000111cd2000f0000
```

An administrator is configuring a storage device as shown in the Exhibit.



What is the expected effect on the stated device after running the command?

- ☐ I/O will rotate on all storage targets regardless of port group state.
- ☒ I/O will rotate on all storage targets that are Active Optimized state only.
- ☐ I/O will rotate on all storage targets that are Active Unoptimized state only.
- ☐ I/O will rotate on all storage targets that are on Available Nodes only.

Explanation:

Explanation:

Command	Description	Options help
storage nmp psp roundrobin deviceconfig set	Allow setting of the Round Robin path options on a given device controlled by the Round Robin Selection Policy.	--bytes -B When the --type option is set to 'bytes' this is the value that will be assigned to the byte limit value for this device. --device -d The device you wish to set the Round Robin settings for. This device must be controlled by the Round Robin Path Selection Policy --help -h Show the help message. --iops -I When the --type option is set to 'iops' this is the value that will be assigned to the I/O operation limit value for this device. --type -t Set the type of the Round Robin path switching that should be enabled for this device. Valid values for type are: bytes: Set the trigger for path switching based on the number of bytes sent down a path. default: Set the trigger for path switching back to default values. iops: Set the trigger for path switching based on the number of I/O operations on a path. --useano -U Set useano to true, to also include non-optimized paths in the set of active paths used to issue I/Os on this device, otherwise set it to false

Reference: https://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vcli.ref.doc_50%2Fesxcli_storage.html

Question 68

An administrator uses the `df -h` command and notices that an NFS datastore is reporting a capacity of 0 Bytes.

◆◆

What condition would cause this to occur?

- ☒ The NFS server on which the datastore resides is down.
- ☐ The datastore was mounted as Read/Write.
- ☐ The datastore was mounted as Read-Only.
- ☐ The datastore was created with NFS version 4.1.

Explanation:

Explanation: Please refer to this link to understand the situation and how to mitigate it: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2073021

Question 69

During a vCenter Server upgrade, an ESXi 6.x host in a High Availability (HA) cluster fails.

◆◆

Which statement is true?

- ☒ HA will fail the virtual machines over to an available host during the vCenter Server upgrade process.
- ☐ HA is unavailable during the vCenter Server upgrade process.
- ☐ HA will fail the virtual machines over to an available host after the vCenter Server upgrade completes.
- ☐ HA will successfully vMotion the virtual machines during the host failure.

Explanation:

Explanation:

During a vCenter Upgrade the VMware HA Agent on all of the ESXi hosts will be updated. ... it's very likely that the HA agent will fail to initialize the first time when the host restarts.

During a vCenter Upgrade the VMware HA Agent on all of the ESXi hosts will be updated. After this happens the "require reboot" flag is set on the hosts, as explained in VMware KB 2034945.

Question 70

Which three statements are correct regarding Fibre Channel over Ethernet (FCoE)? (Choose three.)

- ☒ The network switch must have Priority-based Flow Control (PFC) set to AUTO.
- ☐ The network switch must have Priority-based Flow Control (PFC) set to ON.
- ☐ Each port on the FCoE card must reside on the same vSwitch.
- ☒ Each port on the FCoE card must reside on a separate vSwitch.
- ☒ The ESXi host will require a reboot after moving an FCoE card to a different vSwitch.

Explanation:

Explanation:

Configuration Guidelines for Software FCoE

When setting up your network environment to work with ESXi software FCoE, follow the guidelines and best practices that VMware offers.

Network Switch Guidelines

Follow these guidelines when you configure a network switch for software FCoE environment:

- On the ports that communicate with your ESXi host, disable the Spanning Tree Protocol (STP). Having the STP enabled might delay the FCoE Initialization Protocol (FIP) response at the switch and cause an all paths down (APD) condition.

The FIP is a protocol that FCoE uses to discover and initialize FCoE entities on the Ethernet.

- Turn on Priority-based Flow Control (PFC) and set it to AUTO.
- Make sure that you have a compatible firmware version on the FCoE switch.

Network Adapter Best Practices

If you plan to enable software FCoE adapters to work with network adapters, specific considerations apply.

- Make sure that the latest microcode is installed on the FCoE network adapter.
- If the network adapter has multiple ports, when configuring networking, add each port to a separate vSwitch. This practice helps you to avoid an APD condition when a disruptive event, such as an MTU change, occurs.
- Do not move a network adapter port from one vSwitch to another when FCoE traffic is active. If you need to make this change, reboot your host afterwards.
- If you changed the vSwitch for a network adapter port and caused a failure, moving the port back to the original vSwitch resolves the problem.

Set Up Networking for Software FCoE

Before you activate the software FCoE adapters, you need to create VMkernel network adapters for all physical FCoE NICs installed on your host.

This procedure explains how to create a single VMkernel network adapter connected to a single FCoE physical network adapter through a vSphere standard switch. If your host has multiple network adapters or multiple ports on the adapter, connect each FCoE NIC to a separate standard switch. For more information, see the *vSphere Networking* documentation.

NOTE ESXi supports the maximum of four network adapter ports used for software FCoE.

Procedure

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Click **Actions > Add Networking**.
- 3 Select **VMkernel Network Adapter**, and click **Next**.

Reference: <https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-storage-guide.pdf>

Question 71

An administrator connects to an ESXi 6.x host console in order to shutdown the host.

Which option in the Direct Console User Interface would perform this task?

- ☒ Press the F12 key
- ☐ Press the F2 key
- ☐ Press Alt + F1 simultaneously
- ☐ Press Alt + F2 simultaneously

Explanation:

Explanation:

ESXi 4.x/5.x/6.0

1. From the Direct Console User Interface (DCUI) screen, press F12 to view the shutdown-related options for the ESXi host.

Press F2 to shut down.

Press F11 to reboot.

2. From Local or Remote Tech Support Mode, or from an SSH session, run one of these commands:

Run the reboot command to restart the host.

Run the poweroff command to shut down the host.

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1013193

Question 72

Refer to the Exhibit.

Identity source type:

☐ Active Directory (Integrated Windows Authentication)
☒ Active Directory as a LDAP Server
☐ Open LDAP
☐ Local OS

Identity source settings

Name:

Base DN for users:

Domain name:

Domain alias:

Base DN for groups:

Primary server URL:

Secondary server URL:

Username:

Password:

An administrator is adding an Active Directory over LDAP Identity Source for vCenter Single Sign-On, as indicated in the Exhibit.



What is the correct value to configure for the Domain alias?

- ☒ The domain's NetBIOS name.
- ☐ The fully qualified domain name.
- ☐ vsphere.local
- ☐ A user defined label.

Explanation:

Explanation:

Domain alias

(Optional) The domain's NetBIOS name.

Reference:

<https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.security.doc%2FGUID-B23B1360-8838-4FF2-B074-71643C4CB040.html>

Question 73

Attempting to update an ESXi 6.x host using the following command:

❖❖

```
esxcli software vib update -d update.zip
```

❖❖

Generates the following error:

❖❖

Could not download from depot at /tmp/index.xml, skipping (('tmp/index.xml', '[Errno 4] IOError: <urlopen error [Errno 2] No such file or directory: 'tmp/index.xml'>')) url = /tmp/index.xml

What action should be taken to correct the problem?

- ☒ Add the full file path to the command.
- ☐ Use the update option instead of install.
- ☐ Use the switch -d instead of -v.
- ☐ Replace esxcli with esxupdate.

Explanation:

Explanation:

Put the full path to the .zip even if you're in the same directory

ex: `esxcli software vib install -d /tmp/ESXi600-201504001.zip`

Also make sure you're in maintenance mode

How to perform next/related steps: Check:----

<https://pubs.vmware.com/vsphere-60/index.jsp#com.vmware.vsphere.upgrade.doc/GUID-22A4B153-CB21-47B4-974E-2E5BB8AC6874.html?resultof=%2522>

Question 74

What is the command to list multipathing modules on an ESXi 6.x host?

- ☐ `esxcli storage core list plugin --plugin-class=MP`
- ☐ `esxcli storage core list plugin --class-plugin=MP`
- ☒ `esxcli storage core plugin list --plugin-class=MP`
- ☐ `esxcli storage core plugin list --class-plugin=MP`

Explanation:

Explanation:

Command	Description	Options help
storage core plugin list	List loaded PSA plugins on the system.	--help -h Show the help message. --plugin-class -N Indicate the class of plugin to limit the list to. Allowed values are : Filter: Filter plugins MP: MultiPathing plugins VAAI: VAAI plugins all: All PSA Plugins (default)

Reference:

https://pubs.vmware.com/vsphere50/index.jsp?topic=%2Fcom.vmware.vcli.ref.doc_50%2Fesxcli_storage.html

Question 75

An administrator wants to set a non-default isolation address of 192.168.1.2 for High Availability

Which advanced setting would accomplish this task?

- ☒ Das.isolationaddress0=192.168.1.2
- ☐ Das.useisolationaddress0=192.168.1.2
- ☐ Das.defaultisolationaddress0=192.168.1.2
- ☐ Das.haisolationaddress0=192.168.1.2

Explanation:

Explanation: Details

In VirtualCenter Server 2.0.2 and above, you can specify more than one isolation response address for VMware High Availability (HA). The use of multiple isolation response addresses offers VMware HA a potentially more accurate picture of the network connectivity of a host. There may be situations in which a single isolation address would indicate that a host is in a state of complete isolation from the network, but access to additional isolation addresses would show that only a partial network failure has occurred.

Solution

Multiple isolation response addresses can be specified using the `das.isolationaddress0` through `das.isolationaddress9` options.

◆

To specify the values for these options in Virtual Infrastructure Client:

1. Select the HA cluster.
2. On the Summary tab, click Edit Settings.
3. In the Settings dialog, select VMware HA.
4. Click Advanced Options.

5. In the Advanced Options (HA) dialog, enter the option name and the corresponding value:

Option: `das.isolationaddress0`

Value: A valid IP address other than the default gateway address

Similarly, you can set more isolation response addresses using `das.isolationaddress1` through `das.isolationaddress9`.

6. In the Advanced Options (HA) dialog box, set this option:

Option: `das.usedefaultisolationaddress`

Value: `false`

Notes:

- You must set this option if you are setting multiple isolation response addresses.
- Additionally you should configure `das.usedefaultisolationaddress` to `false` when the default gateway is a device which cannot be pinged.

7. Click **OK**.

8. Click **OK**.

9. Edit settings for the cluster again.

10. Disable HA and click **OK**.

11. Edit settings again.

12. Enable HA and click **OK**.

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1002117

Question 76

VMware vSphere Replication protects virtual machines from partial or complete site failures by replicating the virtual machines between which three sites? (Choose three.)

- ☒ From a source site to a target site.
- ☒ From within a single site from one cluster to another.
- ☒ From multiple source sites to a shared remote target site.
- ☐ From a single source site to multiple remote target sites.
- ☐ From multiple source sites to multiple remote target sites.

Explanation:

Explanation:

Overview of VMware vSphere Replication VMware vSphere Replication is an extension to VMware vCenter Server that provides hypervisor-based virtual machine replication and recovery

vSphere Replication is an alternative to storage-based replication. It protects virtual machines from partial or complete site failures by replicating the virtual machines between the following sites:

From a source site to a target site

Within a single site from one cluster to another

From multiple source sites to a shared remote target site

Reference:

https://pubs.vmware.com/vsphere55/index.jsp?topic=%2Fcom.vmware.vsphere.replication_admin.doc%2FGUID-C987AD18-7C2D-4FA6-B6E4-6B0DDA915A7A.html

Question 77

Which two reasons would prevent Storage Distributed Resource Scheduler (SDRS) from operating on a datastore? (Choose two.)

- ☒ The datastore has Storage I/O Control disabled.
- ☒ The datastore is connected to an unsupported host.
- ☐ The datastore is hosted on an NFS server.
- ☐ The datastore is hosted on an iSCSI server.

Explanation:

Explanation: Storage DRS to datastores is like DRS to hosts. It moves VMDKs from heavily utilized to less utilized datastores according to preset space and I/O metrics, making efficient storage management more achievable.

Storage DRS can run in automatic mode, but most admins are using it in manual mode. This is understandable as we must keep in mind that Storage vMotion is not as light as vMotion, and hence should not be taken lightly.

Storage DRS does use other storage management technologies like Storage I/O Control (SIOC), and can benefit from others like Profile-Driven Storage and vSphere API for Storage Awareness (VASA).

Question 78

Which two identification sources can a vSphere Auto Deploy rule use to identify target hosts? (Choose two.)

- ☐ Processor Serial Number
- ☐ Security Identifier
- ☒ SMBIOS information
- ☒ BIOS UUID

Explanation:

Explanation: You specify the behavior of the Auto Deploy server by using a set of rules written in Power CLI. The Auto Deploy rule engine checks the rule set for matching host patterns to decide which items (image profile, host profile, or vCenter Server location) to provision each host with. The rule engine maps software and configuration settings to hosts based on the attributes of the host. For example, you can deploy image profiles or host profiles to two clusters of hosts by writing two rules, each matching on the network address of one cluster

The rule engine includes rules and rule sets.

Rules

Rules can assign image profiles and host profiles to a set of hosts, or specify the location (folder or cluster) of a host on the target vCenter Server system. A rule can identify target hosts by boot MAC address, SMBIOS information, BIOS UUID, Vendor, Model, or fixed DHCP IP address. In most cases, rules apply to multiple hosts. You create rules by using Auto Deploy PowerCLI cmdlets. After you create a rule, you must add it to a rule set. Only two rule sets, the active rule set and the working rule set, are supported. A rule can belong to both sets, the default, or only to the working rule set. After you add a rule to a rule set, you can no longer change the rule. Instead, you copy the rule and replace either items or patterns. By default, Auto Deploy uses the name of the rule for the copy and hides the original rule.

Reference:

https://pubs.vmware.com/vsphere50/index.jsp?topic=%2Fcom.vmware.vsphere.install.doc_50%2FGUID-EB22DCFF-C606-43CB-A7BC-068AC79A6237.html

Question 79

An administrator has a virtual machine configured with the following settings: ♦

ESXi version: 5.1

CPU: vCPUs 6

Memory: 48GB

Hardware version: 7

VMware Tools: Installed



Which two actions must the administrator take in order to utilize vNUMA? (Choose two.)

- ☐ Upgrade the ESXi host to vSphere 5.5 or later.
- ☒ Upgrade to Virtual Hardware version 8.
- ☐ Configure numa.vcpu.min to 5
- ☒ Configure numa.vcpu.min to 6

Explanation:

Explanation:

Using Virtual NUMA

vSphere 5.0 and later includes support for exposing virtual NUMA topology to guest operating systems, which can improve performance by facilitating guest operating system and application NUMA optimizations.

Virtual NUMA topology is available to hardware version 8 virtual machines and is enabled by default when the number of virtual CPUs is greater than eight. You can also manually influence virtual NUMA topology using advanced configuration options.

You can affect the virtual NUMA topology with two settings in the vSphere Client: number of virtual sockets and number of cores per socket for a virtual machine. If the number of cores per socket (cpuid.coresPerSocket) is greater than one, and the number of virtual cores in the virtual machine is greater than 8, the virtual NUMA node size matches the virtual socket size. If the number of cores per socket is less than or equal to one, virtual NUMA nodes are created to match the topology of the first physical host where the virtual machine is powered on.

Reference:

<https://pubs.vmware.com/vsphere-51/index.jsp#com.vmware.vsphere.resmgmt.doc/GUID-17B629DE-75DF-4C23-B831-08107007FBB9.html>

-----Explanation note-----

Also, Check Page 21,22 https://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.1.pdf

D-) Size your virtual machines so they align with physical NUMA boundaries. For example, if you have a host system with six cores per NUMA node, size your virtual machines with a multiple of six vCPUs (i.e., 6 vCPUs, 12 vCPUs, 18 vCPUs, 24 vCPUs, and so on). Hence answer is D

Also, Check Page 41,42 https://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.1.pdf

Question 80

Which three ports are used by the vSphere Web Client when connecting directly to an ESXi 6.x host? (Choose three.)

- ☒ 443 TCP
- ☒ 902 TCP and UDP
- ☒ 903 TCP
- ☐ 5480 TCP
- ☐ 9443 TCP and UDP

Explanation:

Explanation:

Port	Purpose	Traffic type
443(default)	• HTTPS access • vSphere Client access to vCenter Server • vSphere Client access to ESXi hosts • vSphere Client access to vSphere Update Manager	Incoming tcp
902(default)	vSphere Client access to virtual machine consoles	Incoming and outgoing tcp, Outgoing udp
903	• Remote console traffic generated by user access to virtual machines on a specific host. • vSphere Client access to virtual machine consoles • MKS transactions (xinetd/vmware-authd-mks)	Incoming TCP

Reference: <https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-installation-setup-guide.pdf>

Question 81

An administrator is performing a silent automatic update of VMware Tools on a Windows virtual machine.

◆◆

What syntax needs to be entered into the Advanced Options box?

- ☒ /s /v "/qn" /l "c:\Windows\filename.log"
- ☐ --prefix=/usr/local,/usr/lib,/usr/doc --silent
- ☐ --prefix=c:\Windows,c:\VMtools --silent
- ☐ /fs /v "/qn+" /l "c:\Windows\filename.log"

Explanation:

Explanation:

Procedure

- 1 Select **Automatic Tools Upgrade**.
- 2 (Optional) In the **Advanced Options** field, enter advanced options for the guest operating system.

Option	Description
Microsoft Windows Guest Operating Systems	Enter <code>/s /v "/qn" /I "Microsoft_Windows_location\file-name.log"</code> to perform a silent upgrade of VMware Tools and create a log file in the specified location on the guest operating system.
Linux Guest Operating Systems	■ Enter <code>--default</code> to perform the default behavior. Perform upgrade of VMware Tools. Install tools bin, lib and doc file fault /usrdirectory. ■ Enter <code>--prefix=binary_location,lib_location,doc_location</code> to perform a silent upgrade of VMware Tools and install the binary, and document files in the specified locations.


- 3 Click **OK**.

The **VMware Tools** label on the **Summary** tab changes to **OK**.

Reference: <https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.upgrade.doc%2FGUID-B3989445-72A9-4749-A996-CDDCD1FA634D.html>

Question 82

Refer to the Exhibit.



An administrator wants to increase the capacity of a VMFS5 datastore; but the Increase Size by slider is not available, as shown in the Exhibit.

How should the administrator resolve this problem?

- ☒ Select a valid partition configuration from the drop-down menu.
- ☐ Use fdisk utility to manually resize the partition.
- ☐ Use the mouse to resize the partition indicated in the Partition Layout.
- ☐ Clicking Next will prompt for available usable space.

Explanation:

Explanation:

Procedure

1. Select the datastore to grow and click the Increase Datastore Capacity icon.
2. Select a device from the list of storage devices.

Your selection depends on whether an expandable storage device is available.

Choose the options from Drop down menu list:

Check the link:

<https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.storage.doc%2FGUID-D57FEF5D-75F1-433D-B337-E760732282FC.html>

Question 83

Deletion of an NFS datastore generates the following error:

Sysinfo error on operation returned the following status: Busy

Which step can be performed that will allow the deletion to complete successfully?

- ☒ Storage vMotion any virtual machines on the datastore to another location.
- ☐ Remove the datastore from the Storage DRS cluster in which it resides.
- ☐ Verify the value of the NFS.HeartbeatDelta is not set to 0.
- ☐ Suspend any running virtual machines and unmount the NFS datastore.

Explanation:

Explanation:

To resolve this issue, temporarily stop the VSA plug-in so that the locks on the datastore are released.



To temporarily stop the VSA plug-in:

1. Ensure that no process is locking any files on the datastore:

- a. Connect to the host using SSH. For more information, see Using ESXi Shell in ESXi 5.x (2004746).
- b. Check for file locks on the datastore, see Investigating virtual machine file locks on ESXi/ESX (10051).

If the datastore is free of locks continue to step 2. If there are active files locks in place, resolve them and attempt to unmount the NFS mount again.

2. Disable the VSA plug-in from the vSphere Web Client:

- a. From the Home page of the vSphere Web Client, click Administration.
- b. Under Solutions, click Client Plug-Ins.
- c. Right-click the VSA plug-in and click Disable.

3. Stop the VSA service on vCenter Server.

4. Restart the management agents on the host. For more information, see Restarting the Management agents on an ESXi or ESX host (1003490).

5. Unmount the datastore using the vSphere Client.

6. Start the VSA service on vCenter Server.

7. Re-enable the VSA plug-in in the vSphere Web Client:

- a. From the Home page of the vSphere Web Client, click Administration.
- b. Under Solutions, click Plug-In Management.
- c. Right-click the VSA plug-in and click Enable.

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2102598

Question 84

What are two likely causes for a DRS cluster to become unbalanced? (Choose two.)

- ☐ Migration threshold is too low.
- ☒ Affinity rules are preventing virtual machines from being moved.
- ☒ A device is mounted to a virtual machine preventing vMotion.
- ☐ Migration cost is too low.

Explanation:

Explanation:

Load Imbalance on Cluster

A cluster has a load imbalance of resources.

Problem

A cluster might become unbalanced because of uneven resource demands from virtual machines and unequal capacities of hosts.

Cause

The following are possible reasons why the cluster has a load imbalance:

The migration threshold is too high.

A higher threshold makes the cluster a more likely candidate for load imbalance.

VM/VM or VM/Host DRS rules prevent virtual machines from being moved.

DRS is disabled for one or more virtual machines.

A device is mounted to one or more virtual machines preventing DRS from moving the virtual machine in order to balance the load.

Virtual machines are not compatible with the hosts to which DRS would move them. That is, at least one of the hosts in the cluster is incompatible for the virtual machines that would be migrated. For example, if host A's CPU is not vMotion-compatible with host B's CPU, then host A becomes incompatible for powered-on virtual machines running on host B.

It would be more detrimental for the virtual machine's performance to move it than for it to run where it is currently located. This may occur when loads are unstable or the migration cost is high compared to the benefit gained from moving the virtual machine.

vMotion is not enabled or set up for the hosts in the cluster.

Reference:

<https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.troubleshooting.doc%2FGUID-0C132C85-E7BF-445E-925D-31420A06B01C.html>

Question 85

An administrator has upgraded a Distributed vCenter Server environment from 5.5 to 6.0.

◆◆

What is the next step that should be taken?

- ☒ vCenter Inventory Service must be manually stopped and removed.
- ☐ vCenter Inventory Service must be changed from manual to automatic.
- ☐ vCenter Inventory Service must be manually stopped and restarted.
- ☐ vCenter Inventory Service must be changed from automatic to manual.

Explanation:

Explanation:

vCenter Server 5.x Distributed Service Migration During Upgrade			
Service Name	Service Location Before Upgrade	Service Location After Upgrade	Post Upgrade Actions
vCenter Inventory Service	Not installed on the vCenter Serversystem	Installed on the vCenter Serversystem	vCenter Inventory Service 5.x data is copied to the Inventory Service 6.0 instance that is installed with vCenter Server 6.0. You do not need to copy it manually. vCenter Inventory Service 5.x is still running but no longer used. It must be manually stopped and removed.

Reference: <https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.upgrade.doc%2FGUID-4BFB12D8-9FCA-4AB1-A44F-2986966F0AD5.html>

Question 86

Which two features are deprecated in Network I/O Control 3 (NIOC3)? (Choose two.)

- ☒ Class Of Service (COS) Tagging
- ☐ Bandwidth Allocation
- ☒ User-defined network resource pools
- ☐ Admission control

Explanation:

Explanation : The upgrade of a NIOC to version 3 is disruptive. Certain functionality is available only in NIOC v2 and is removed during the upgrade to version 3:

User-defined network resource pools including all associations between them and existing distributed port groups

You can preserve certain resource allocation settings by transferring the shares from the user-defined network resource pools to shares for individual network adapters.

Existing associations between ports and user-defined network resource pools

In Network I/O Control version 3, you cannot associate an individual distributed port to a network resource pool that is different from the pool assigned to the parent port group.

CoS tagging of the traffic that is associated with a network resource pool

Network I/O Control version 3 does not support marking traffic that has higher QoS demands with CoS tags. After the upgrade, to restore CoS tagging of traffic that was associated with a user-defined network resource pool, use the traffic filtering and marking networking policy.

Reference: <http://www.virtten.net/2015/09/vcp6-delta-part-7-network-enhancements/>

Question 87

What are three reasons why a virtual machine might fail to power on? (Choose three.)

- ☒ The virtual machine is running on an ESXi host which has an expired license.
- ☒ The virtual machine is running on a datastore which has insufficient disk space for the .vswp file.
- ☒ The virtual machine is in a cluster with vSphere HA Admission control enabled.
- ☐ The virtual machine has a disconnected network adapter.
- ☐ The virtual machine does not have a Virtual Hard Disk assigned.

Explanation:

Explanation:

When a virtual machine fails to power on, a reason may be logged to the `vmware.log` file for the virtual machine, to the management agent logs, or presented in the client. Review any messages and consider these points:

1. The virtual machine monitor may be asking a question to be answered during startup. A virtual machine may pause the power-on task at 95% to obtain additional information from the administrator. For more information, see [Powering on a virtual machine pauses at 95% while waiting for a question to be answered \(1027096\)](#).

2. Creating a new power-on task may fail if another task for the virtual machine or other component is already in progress, and multiple concurrent tasks on the object are not permitted. For more information, see:

Collecting information about tasks in VMware ESX and ESXi (1013003)

vCenter operation times out with the error: Operation failed since another task is in progress (1004790).

3. A virtual machine may fail to power on if licensing requirements are not met. For more information, see:

Cannot Power on Virtual Machines, "Not enough licenses installed to perform the operation" Error Message (7114568)

The [Managing ESX/ESXi and vCenter Server Licenses](#) section of the [Datacenter Administration Guide for vSphere 4 or higher](#)

The [VirtualCenter and ESX Server Licensing Overview](#) section of the [Installation Guide for VMware Infrastructure 3](#)

4. The virtual machine may be configured to reserve physical memory on the host, but the host memory is over-committed and the required memory is unavailable. For more information, see:

Virtual machine does not power on when there is high CPU reservation (1001637)

Powering on a virtual machine fails with the error: `memoryAllocation.reservation` (1036914).

5. The virtual machine may be starting in a VMware High Availability cluster with strict admission control enabled, and there are insufficient resources to guarantee failover for all virtual machines. For more information, see:

Implications of enabling or disabling VMware HA strict admission control when using DRS and VMware DPM (1007006)

Diagnosing insufficient fail over capacity on a VMware High Availability Cluster (1003717)

Automating High Availability (HA) Services with VMware HA whitepaper

6. A file required for starting the virtual machine, such as a virtual disk or swap file, may be unavailable or missing. For more information, see [Investigating virtual machine file locks in ESX/ESXi \(10051\)](#).

7. The virtual machine may have been previously suspended and making use of CPU features which are unavailable or incompatible with the CPU features available on this host. The virtual machine cannot be started without the required features. For more information, see:

Powering on a virtual machine from a suspend state or reverting to a snapshot fails (1038218)

Virtual machines fail to power on due to a suspend state from an incompatible CPU type (1000241)

To retain the suspended state, move the virtual machine back to the host it was originally suspended on and power-on the virtual machine there.◆

To discard the suspend state, and power on the virtual machine in a crash-consistent manner, see◆Unable to power on a suspended virtual machine using vSphere Client (1004606).

8. The virtual machine may require both a VT-capable CPU and the VT feature to be enabled in the host system's BIOS. This is true for all 64-bit virtual machines. If the VT feature is unavailable, the virtual machine may◆produce the message◆msg.cpuid.noLongmode. For more information, see◆Enabling VT on Intel EM64T Systems for ESX Server 3 (3282933).

9. The virtual machine may require another CPU feature which is unavailable on this host. The virtual machine may◆produce a message similar to◆msg.cpuid.<FeatureName>, identifying the specific feature it has been configured to require. Move the virtual machine back to the host which has the required CPU features, or edit the virtual machine's configuration to remove the requirement.

10. The virtual machine may start, but quickly fail with an error during startup. Review the contents of the vmware.log file in the virtual machine's directory for any errors or warnings, and search the Knowledge Base for the error or warning. Base your troubleshooting on the specific messages seen in the logs. For more information, see:◆

Determining why a virtual machine was powered off or restarted (1019064)

Interpreting virtual machine monitor and executable failures (1019471)

11. If the virtual machine does successfully power on, but the guest OS doesn't start correctly, there may be an incompatibility between the virtual hardware and drivers within the guest OS. For example, a missing SCSI driver may be required for booting.◆For more information, see◆Windows virtual machine configured to use a BusLogic SCSI controller reports that the operating system does not support the controller (2007603).

12. If the guest OS, or a driver or application within the virtual machine experiences a problem during startup, the guest OS may become unresponsive. Continue troubleshooting.◆For more information, see◆Troubleshooting unresponsive guest operating system issues (1007818).

Reference:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2001005

Question 88

In a vSphere High Availability cluster, what is the VM Monitoring I/O stats interval default value?

- ☐ 60 seconds
- ☐ 90 seconds
- ☒ 120 seconds
- ☐ 180 seconds

Explanation:

Explanation: The I/O stats interval determines if any disk or network activity has occurred for the virtual machine during the previous two minutes (120 seconds). If not, the virtual machine is reset. This default value (120 seconds) can be changed using the advanced attribute das.iostatsinterval.

Reference: https://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vsphere.avail.doc_50%2FGUID-62B80D7A-C764-40CB-AE59-752DA6AD78E7.html

Question 89

An administrator is attempting to enable Fault Tolerance on a virtual machine with 4 vCPUs and observes the following in the vSphere Web Client:

❖❖

Fault Tolerance has not been licensed on host <hostname>.

❖❖

What is the minimum licensed edition that supports this configuration?

- ☐ Enterprise
- ☒ Enterprise Plus
- ☐ Standard
- ☐ Essentials Plus

Explanation:

Explanation:

Licensing

The number of vCPUs supported by a single fault tolerant VM is limited by the level of licensing that you have purchased for vSphere. Fault Tolerance is supported as follows:

vSphere Standard and Enterprise. Allows up to 2 vCPUs

vSphere Enterprise Plus. Allows up to 4 vCPUs

Reference:

<https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.avail.doc%2FGUID-57929CF0-DA9B-407A-BF2E-E7B72708D825.html>

Question 90

An administrator subscribes to the vCloud Air Disaster Recovery service.



Which replicated objects can be directly monitored and managed?

- ☐ Virtual machine Snapshots
- ☐ vApps
- ☒ Virtual machines
- ☐ ESXi Hosts

Explanation:

Explanation:

VMware provides a disaster recovery service that lets you protect your workloads from a disaster or disruptive event by replicating them to the cloud. You can replicate virtual machines from your vSphere-based source site to vCloud Air, and recover and use the replicated virtual machines in the cloud if the source site is unavailable.

The Disaster Recovery service is based on vSphere Replication, an extension to VMware vSphere that provides hypervisor-based virtual machine replication and recovery.

You use vSphere Replication at your source site to configure your environment and replicate virtual machines to vCloud Air. You can use vSphere Replication, the vCloud Air portal, or the vCloud Air plug-in to monitor and manage the replicated virtual machines in the cloud.

Reference:

<http://pubs.vmware.com/vchspplugin-15/index.jsp?topic=%2Fcom.vmware.vchspplugin.doc%2FGUID-92B7E03C-8DDE-4F91-A64C-8E5AF7BD0F66.html>

Question 91

Refer to the Exhibit.

```
7:03:33pm up 17 days 6:15, 750 worlds, 21 VMs, 47 vCPUs; CPU load average: 0.17, 0.24, 0.28
Power Usage: 158W, Power Cap: N/A
PSTATE MHZ:
CPU %USED %UTIL %C0 %C1 %C2
0 8.0 17.5 18 52 30
1 4.9 11.9 12 49 39
2 7.7 15.8 16 63 21
3 3.3 8.0 8 56 36
4 9.3 19.0 19 64 17
5 3.8 9.4 9 65 26
6 5.2 11.6 11 60 28
7 6.8 14.1 14 54 32
8 4.8 10.6 11 77 12
9 4.4 9.4 9 55 35
10 3.2 7.6 9 76 14
11 3.5 8.0 8 59 33
12 13.8 25.8 24 34 41
13 7.4 14.9 15 35 50
14 13.6 25.6 26 26 48
15 3.6 7.9 8 44 48
16 6.8 13.3 13 23 63
17 3.5 7.5 7 31 62
18 6.4 12.5 12 33 55
19 3.4 7.2 7 29 64
20 2.7 5.7 6 37 58
21 5.2 9.6 10 19 72
22 4.5 8.7 9 33 58
23 4.4 8.5 8 41 51
```

An administrator is troubleshooting intermittent poor performance of virtual machines in a vSphere 6.x cluster. Investigating esxtop data shows that the only statistic that stands out is %CSTP as depicted in Exhibit 1:

```

[00:00]sign up 66 days 7:56, 982 nodes, 1 Vm, 8 vCPUs; CPU 100% averages 1.31, 1.36, 1.34
[00:00]RIBZ(1) : 00 00 00 00 AVG: 00
[00:00]RIBZ(1) : 00 00 00 00 AVG: 00

```

ID	BID NAME	WALL	VMEM	MEM	VSTX	WALLT	VMEMT2	VSTXT	RIBL	ROWBL	RIBP
00470102	00470102 PARTN	10	194.25	20.71	0.00	402.31	8.00	67.79	0.00	0.62	40.05
00470102	00470102 BOLD	10	197.18	20.71	0.00	401.34	6.16	54.29	0.00	0.64	40.43

The administrator proceeds to switch to the **Power Management** screen and observes the data depicted in Exhibit 2:

```

T:00:33pm up 17 days 6:15, 750 users, 21 VMs, 47 vCPUs: CPU load average: 0.17, 0.24, 0.28
Power Usage: 150W, Power Cap: N/A
ESTATE MIB:

CPU MIBS:
MIB1 MIB2 MIB3 MIB4 MIB5 MIB6 MIB7 MIB8 MIB9 MIB10 MIB11 MIB12 MIB13 MIB14 MIB15 MIB16 MIB17 MIB18 MIB19 MIB20 MIB21 MIB22 MIB23 MIB24 MIB25 MIB26 MIB27 MIB28 MIB29 MIB30 MIB31 MIB32 MIB33 MIB34 MIB35 MIB36 MIB37 MIB38 MIB39 MIB40 MIB41 MIB42 MIB43 MIB44 MIB45 MIB46 MIB47 MIB48 MIB49 MIB50 MIB51 MIB52 MIB53 MIB54 MIB55 MIB56 MIB57 MIB58 MIB59 MIB60 MIB61 MIB62 MIB63 MIB64 MIB65 MIB66 MIB67 MIB68 MIB69 MIB70 MIB71 MIB72 MIB73 MIB74 MIB75 MIB76 MIB77 MIB78 MIB79 MIB80 MIB81 MIB82 MIB83 MIB84 MIB85 MIB86 MIB87 MIB88 MIB89 MIB90 MIB91 MIB92 MIB93 MIB94 MIB95 MIB96 MIB97 MIB98 MIB99 MIB100 MIB101 MIB102 MIB103 MIB104 MIB105 MIB106 MIB107 MIB108 MIB109 MIB110 MIB111 MIB112 MIB113 MIB114 MIB115 MIB116 MIB117 MIB118 MIB119 MIB120 MIB121 MIB122 MIB123 MIB124 MIB125 MIB126 MIB127 MIB128 MIB129 MIB130 MIB131 MIB132 MIB133 MIB134 MIB135 MIB136 MIB137 MIB138 MIB139 MIB140 MIB141 MIB142 MIB143 MIB144 MIB145 MIB146 MIB147 MIB148 MIB149 MIB150 MIB151 MIB152 MIB153 MIB154 MIB155 MIB156 MIB157 MIB158 MIB159 MIB160 MIB161 MIB162 MIB163 MIB164 MIB165 MIB166 MIB167 MIB168 MIB169 MIB170 MIB171 MIB172 MIB173 MIB174 MIB175 MIB176 MIB177 MIB178 MIB179 MIB180 MIB181 MIB182 MIB183 MIB184 MIB185 MIB186 MIB187 MIB188 MIB189 MIB190 MIB191 MIB192 MIB193 MIB194 MIB195 MIB196 MIB197 MIB198 MIB199 MIB200 MIB201 MIB202 MIB203 MIB204 MIB205 MIB206 MIB207 MIB208 MIB209 MIB210 MIB211 MIB212 MIB213 MIB214 MIB215 MIB216 MIB217 MIB218 MIB219 MIB220 MIB221 MIB222 MIB223 MIB224 MIB225 MIB226 MIB227 MIB228 MIB229 MIB230 MIB231 MIB232 MIB233 MIB234 MIB235 MIB236 MIB237 MIB238 MIB239 MIB240 MIB241 MIB242 MIB243 MIB244 MIB245 MIB246 MIB247 MIB248 MIB249 MIB250 MIB251 MIB252 MIB253 MIB254 MIB255 MIB256 MIB257 MIB258 MIB259 MIB260 MIB261 MIB262 MIB263 MIB264 MIB265 MIB266 MIB267 MIB268 MIB269 MIB270 MIB271 MIB272 MIB273 MIB274 MIB275 MIB276 MIB277 MIB278 MIB279 MIB280 MIB281 MIB282 MIB283 MIB284 MIB285 MIB286 MIB287 MIB288 MIB289 MIB290 MIB291 MIB292 MIB293 MIB294 MIB295 MIB296 MIB297 MIB298 MIB299 MIB300 MIB301 MIB302 MIB303 MIB304 MIB305 MIB306 MIB307 MIB308 MIB309 MIB310 MIB311 MIB312 MIB313 MIB314 MIB315 MIB316 MIB317 MIB318 MIB319 MIB320 MIB321 MIB322 MIB323 MIB324 MIB325 MIB326 MIB327 MIB328 MIB329 MIB330 MIB331 MIB332 MIB333 MIB334 MIB335 MIB336 MIB337 MIB338 MIB339 MIB340 MIB341 MIB342 MIB343 MIB344 MIB345 MIB346 MIB347 MIB348 MIB349 MIB350 MIB351 MIB352 MIB353 MIB354 MIB355 MIB356 MIB357 MIB358 MIB359 MIB360 MIB361 MIB362 MIB363 MIB364 MIB365 MIB366 MIB367 MIB368 MIB369 MIB370 MIB371 MIB372 MIB373 MIB374 MIB375 MIB376 MIB377 MIB378 MIB379 MIB380 MIB381 MIB382 MIB383 MIB384 MIB385 MIB386 MIB387 MIB388 MIB389 MIB390 MIB391 MIB392 MIB393 MIB394 MIB395 MIB396 MIB397 MIB398 MIB399 MIB400 MIB401 MIB402 MIB403 MIB404 MIB405 MIB406 MIB407 MIB408 MIB409 MIB410 MIB411 MIB412 MIB413 MIB414 MIB415 MIB416 MIB417 MIB418 MIB419 MIB420 MIB421 MIB422 MIB423 MIB424 MIB425 MIB426 MIB427 MIB428 MIB429 MIB430 MIB431 MIB432 MIB433 MIB434 MIB435 MIB436 MIB437 MIB438 MIB439 MIB440 MIB441 MIB442 MIB443 MIB444 MIB445 MIB446 MIB447 MIB448 MIB449 MIB450 MIB451 MIB452 MIB453 MIB454 MIB455 MIB456 MIB457 MIB458 MIB459 MIB460 MIB461 MIB462 MIB463 MIB464 MIB465 MIB466 MIB467 MIB468 MIB469 MIB470 MIB471 MIB472 MIB473 MIB474 MIB475 MIB476 MIB477 MIB478 MIB479 MIB480 MIB481 MIB482 MIB483 MIB484 MIB485 MIB486 MIB487 MIB488 MIB489 MIB490 MIB491 MIB492 MIB493 MIB494 MIB495 MIB496 MIB497 MIB498 MIB499 MIB500 MIB501 MIB502 MIB503 MIB504 MIB505 MIB506 MIB507 MIB508 MIB509 MIB510 MIB511 MIB512 MIB513 MIB514 MIB515 MIB516 MIB517 MIB518 MIB519 MIB520 MIB521 MIB522 MIB523 MIB524 MIB525 MIB526 MIB527 MIB528 MIB529 MIB530 MIB531 MIB532 MIB533 MIB534 MIB535 MIB536 MIB537 MIB538 MIB539 MIB540 MIB541 MIB542 MIB543 MIB544 MIB545 MIB546 MIB547 MIB548 MIB549 MIB550 MIB551 MIB552 MIB553 MIB554 MIB555 MIB556 MIB557 MIB558 MIB559 MIB560 MIB561 MIB562 MIB563 MIB564 MIB565 MIB566 MIB567 MIB568 MIB569 MIB570 MIB571 MIB572 MIB573 MIB574 MIB575 MIB576 MIB577 MIB578 MIB579 MIB580 MIB581 MIB582 MIB583 MIB584 MIB585 MIB586 MIB587 MIB588 MIB589 MIB590 MIB591 MIB592 MIB593 MIB594 MIB595 MIB596 MIB597 MIB598 MIB599 MIB600 MIB601 MIB602 MIB603 MIB604 MIB605 MIB606 MIB607 MIB608 MIB609 MIB610 MIB611 MIB612 MIB613 MIB614 MIB615 MIB616 MIB617 MIB618 MIB619 MIB620 MIB621 MIB622 MIB623 MIB624 MIB625 MIB626 MIB627 MIB628 MIB629 MIB630 MIB631 MIB632 MIB633 MIB634 MIB635 MIB636 MIB637 MIB638 MIB639 MIB640 MIB641 MIB642 MIB643 MIB644 MIB645 MIB646 MIB647 MIB648 MIB649 MIB650 MIB651 MIB652 MIB653 MIB654 MIB655 MIB656 MIB657 MIB658 MIB659 MIB660 MIB661 MIB662 MIB663 MIB664 MIB665 MIB666 MIB667 MIB668 MIB669 MIB670 MIB671 MIB672 MIB673 MIB674 MIB675 MIB676 MIB677 MIB678 MIB679 MIB680 MIB681 MIB682 MIB683 MIB684 MIB685 MIB686 MIB687 MIB688 MIB689 MIB690 MIB691 MIB692 MIB693 MIB694 MIB695 MIB696 MIB697 MIB698 MIB699 MIB700 MIB701 MIB702 MIB703 MIB704 MIB705 MIB706 MIB707 MIB708 MIB709 MIB710 MIB711 MIB712 MIB713 MIB714 MIB715 MIB716 MIB717 MIB718 MIB719 MIB720 MIB721 MIB722 MIB723 MIB724 MIB725 MIB726 MIB727 MIB728 MIB729 MIB730 MIB731 MIB732 MIB733 MIB734 MIB735 MIB736 MIB737 MIB738 MIB739 MIB740 MIB741 MIB742 MIB743 MIB744 MIB745 MIB746 MIB747 MIB748 MIB749 MIB750 MIB751 MIB752 MIB753 MIB754 MIB755 MIB756 MIB757 MIB758 MIB759 MIB760 MIB761 MIB762 MIB763 MIB764 MIB765 MIB766 MIB767 MIB768 MIB769 MIB770 MIB771 MIB772 MIB773 MIB774 MIB775 MIB776 MIB777 MIB778 MIB779 MIB780 MIB781 MIB782 MIB783 MIB784 MIB785 MIB786 MIB787 MIB788 MIB789 MIB790 MIB791 MIB792 MIB793 MIB794 MIB795 MIB796 MIB797 MIB798 MIB799 MIB800 MIB801 MIB802 MIB803 MIB804 MIB805 MIB806 MIB807 MIB808 MIB809 MIB810 MIB811 MIB812 MIB813 MIB814 MIB815 MIB816 MIB817 MIB818 MIB819 MIB820 MIB821 MIB822 MIB82
```

◆Based on the information in the exhibits, which two configurations are probable causes of the poor performance? (Choose two.)

- ☒ The active power policy is set to Low Power.
- ☒ The host has active Sleep States configured in the BIOS.
- ☐ The active power policy is set to High Performance.
- ☐ The host has active Power States configured in the BIOS.

Explanation:

Explanation:

A and B

Analyzing esxtop columns

Refer to this table for relevant columns and descriptions of these values:

Column Description

CMDS/s

This is the total amount of commands per second and includes IOPS (Input/Output Operations Per Second) and other SCSI commands such as SCSI reservations, locks, vendor string requests, unit attention commands etc. being sent to or coming from the device or virtual machine being monitored.

In most cases, CMDS/s = IOPS unless there are a lot of metadata operations (such as SCSI reservations)

DAVG/cmd This is the average response time in milliseconds per command being sent to the device.

KAVG/cmd This is the amount of time the command spends in the VMkernel.

GAVG/cmd This is the response time as it is perceived by the guest operating system. This number is calculated with the formula: DAVG + KAVG = GAVG

Link:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1008205

And, BIOS:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1018206

Question 92

What are three recommended prerequisites before upgrading virtual machine hardware? (Choose three.)

- ☒ Create a backup or snapshot of the virtual machine.
- ☒ Upgrade VMware Tools to the latest version.
- ☒ Verify that the virtual machine is stored on VMFS3, VMFS5, or NFS datastores
- ☐ Detach all CD-ROM/ISO images from the virtual machines.
- ☐ Set the Advanced Parameter virtualHW.version = 11

Explanation:

Explanation: Before you upgrade the virtual hardware:

1. Create a backup or snapshot of the virtual machine. For more information, see:

Take a Snapshot in the vSphere Web Client section in the vSphere 5.1 Virtual Machine Administration guide.

Take a Snapshot in the vSphere Client section in the vSphere 5.1 Virtual Machine Administration guide.

2. Upgrade VMware Tools. On Microsoft Windows virtual machines, if you upgrade the virtual hardware before you upgrade VMware Tools, the virtual machine might lose its network settings.

3. Verify that all .vmdk files are available to the ESXi/ESX hosts on a VMFS 3, VMFS 5, or NFS datastore.

4. Verify that the virtual machines are stored on VMFS 3, VMFS 5 or NFS datastores.

5. Determine the version of the virtual hardware by selecting the virtual machine from the vSphere Client or vSphere Web Client and clicking the Summary tab. The VM Version label in the Compatibility field displays the virtual hardware version

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1010675

Question 93

An administrator tries to capture network traffic for a virtual machine, but cannot see the expected traffic in the packet capture tool.



Which step can resolve the problem?

- ☐ Migrate the◆virtual machine◆to a Distributed Virtual Switch.
- ☒ Enable Promiscuous Mode on the relevant port group.
- ☐ Modify the default value of MAC Address changes.
- ☐ Enable Forged Transmits on the virtual machine.

Explanation:

Explanation: When promiscuous mode is enabled at the portgroup level, objects defined within that portgroup have the option of receiving all incoming traffic on the vSwitch. Interfaces and virtual machines within the portgroup will be able to see all traffic passing on the vSwitch, but all other portgroups within the same virtual switch do not.

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1002934

Question 94

An administrator is experiencing network connectivity issues between virtual machines. The virtual machines and hosts are◆configured◆as follows:◆

VM1 is◆running◆on Host1

VM2 is◆running◆on Host2◆

Both Host1 and Host2 are attached to◆the vSphere◆Distributed Switch dvSwitch1

Both Host1 and Host2 are using vmnic0 and vmnic1 on dvSwitch1

Both virtual machines are using the default portgroup for network traffic



What are three settings the administrator should investigate while troubleshooting◆the◆connectivity issue? (Choose three.)

- ☒ VLANs of the physical NICs
- ☒ Failover order of the uplinks
- ☒ Virtual NIC connectivity to the dvSwitch
- ☐ Security policy of the portgroup
- ☐ Traffic shaping on the portgroup

Explanation:

Explanation: Please find link to understand more about the connectivity issues https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1003893

Question 95

An administrator notices that there is an all paths down (APD) event occurring for the software FCoE storage.



What◆is a likely◆cause?

- ☒ Spanning Tree Protocol is enabled on the network ports.
- ☐ Spanning Tree Protocol is disabled on the network ports.
- ☐ Spanning Tree Protocol is enabled on the storage processors.
- ☐ Spanning Tree Protocol is disabled on the storage processors.

Explanation:

Explanation : Configuration Guidelines for Software FCoE

When setting up your network environment to work with ESXi software FCoE, follow the guidelines and best practices that VMware offers.

Network Switch Guidelines

Follow these guidelines when you configure a network switch for software FCoE environment:

On the ports that communicate with your ESXi host, disable the Spanning Tree Protocol (STP). Having the STP enabled might delay the FCoE Initialization Protocol (FIP) response at the switch and cause an all paths down (APD) condition. The FIP is a protocol that FCoE uses to discover and initialize FCoE entities on the Ethernet.

Turn on Priority-based Flow Control (PFC) and set it to AUTO.

Make sure that you have a compatible firmware version on the FCoE switch.

Reference: <https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-storage-guide.pdf>

Question 96

Which two◆considerations◆should an administrator◆keep in mind when◆booting from Software Fiber Channel over Ethernet◆(FCoE)? (Choose two.)

- ☐ Software FCoE boot configuration can be changed from within◆ESXi.
- ☐ Software FCoE boot firmware cannot export information in FBFT format.
- ☒ Multipathing is not supported at pre-boot.
- ☒ Boot LUN cannot be shared with other hosts even on shared storage.

Explanation:

Explanation : Considerations

You cannot change software FCoE boot configuration from within ESX

Coredump is not supported on any software FCoE LUNs, including the boot LUN.

Multipathing is not supported at pre-boot.

Boot LUN cannot be shared with other hosts even on shared storage

Reference: <https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-storage-guide.pdf>

Question 97

Where is a Virtual SAN Fault Domain configured?

- ☒ VMware Virtual SAN Cluster configuration
- ☐ VMware High Availability Cluster configuration
- ☐ Distributed Resource Scheduler configuration
- ☐ Datacenter Advanced Settings configuration

Explanation:

Explanation:

Procedure

- 1 Navigate to the Virtual SAN cluster in the vSphere Web Client.
- 2 Click the Manage tab and click Settings.
- 3 Under Virtual SAN, click Fault Domains and Stretched Cluster.
- 4 Click the Create a new fault domain icon (+).
- 5 Type the fault domain name.
- 6 From the Show drop-down menu, select Hosts not in fault domain to view the list of hosts that are not assigned to a fault domain or select Show All Hosts to view all hosts in the cluster.
- 7 Select one or more hosts to add to the fault domain.

A fault domain cannot be empty. You must select at least one host to include in the fault domain.

- 8 Click OK.

The selected hosts appear in the fault domain

Reference: <https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.virtualsan.doc%2FGUID-C365ACE8-7485-4463-A12C-71D1917A4930.html>

Question 98

An administrator is using Auto Deploy to create several ESXi 6.x hosts that will be connected to a vSphere Distributed Switch. After the hosts are deployed, the administrator notices that LACP packets are not being sent between them.

◆◆

Which statement best describes why this issue is occurring?

- ☐ LACP is not enabled on the vCenter Server.
- ☒ The LACP support settings do not exist in the host profile.
- ☐ The LACP installation bundle is not included in the image profile.
- ☐ LACP has not been configured on the Auto Deploy server.

Explanation:

Explanation: LACP Limitations on a vSphere Distributed Switch

Link Aggregation Control Protocol (LACP) on a vSphere distributed switch allows network devices to negotiate automatic bundling of links by sending LACP packets to a peer. However, there are some limitations when using LACP with a vSphere distributed switch.

LACP only works with IP Hash load balancing and Link Status Network failover detection.

LACP is not compatible with iSCSI software multipathing.

vSphere only supports one LACP group per distributed switch, and only one LACP group per host.

LACP settings do not exist in host profiles.

LACP between two nested ESXi hosts is not possible.

LACP does not work with port mirroring.

Reference: <https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.networking.doc%2FGUID-3FDE1E96-9217-4FE6-8B76-6E3A64766828.html>

Question 99

An administrator wants to configure a High Availability cluster to allow virtual machines a 10 minute window to shut down in the event of a Host Isolation incident.

◆◆

What two configuration settings would satisfy this requirement? (Choose two.)

- ☐ Set the advanced option das.isolationshutdowntimeout = 10.
- ☒ Set the advanced option das.isolationshutdowntimeout = 600.
- ☒ Configure Host Isolation Response to Shut Down and Restart VMs.
- ☐ Configure Host Isolation Response to Power Off and Restart VMs.

Explanation:

Explanation:

das.isolationShutdownTimeout

The number of seconds an FDM waits for a virtual machine to power off after initiating a guest shutdown before the FDM issues a power off. If the option is unset, 300s is used.

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2033250

Question 100

An administrator notices that one virtual machine is in an orphaned state.

◆◆

What are two reasons that a virtual machine can appear as orphaned? (Choose two.)

- ☒ A VMware High Availability host failure has occurred.
- ☒ The virtual machine was unregistered directly on the host.
- ☐ The ESXi host is disconnected.
- ☐ The user does not have privilege to access the virtual machine.

Explanation:

Explanation: Please refer to the following link to understand more about the Orphaned State: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1003742

Question 101

An administrator created a six node Virtual SAN cluster, created a fault domain, and moved three of the six nodes into that domain.

◆

A node that is a member of the fault domain fails.

◆◆

What is the expected result?

- ☒ The remaining two fault domain members are treated as failed.
- ☐ The remaining two fault domain members stay protected by the domain.
- ☐ One of the non-member nodes will be automatically added to the fault domain.
- ☐ VMware High Availability will restart virtual machines on remaining nodes in the domain.

Explanation:

Explanation:

Defines the number of host and device failures a virtual machine object can tolerate. For n failures tolerated, n+1 copies of the virtual machine object are created and 2*n+1 hosts contributing storage

are required. When provisioning a virtual machine, if you do not choose a storage policy, Virtual SAN assigns this policy as the default virtual machine storage policy.

Default value is 1. Maximum value is 3. If fault domains are configured, 2n+1 fault domains with hosts

contributing capacity are required. A host, which is not part of any fault domain is considered as its own single host fault domain.

Default value is 1. Maximum value is 3.

NOTE If you do not want Virtual SAN to protect a single mirror copy of virtual machine objects, you can specify the Number of failures to tolerate=0. However, the host might experience unusual delays when

entering maintenance mode. The delay occurs because Virtual SAN has to evacuate the object from the host for the maintenance operation to complete successfully. Setting the Number of failures to tolerate=0 means that your data is unprotected, and you might lose data when the Virtual SAN cluster encounters a device failure.

NOTE When creating a new storage policy, if you do not specify any value for Number of failures to tolerate, by default, Virtual SAN creates a single mirror copy of the virtual machine objects and tolerates only one failure. However, in the event of a multiple component failures your data might be at risk.

link: <https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/virtual-san-60-administration-guide.pdf>

Question 102

Which two statements are true regarding Storage Multipathing Plug-Ins? (Choose two.)

- ☐ The default Path Selection Policy is VMW_PSP_MRU for iSCSI or FC devices.
- ☒ The default Path Selection Policy is VMW_PSP_FIXED for iSCSI or FC devices.
- ☒ VMW_PSP_MRU is typically selected for ALUA arrays by default.
- ☐ VMW_PSP_FIXED is typically selected for ALUA arrays by default.

Explanation:

Explanation : By default, the VMware NMP supports the following PSPs:

VMW_PSP_MRU

The host selects the path that it used most recently. When the path becomes unavailable, the host selects an alternative path. The host does not revert back to the original path when that path becomes available again. There is no preferred path setting with the MRU policy. MRU is the default policy for most active-passive storage devices.

Displayed in the vSphere Client as the Most Recently Used (VMware) path selection policy.

VMW_PSP_FIXED

The host uses the designated preferred path, if it has been configured. Otherwise, it selects the first working path discovered at system boot time. If you want the host to use a particular preferred path, specify it manually. Fixed is the default policy for most active-active storage devices.

Note

If the host uses a default preferred path and the path's status turns to Dead, a new path is selected as preferred. However, if you explicitly designate the preferred path, it will remain preferred even when it becomes inaccessible.

Reference: <https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.storage.doc%2FGUID-B7AD0CA0-CBE2-4DB4-A22C-AD323226A257.html>

Question 103

Which two statements regarding Distributed Resource Scheduler (DRS) affinity rules are true? (Choose two.)

- ☒ When two VM-VM affinity rules conflict, the older one takes precedence and the newer rule is disabled.
- ☐ Using Specify Failover Hosts admission control policy, VM-VM affinity rules are not supported.
- ☒ DRS gives higher precedence to preventing violations of anti-affinity rules than violations of affinity rules.
- ☐ It is not possible to create an affinity rule that conflicts with the other rules being used.

Explanation:

Explanation:

VM-VM Affinity Rule Conflicts

You can create and use multiple VM-VM affinity rules, however, this might lead to situations where the rules conflict with one another.


If two VM-VM affinity rules are in conflict, you cannot enable both. For example, if one rule keeps two virtual machines together and another rule keeps the same two virtual machines apart, you cannot enable both rules. Select one of the rules to apply and disable or remove the conflicting rule. When two VM-VM affinity rules conflict, the older one takes precedence and the newer rule is disabled. DRS only tries to satisfy enabled rules and disabled rules are ignored. DRS gives higher precedence to preventing violations of anti-affinity rules than violations of affinity rules.

Reference:

<https://pubs.vmware.com/vsphere-51/index.jsp#com.vmware.vsphere.resmgmt.doc/GUID-69C738B6-5FC8-4189-9CB5-DD90A5A05979.html>

Question 104

Refer to the Exhibit.




An administrator receives vRealize Operations alerts for the Health of the virtual machine PVMAPP_0 as shown in Exhibit 1:

◆
◆

The administrator clicks on the Analysis pane, as shown in Exhibit 2:◆

◆◆



Based on the exhibits, what action would correctly address the performance problems observed?

- ☒ Increase the number of vCPUs for PVMAPP_0.
- ☐ Increase the allocation of memory for PVMAPP_0.
- ☐ Increase the Memory limit for PVMAPP_0.
- ☐ Increase the reservation in MHz for vCPUs for PVMAPP_0.

Explanation:

Explanation:

◆ The default Hardware Version of the vRealize Operations Manager virtual machine is 7, limiting the maximum vCPUs selection to 8. You can upgrade the Hardware Version of the Analytics machine to increase the vCPU limit.

Theoretically, the performance of PVM should be similar or better than HVM. But, this trend tends to break when the vCPU count increases from 20 vCPUS. This problem is counter intuitive and is only observed in case of vCPUs being greater than 20 physical cores. We classify this problem as ◆sleepy lock anomaly, which occurs due to the usage of paravirtual interface in the ticket spinlock implementation that has been introduced to solve the ◆Lock Holder Preemption problem ◆without the architectural support (i.e. ◆Pause Loop Exiting).

Take ◆steps to analyze the scalability behavior of VMs with high core count (till 80 core). Our preliminary study suggests that besides cache contention bottleneck, the usage of ticket spinlock is another culprit for the degradation. We will take a step further and will try to analyze the scalability characteristics of these monster VMs by running other benchsuitses from Mosbench.

<https://www.vmware.com/support/pubs/vrealize-automation-pubs.htm>

Question 105

Which file determines the location of the installation script during a scripted upgrade?

- ☒ boot.cfg
- ☐ ks.cfg
- ☐ script.cfg
- ☐ upgrade.cfg

Explanation:

Explanation:

About the boot.cfg File The boot loader configuration file `boot.cfg` specifies the kernel, the kernel options, and the boot modules that the `mboot.c32` boot loader uses in an ESXi installation.

The `boot.cfg` file is provided in the ESXi installer. You can modify the `kernelopt` line of the `boot.cfg` file to specify the location of an installation script or to pass other boot options.

The `boot.cfg` file has the following syntax:

```
# boot.cfg -- mboot configuration file

#

# Any line preceded with '#' is a comment.
```

```
title=STRING
```

```
kernel=FILEPATH
```

```
kernelopt=STRING
```

```
modules=FILEPATH1 --- FILEPATH2... --- FILEPATHn
```

```
# Any other line must remain unchanged.
```

The commands in `boot.cfg` configure the boot loader.

Reference: <https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.upgrade.doc%2FGUID-1DE4EC58-8665-4F14-9AB4-1C62297D866B.html>

Question 106

After successfully adding a new ESXi 6.x host to vCenter Server, an administrator sees it as Not Responding in the vSphere Web Client interface a few minutes later.

If the issue is caused by a network firewall blocking traffic, which port must be opened to correct this specific problem?

- ☐ 443 (TCP)
- ☐ 443 (UDP)
- ☐ 902 (TCP)
- ☒ 902 (UDP)

Explanation:

Explanation:

Port: 902 (Default)

Purpose:

Host access to other hosts for migration and provisioning

Authentication traffic for ESXi and remote console traffic (xinetd/vmware-authd)

vSphere Client access to virtual machine consoles

(UDP) Status update (heartbeat) connection from ESXi to vCenter Server

Traffic Type

Incoming and outgoing TCP, outgoing UDP

Reference:

https://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vsphere.security.doc_50%2FGUID-ECEA77F5-D38E-4339-9B06-FF9B78E94B68.html

Question 107

Which two statements are true regarding iSCSI adapters? (Choose two.)

- ☒ Software iSCSI adapters require vmkernel networking.
- ☒ Independent Hardware iSCSI adapters offload processing from the ESXi host.
- ☐ Dependent Hardware iSCSI adapters do not require vmkernel networking.
- ☐ Independent Hardware iSCSI adapters require vmkernel networking.

Explanation:

Explanation : Software iSCSI Adapter

A software iSCSI adapter is a VMware code built into the VMkernel. It allows your host to connect to the iSCSI storage device through standard network adapters. The software iSCSI adapter handles

iSCSI processing while communicating with the network adapter. With the software iSCSI adapter, you can use iSCSI technology without purchasing specialized hardware.

Hardware iSCSI Adapter A hardware iSCSI adapter is a third-party adapter that offloads iSCSI and network processing from your host. Hardware iSCSI adapters are divided into categories.

Dependent Hardware iSCSI Adapter

Depends on VMware networking, and iSCSI configuration and management interfaces provided by VMware. This type of adapter can be a card that presents a standard network adapter and iSCSI offload functionality for the same port. The iSCSI offload functionality depends on the host's network configuration to obtain the IP, MAC, and other parameters used for iSCSI sessions. An example of a dependent adapter is the iSCSI licensed Broadcom 5709 NIC.

Independent Hardware iSCSI Adapter


Implements its own networking and iSCSI configuration and management interfaces. An example of an independent hardware iSCSI adapter is a card that either presents only iSCSI offload functionality or iSCSI offload functionality and standard NIC functionality. The iSCSI offload functionality has independent configuration management that assigns the IP, MAC, and other parameters used for the iSCSI sessions. An example of a independent adapter is the QLogic QLA4052 adapter.

Hardware iSCSI adapters might need to be licensed. Otherwise, they will not appear in the client or vSphere CLI. Contact your vendor for licensing information

Reference: <https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-storage-guide.pdf>

Question 108

Refer to the Exhibit.



An administrator would like to add Challenge Handshake Authentication Protocol (CHAP) to an iSCSI adapter. The administrator accesses the Storage Adapters menu as shown in the Exhibit

In which tab can the task be accomplished?

- ☒ Properties
- ☐ Advanced Options
- ☐ Targets
- ☐ Devices

Explanation:

Explanation:

Procedure

- 1 Access the iSCSI Initiator Properties dialog box.
- 2 On the General tab, click CHAP.
- 3 To configure one-way CHAP, under CHAP specify the following:
 - a. Select the CHAP security level.

Do not use CHAP unless required by target (software and dependent hardware iSCSI only)

Use CHAP unless prohibited by target

Use CHAP (software and dependent hardware iSCSI only). To configure mutual CHAP, you must select this option.

- b. Specify the CHAP name.

Make sure that the name you specify matches the name configured on the storage side.

To set the CHAP name to the iSCSI initiator name, select Use initiator name.

To set the CHAP name to anything other than the iSCSI initiator name, deselect Use initiator name and type a name in the Name text box.

c. Enter a one-way CHAP secret to be used as part of authentication. Use the same secret that you enter on the storage side.

4 To configure mutual CHAP, first configure one-way CHAP by following the directions in Step 3.

Make sure to select Use CHAP as an option for one-way CHAP. Then, specify the following under Mutual CHAP:

a. Select Use CHAP.

b. Specify the mutual CHAP name.

c. Enter the mutual CHAP secret. Make sure to use different secrets for the one-way CHAP and mutual CHAP.

5 Click OK.

6 Rescan the initiator.

Reference: <https://pubs.vmware.com/vsphere-51/index.jsp#com.vmware.vsphere.storage.doc/GUID-488A90C3-4826-4EB7-BAA4-E9C799AA2C02.html>

Question 109

An administrator is observing slow performance of the vCenter Inventory Service and observes the entries from the wrapper.log file:

◆

Exception in thread "tomcat-exec-2" java.lang.OutOfMemoryError: Java heap space
Exception in thread "http-bio-0.0.0.0-10443-Acceptor-0" java.lang.OutOfMemoryError: Java heap space

What should the administrator do to resolve the problem?

- ☒ Increase the memory resources of the vCenter Server.
- ☐ Increase the values using cloudvm-ram-size.
- ☐ Increase the memory resources of the Platform Services Controller.
- ☐ Increase the wrapper.java.maxmemory value in wrapper.conf

Explanation:

Explanation:

Cause

This issue occurs if the allocated JVM maximum memory heap is too low for the VMware Inventory service, VMware Profile-Driven Storage service, or VMware VirtualCenter Web Management service.

Several core vCenter Server services use individual configuration files (wrapper.conf) to determine the JVM maximum heap size that each service can use. The VMware Inventory service, VMware Profile-Driven Storage service, and VMware VirtualCenter Web Management service rely on the JVM maximum memory heap configured in their individual wrapper.conf file and play a critical role in the functionality of vCenter Server. ♦

Note: If vCenter Server 5.5 has recently been updated beyond 5.5 U2, see ♦Upgrading VMware vCenter Server 5.5 to 5.5 U2 or later reduces the JVM maximum memory heap size (2114669).

Resolution

This is a known issue affecting vCenter Server 5.x.

To resolve this issue, manually set the inventory size.

To change the JVM maximum memory heap size for a Windows installed vCenter Server 5.x, see, see ♦Configuring Tomcat server settings in VMware vCenter Server 5.1 and 5.5 (2021302).

To manually set the inventory size in vCenter Server Appliance 5.x:

1. Log in to vCenter Server Appliance administration interface (VAMI) URL:

https://vCenter_Server_FQDN\IP:5480

2. Click the ♦Services♦ tab.

3. Change the Inventory Size to Medium or Large, depending on your environment.

Note: The system memory may need to be increased when changing this setting. ♦For more information on environment sizing, see the Configure Inventory Size for the VMware vCenter Server Appliance♦ section of the ♦Configuring Hosts and vCenter Server♦ guide.

4. Restart the VMware Profile-Driven Storage Service and VMware Inventory Service. For more information, see ♦Stopping, starting, or restarting vCenter Server Appliance services (2054085).

Reference:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2126282

Question 110

An administrator wants to prevent a priority virtual machine that resides in a vSphere DRS cluster from migrating to other hosts in the cluster.



What configuration step will accomplish this without affecting other virtual machines?

- ☐ Set VM Overrides to Partially Automated.
- ☒ Set VM Overrides to Manual.
- ☐ Configure a VM/Host Rule with the setting Must run on hosts in group.
- ☐ Configure a VM/Host Rule with the setting Should run on hosts in group.

Explanation:

Explanation:

Option

Manual

Description

Placement and migration recommendations are displayed, but do not run until you manually apply the recommendation.

Reference: <https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.resmgmt.doc%2FGUID-E83A6B62-37F8-46D0-8D0A-D6A9181979F2.html>

Question 111

A vSphere 6 High Availability cluster has been configured with default settings. Four◆virtual machines◆in the cluster◆have◆been configured with these◆priorities:

Prod-DB. High

Prod-Email: High

Prod-VC. Medium

Dev-VDI: Low



How many VM Overrides would need to be defined at the cluster level to meet the restart priorities?

- ☐ 1
- ☒ 2
- ☐ 3
- ☐ 4

Explanation:

Explanation:


After you create a vSphere HA cluster, you can override the default cluster settings for Restart Priority and Isolation Response for specific virtual machines. Such overrides are useful for virtual machines that are used for special tasks. For example, virtual machines that provide infrastructure services like DNS or DHCP might need to be powered on before other virtual machines in the cluster.

VM Overrides the default or defined to meet the restart priorities : 2 are HIGH and medium, low respectively hence its 2 overrides at cluster level.

<https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-availability-guide.pdf>

Question 112

Refer to the Exhibit.



An administrator is reviewing a vSphere Distributed Resource Scheduler (DRS) enabled Cluster and observes unexpected behavior as shown in the Exhibit.

What are three potential causes of the cluster imbalance? (Choose three.)

- ☒ A local device is mounted to one or more virtual machines.
- ☒ DRS rules prevent virtual machines from being moved.
- ☒ vMotion is not configured and enabled.
- ☐ There are insufficient cluster resources to perform the migration.
- ☐ DRS has been configured for a conservative migration threshold.

Explanation:

Explanation:

A-) A device is mounted to one or more virtual machines preventing DRS from moving the virtual machine in order to balance the load.

B-) The migration threshold is too high.

A higher threshold makes the cluster a more likely candidate for load imbalance.

VM/VM or VM/Host DRS rules prevent virtual machines from being moved.

C-) It would be more detrimental for the virtual machine's performance to move it than for it to run where it is currently located. This may occur when loads are unstable or the migration cost is high compared to the benefit gained from moving the virtual machine.

Observe that vMotion is not enabled or set up for the hosts in the cluster, DRS does not move any virtual machines from a host. ... from this host would violate a VM/VM DRS rule or VM/Host DRS rule. <https://pubs.vmware.com/>

Question 113

An administrator is troubleshooting network communications between the vCenter Server and the ESXi 6.x host.

◆◆

Which log shows the interaction events between these components?

- ☐ /var/log/syslog.log
- ☐ /var/log/hostd.log
- ☒ /var/log/vpxa.log
- ☐ /var/log/fdm.log

Explanation:

Explanation:

When an ESXi 5.1 / 5.5 host is managed by vCenter Server 5.1 and 5.5, two components are installed, each with its own logs:

/var/log/vpxa.log: vCenter Server vpxa agent logs, including communication with vCenter Server and the Host Management hostd agent.

/var/log/fdm.log: vSphere High Availability logs, produced by the fdm service. For more information, see the vSphere HA Security section of the vSphere Availability Guide.

Reference:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2032076

Question 114

An administrator is troubleshooting a virtual machine that has unexpectedly powered off

◆◆

Which two logs should be used to troubleshoot the issue? (Choose two.)

- ☒ vmware.log
- ☒ hostd.log
- ☐ syslog.log
- ☐ shell.log

Question 115

Which three troubleshooting actions should an administrator take to address slow performance when deploying a virtual machine template? (Choose three.)

- ☒ Increase network throughput by adding additional uplinks to the vSwitch.
- ☒ Change the destination datastore or volume for the virtual machine template.
- ☒ Configure a Provisioning Traffic vmkernel port to perform the deployment operation.
- ☐ Reduce the size of the virtual machine template's virtual disk.
- ☐ Deploy the virtual machine template to the cluster and allow Distributed Resource Scheduler to register the virtual machine.

Explanation:

Explanation:

Validate that each troubleshooting step below is true for your environment. Each step will provide instructions or a link to a document, in order to eliminate possible causes and take corrective action as necessary. The steps are ordered in the most appropriate sequence to isolate the issue and identify the proper resolution. Do not skip a step.

1. Verify if the slow deployment is specific to one template or if it affects all templates. To properly test this, VMware recommends to create a brand new template and test the deployment. This provides a clean test of the environment. For more information if template deployment is slow for a single template, see [Deploying a single template is slow in vCenter Server \(1004028\)](#).
2. If you are using ESX (not applicable to ESXi), verify that no processes are over utilizing the resources on the ESX Service Console. For more information, see [Checking for resource starvation of the ESX Service Console \(1003496\)](#).
3. Verify that the network configuration on the ESX/ESXi host is optimized for the best performance, including speed and duplex settings. For more information, see [Configuring the speed and duplex of an ESX/ESXi host network adapter \(1004089\)](#).
4. Verify that the firmware on the RAID controller or HBA is up to date. For more information, see [Checking your firmware and BIOS levels to ensure compatibility with ESX/ESXi \(1037257\)](#).
5. Verify that the local storage or SAN array is configured correctly. For more information, see [Slow ESX/ESXi performance caused by misconfigured local storage or SAN array \(1006602\)](#).
6. Verify on ESX/ESXi 4.1 and above if the storage array devices in the environment support the hardware acceleration functionality and if they are responding correctly to VAAI primitives. If there is no VAAI support on the array cloning or Storage vMotion may fail at 18%. For more information see [Cloning or Storage vMotion fails at 18% with the error: Failed to clone: Connection timed out \(1029244\)](#).

◆

For further information on VAAI, see:

[vStorage APIs for Array Integration FAQ \(1021976\)](#)

Disabling the VAAI functionality in ESX/ESXi (1033665)

Note: ♦ If your problem still exists after trying the steps in this article, please:

Gather the VMware Support Script Data. ♦ For more information, see ♦ Collecting diagnostic information for VMware products (1008524).

File a support request with VMware Support and note this ♦ KB Article ID in the problem description. ♦ For more information, see ♦ How to Submit a Support Request.

Link: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1004002

Question 116

Refer to the Exhibit.

Device Details	
Properties	Paths
Owner	NMP
Partition Details	
Partition Format	GPT
▶ Primary Partitions	1
▶ Logical Partitions	0
Multipathing Policies	
▶ Path Selection Policy	Fixed (VMware)
Storage Array Type Policy	VMW_SATP_EQL

An administrator wishes to provide Load Balanced I/O for the device shown in the Exhibit.



To meet this requirement, which setting should be changed?

- ☐ Storage Array Type Policy = VMW_NMP_RR
- ☒ Path Selection Policy = Round Robin (VMware)
- ☐ Storage Array Type Policy = VMW_SATP_RR
- ☐ Path Selection Policy = MRU (VMware)

Explanation:

Explanation: Round Robin (RR): The VMW_PSP_RR policy uses an automatic path selection, rotating through all available paths, enabling the distribution of load across the configured paths.

For Active/Passive storage arrays, only the paths to the active controller will be used in the Round Robin policy.

For Active/Active storage arrays, all paths will be used in the Round Robin policy.

Note: For logical units associated with Microsoft Cluster Service (MSCS) and Microsoft Failover Clustering virtual machines, the Round Robin pathing policy is supported only on ESXi 5.5 and later.

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=display&externalId=1011340

Question 117


An administrator has recently configured HA on a cluster. After reviewing the summary tab on one of the hosts, the warning in Exhibit 1 is displayed:



This host currently has no management network redundancy



The administrator proceeds to view the management network port group data shown in Exhibit 2:



The administrator then views the management network vSwitch as shown in Exhibit 3:



Based on the exhibits, which two steps should be taken to ensure redundancy on the management network? (Choose two.)

- ☒ Move vmnic1 to Standby adapters.
- ☐ Add an additional vmknic to the Network Adapters and move it to Active adapters.
- ☐ Set the advanced HA configuration parameter das.ignoreRedundantNetWarning to True.
- ☒ Uncheck the Override Failover Checkbox on the management network port group.

Explanation:

Explanation: A and D

Use Load Balancing and Failover policies to determine how network traffic is distributed between adapters and how to reroute traffic in the event of an adapter failure.

The Failover and Load Balancing policies include the following parameters:

Load Balancing policy: The Load Balancing policy determines how outgoing traffic is distributed among the network adapters assigned to a standard switch. Incoming traffic is controlled by the Load Balancing policy on the physical switch.

Failover Detection: Link Status/Beacon Probing

Network Adapter Order (Active/Standby)

In some cases, you might lose standard switch connectivity when a failover or failback event occurs. This causes the MAC addresses used by virtual machines associated with that standard switch to appear on a different switch port than they previously did. To avoid this problem, put your physical switch in portfast or portfast trunk mode.

Link:

<https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.networking.doc%2FGUID-D5EA6315-5DCD-463E-A701-B3D8D9250FB5.html>

Question 118

A virtual machine is exhibiting these symptoms: ♦♦

Memory usage is constantly high (94% or greater) or constantly low (24% or less).

Free memory is consistently 6% or less and swapping frequently occurs

Which three solutions could correct this problem? (Choose three.)

- ☒ Verify that VMware Tools is installed on each virtual machine.
- ☒ Decrease the memory reservation setting, if higher than active memory.
- ☒ Add physical memory to the host.
- ☐ Disable the balloon driver in each virtual machine.
- ☐ Create a memory limit for each virtual machine.

Explanation:

Explanation:

Problem

Memory usage is constantly high (94% or greater) or constantly low (24% or less).

Free memory consistently is 6% or less and swapping frequently occurs.

Cause

The host probably is lacking the memory required to meet the demand. The active memory size is the same as the granted memory size, which results in memory resources that are not sufficient for the workload. Granted memory is too much if the active memory is constantly low.

Host machine memory resources are not enough to meet the demand, which leads to memory reclamation and degraded performance.

The active memory size is the same as the granted memory size, which results in memory resources that are not sufficient for the workload.

Solution

Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.

Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.

Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.

If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.

Migrate one or more virtual machines to a host in a DRS cluster.

Add physical memory to the host.

Reference:

<https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.monitoring.doc%2FGUID-115861E6-810A-43BB-8CDB-EE99CF8F3250.html>

Question 119

When attempting to connect to a vCenter Server, an administrator observes the following at the top of the vSphere Web Client:



Could not connect to one or more vCenter Server Systems: <https://vCenter.corp.com:443/sdk>



What three reasons could be preventing the vSphere Web Client from communicating with this vCenter Server? (Choose three.)

- ☒ The vCenter Server machine is not responding via the network.
- ☒ An incorrect entry for this vCenter Server exists in the Single Sign-On service.
- ☒ The SSL certificates do not match the FQDN address for the server.
- ☐ The Platform Services Controller is external to this vCenter Server.
- ☐ The DNS entry for the vCenter Server is incorrect.

Explanation:

Explanation: This issue occurs in these situations:

During the re-installation of vCenter Server, it is possible to have the same vCenter Server registered more than once to Single Sign-On (SSO).

With a previous install of vCenter Server, SSL certificates are not overwritten or removed properly during an upgrade or re-installation.


Note: If there are previous issues with the certificates, they may not be exposed until the installation and use of the 5.1 Web Client.

Before proceeding with the steps in the Resolution section, ensure that you are not experiencing the issue identified in vSphere Web Client 5.1 reports this SSL warning after an installation or upgrade: Failed to verify the SSL certificate for one or more vCenter Server Systems (2036505).

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2050273

Question 120

Refer to the Exhibit.



Which tab shows the Hardware Acceleration support status?

- ☒ Devices
- ☐ Properties
- ☐ Paths
- ☐ Advanced Options

Explanation:

Explanation:

To know the detailed info about the tab:

If you go to **Host > Configuration > Storage**, you can see the Hardware Acceleration Status in the panel on the right side.

For each storage device and datastore, the vSphere Client displays the hardware acceleration support status in the Hardware Acceleration column of the Devices view and the Datastores view.

The status values are **Unknown**, **Supported**, and **Not Supported**. The initial value is **Unknown**. The status changes to **Supported** after the host successfully performs the offload basic operations. If the offload operation fails, the status changes to **Not Supported**.

Check link KB: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1021976

Question 121

An administrator wants to upgrade to vCenter Server 6.x.



The vCenter Server:

Is hosted on a virtual machine server running Microsoft Windows Server 2008 R2, with 8 vCPUs and 16GB RAM.

Will have an embedded Platform Services Controller.

Hosts a Large Environment with 1,000 ESXi hosts and 10,000 Virtual Machines.



Why does the vCenter Server not meet the minimum requirements?

- ☐ Windows Server 2008 R2 is not a supported Operating System for vCenter Server.
- ☒ The virtual machine has insufficient resources for the environment size.
- ☐ The environment is too large to be managed by a single vCenter Server.
- ☐ The Platform Services Controller must be changed to an External deployment.

Explanation:

Explanation:

vCenter Server for Windows Hardware Requirements

When you install vCenter Server on a virtual machine or physical server running Microsoft Windows server, your system must meet specific hardware requirements.

<https://kb.vmware.com/kb/2107948>

VMware vCenter Server 6.0 Deployment Guide - White Paper ...

<https://www.vmware.com/files/pdf/.../vmware-vcenter-server6-deployment-guide.pdf>

Question 122

Refer to the Exhibit.

The number of vSphere HA heartbeat datastores for this host is 1, which is less than required: 2

An administrator receives the error message shown in the Exhibit.

Which two actions can be taken to clear the warning? (Choose two.)

- ☐ Add a Virtual SAN datastore and configure it for High Availability heartbeating.
- ☐ Set the advanced High Availability parameter Das.heartbeatdsperhost to 1.
- ☒ Set the advanced High Availability parameter Das.ignoreInsufficientHbDatastore to true.
- ☒ Add a shared datastore and reconfigure High Availability

Explanation:

Explanation:

To disable the HA error message:

Log in to vCenter Server.

Right-click the cluster and click Edit Settings.

Click vSphere HA > Advanced Options.

Under Option, add an entry for das.ignoreInsufficientHbDatastore.

Under Value, type true.

Click Cluster Features.

De-select Turn on vSphere HA and click OK.

Wait for all the hosts in the cluster to deconfigure HA, then right-click the cluster and click Edit Settings.

Click Cluster Features.

Click Turn on vSphere HA.

Click OK.

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2004739

Question 123

An administrator is troubleshooting a CPU performance issue for a virtual machine.



Which three esxrtop counters may demonstrate CPU contention? (Choose three.)

- ☒ %RDY
- ☐ %RUN
- ☒ %MLMTD
- ☐ %WAIT
- ☒ %CSTP

Explanation:

Explanation:

"%RDY"

The percentage of time the world was ready to run.



A world in a run queue is waiting for CPU scheduler to let it run on a PCPU. %RDY accounts the percentage of this time. So, it is always smaller than 100%.



+Q: How do I know CPU resource is under contention?+

+A: %RDY is a main indicator. But, it is not sufficient by itself.+



+If a "CPU Limit" is set to a VM's resource settings, the VM will be deliberately held from scheduled to a PCPU when it uses up its allocated CPU resource. This may happen even when there is plenty of free CPU cycles. This time deliberately held by scheduler is shown by "%MLMTD", which will be describe next. Note that %RDY includes %MLMTD. For, for CPU contention, we will use "%RDY - %MLMTD". So, if "%RDY - %MLMTD" is high, e.g., larger than 20%, you may experience CPU contention.+



+What is the recommended threshold? Well, it depends. As a try, we could start with 20%. If your application speed in the VM is OK, you may tolerate higher threshold. Otherwise, lower.+



+Q: How do we break down 100% for the world state times?+

+A: A world can be in different states, either scheduled to run, ready to run but not scheduled, or not ready to run (waiting for some events).+

100% = %RUN + %READY + %CSTP + %WAIT.

+Check the description of %CSTP and %WAIT below.+



+Q: What does it mean if %RDY of a VM is high?+

+A: It means the VM is possibly under resource contention. Check "%MLMTD" as well. If "%MLMTD" is high, you may raise the "CPU limit" setting for the VM. If "%RDY - %MLMTD" is high, the VM is under CPU contention.+



"%MLMTD"

The percentage of time the world was ready to run but deliberately wasn't scheduled because that would violate the "CPU limit" settings.



Note that %MLMTD is included in %RDY.



+Q: What does it mean if %MLMTD of a VM is high?+

+A: The VM cannot run because of the "CPU limit" setting. If you want to improve the performance of this VM, you may increase its limit. However, keep in mind that it may reduce the performance of others.+



"%CSTP"

The percentage of time the world spent in ready, co-deschedule state. This co-deschedule state is only meaningful for SMP VMs. Roughly speaking, ESX CPU scheduler deliberately puts a VCPU in this state, if this VCPU advances much farther than other VCPUs.



+Q: What does it mean if %CSTP is high?+

+A: It usually means the VM workload does not use VCPUs in a balanced fashion. The VCPU with high %CSTP is used much more often than the others. Do you really need all those VCPUs? Do you pin the guest application to the VCPUs?+

Reference:

<https://communities.vmware.com/docs/DOC-9279>

Question 124

An administrator needs two vCenter Servers to be visible within a single vSphere Web Client session.

Which two vCenter Server and Platform Services Controller (PSC) configurations would accomplish this? (Choose two.)

- ☐ Install a single PSC with two vCenter Servers registered to it.
- ☒ Install two PSCs in the same Single Sign-On domain with one vCenter Server registered to each PSC.
- ☒ Install a single PSC with two vCenter Servers registered to it and configure Linked Mode.
- ☐ Install two PSCs in the same Single Sign-On domain with one vCenter Server registered to each PSC and configure Linked Mode.

Question 125

Which two statements are true regarding VMFS3 volumes in ESXi 6.x? (Choose two.)

- ☒ Creation of VMFS3 volumes is unsupported.
- ☒ Upgrading VMFS3 volumes to VMFS5 is supported.
- ☐ Existing VMFS3 volumes are unsupported.
- ☐ Upgrading VMFS3 volumes to VMFS5 is unsupported.

Explanation:

Explanation:

Understanding VMFS Datastores

To store virtual disks, ESXi uses datastores, which are logical containers that hide specifics of physical storage from virtual machines and provide a uniform model for storing virtual machine files. Datastores that you deploy on block storage devices use the vSphere VMFS format, a special high-performance file system format that is optimized for storing virtual machines.

Several versions of the VMFS file system have been released since its introduction. The following table shows host-to-VMFS version relationships.

Table 16-1. Host access to VMFS version

VMFS	ESX/ESXi 3.x host	ESX/ESXi 4.x host	ESXi 5.x host	ESXi 6.0 host
VMFS2	RO	RO	N	N
VMFS3	RW	RW	RW	RW
				NOTE You can continue to use existing VMFS3 datastores, but you cannot create new ones. If you have existing VMFS3 datastores, upgrade them to VMFS5.
VMFS5	N	N	RW	RW

- RW: Complete read and write support. You can create and power on virtual machines.
- RO: Read only support. You cannot create or power on virtual machines.
- N: No access. ESXi 5.x and later hosts do not support VMFS2. If your datastore was formatted with VMFS2, first upgrade the datastore to VMFS3 using legacy hosts.

Use the vSphere Web Client to set up a VMFS datastore in advance on a block-based storage device that your ESXi host discovers. A VMFS datastore can be extended to span several physical storage extents, including SAN LUNs and local storage. This feature allows you to pool storage and gives you flexibility in creating the datastore necessary for your virtual machines.

NOTE Pooling ATS-capable hardware creates a spanned VMFS datastore that can use ATS-only locking mechanism. If any device is not ATS-capable, the datastore cannot be ATS-only, but uses ATS+SCSI locking.

You can increase the capacity of a datastore while virtual machines are running on the datastore. This ability lets you add new space to your VMFS datastores as your virtual machine requires it. VMFS is designed for concurrent access from multiple physical machines and enforces the appropriate access controls on virtual machine files.

Reference: <https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-storage-guide.pdf>

Question 126

An administrator is upgrading a vCenter Server Appliance and wants to ensure that all the prerequisites are met.

◆◆

What action must be taken before upgrading the vCenter Server Appliance?

- ☒ Install the Client Integration Plug-in.
- ☐ Install the database client.
- ☐ Install the ODBC connector.
- ☐ Install the Update Manager Plug-in.

Explanation:

Explanation:

Prerequisites

- ? Verify that the clocks of all machines on the vSphere network are synchronized. See [Synchronizing Clocks on the vSphere Network](#).
- ? Verify that the target ESXi host on which you deploy the vCenter Server Appliance is not in lockdown or maintenance mode.
- ? Verify that you have sufficient free disk space on the vCenter Server Appliance that you want to upgrade to accommodate the data for the upgrade.
- ? Verify that port 22 is open on the vCenter Server Appliance that you want to upgrade. The upgrade process establishes an inbound SSH connection to download the exported data from existing appliance.
- ? Verify that port 443 is open on the source ESXi host on which the vCenter Server Appliance that you want to upgrade resides. The upgrade process establishes an HTTPS connection to the source ESXi host to verify that the vCenter Server Appliance is ready for upgrade and to set up an SSH connection between the new and the existing appliance.
- ? Verify that the [vCenter Server SSL certificate](#) for your existing vCenter Server Appliance is configured correctly. See VMware Knowledge Base article [2057223](#).
- ? If you use an external database, back up the vCenter Server Appliance database.
- ? Create a snapshot of the vCenter Server Appliance that you want to upgrade.
- ? Install the new version of the Client Integration Plug-In. See [Install the Client Integration Plug-In](#).

Reference: <https://pubs.vmware.com/vsphere60/index.jsp?topic=%2Fcom.vmware.vsphere.upgrade.doc%2FGUID-6A5C596D-103E-4024-9353-5569263EB427.html>

Question 127

Which three traffic types can be configured for dedicated VMkernel adapters? (Choose three.)

- ☐ Discovery traffic
- ☒ vMotion traffic
- ☒ vSphere Replication NFC traffic
- ☒ Provisioning traffic
- ☐ vSphere Custom traffic

Explanation:

Explanation:

Securing System Traffic

Take appropriate security measures to prevent unauthorized access to the management and system traffic in your vSphere environment. For example, isolate the vMotion traffic in a separate network that includes only the ESXi hosts that participate in the migration. Isolate the management traffic in a network that only network and security administrators are able to access. For more information, see [vSphere Security](#) and [vSphere Installation and Setup](#).

System Traffic Types

You should dedicate a separate VMkernel adapter for every traffic type. For distributed switches, dedicate a separate distributed port group for each VMkernel adapter.

Management traffic

Carries the configuration and management communication for ESXi hosts, vCenter Server, and host-to-host High Availability traffic. By default, when you install the ESXi software, a vSphere Standard switch is created on the host together with a VMkernel adapter for management traffic. To provide redundancy, you can connect two or more physical NICs to a VMkernel adapter for management traffic.

vMotion traffic

Accommodates vMotion. A VMkernel adapter for vMotion is required both on the source and the target hosts. The VMkernel adapters for vMotion should handle only the vMotion traffic. For better performance, you can configure multiple NIC vMotion. To have multi NIC vMotion, you can dedicate two or more port groups to the vMotion traffic, respectively every port group must have a vMotion VMkernel adapter associated with it. Then you can connect one or more physical NICs to every port group. In this way, multiple physical NICs are used for vMotion, which results in greater bandwidth.

Note

vMotion network traffic is not encrypted. You should provision secure private networks for use by vMotion only.

Provisioning traffic

Handles the data that is transferred for virtual machine cold migration, cloning, and snapshot creation.

IP storage traffic and discovery

Handles the connection for storage types that use standard TCP/IP networks and depend on the VMkernel networking. Such storage types are software iSCSI, depended hardware iSCSI, and NFS. If you have two or more physical NICs for iSCSI, you can configure iSCSI multipathing. ESXi hosts support only NFS version 3 over TCP/IP. To configure a software FCoE (Fibre Channel over Ethernet) adapter, you must have a dedicated VMkernel adapter. Software FCoE passes configuration information

though the Data Center Bridging Exchange (DCBX) protocol by using the Cisco Discovery Protocol (CDP) VMkernel module.

Fault Tolerance traffic

Handles the data that the primary fault tolerant virtual machine sends to the secondary fault tolerant virtual machine over the VMkernel networking layer. A separate VMkernel adapter for Fault Tolerance logging is required on every host that is part of a vSphere HA cluster.

vSphere Replication traffic

Handles the outgoing replication data that the source ESXi host transfers to the vSphere Replication server. Dedicate a VMkernel adapter on the source site to isolate the outgoing replication traffic.

vSphere Replication NFC traffic

Handles the incoming replication data on the target replication site.

Virtual SAN traffic

Every host that participates in a Virtual SAN cluster must have a VMkernel adapter to handle the Virtual SAN traffic.

Reference:

<https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.networking.doc%2FGUID-D4191320-209E-4CB5-A709-C8741E713348.html>

Question 128

Which two are valid compliance results that indicate the need to apply a Host Profile? (Choose two.)

- ☒ Non-compliant
- ☐ Inconsistent
- ☒ Unknown
- ☐ Disconnected

Explanation:

Explanation: Check Compliance You can confirm the compliance of a host or cluster to its attached Host Profile and determine which, if any, configuration parameters on a host are different from those specified in the Host Profile.

Procedure

1. Navigate to a Host Profile. The Objects tab lists all Host Profiles, the number of hosts attached to that Host Profile, and summarized results of the last compliance check.


2. Click the Check Host Profile Compliance icon ().

In the Objects tab, the compliance status is updated as Compliant, Unknown, or Non-compliant. A non-compliant status indicates a discovered and specific inconsistency between the profile and the host. To resolve this, you should remediate the host. And unknown status indicates that the compliance of the host could not be verified; to resolve the issue, remediate the host through the Host Profile.

Reference: <https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-host-profiles-guide.pdf>

Question 129

Refer to the Exhibit.



-- Exhibit --

An administrator has created the DRS cluster shown in the Exhibit.



Based on the exhibit, which statement is true?

- ☒ Under CPU contention, Prod-VM1 receives four times the CPU resources than Test-VM1.
- ☐ The Prod-VM1 will always have more CPU resources than all other virtual machines.
- ☐ The Test-VM2 will always have less CPU resources than all other virtual machines.
- ☐ Under CPU contention, Test-VM1 will receive 25% of the total CPU resources.

Explanation:

Explanation:

To understand limits and theory check the information given below:

CPU is calculated like this: [Cluster CPU Cores] * [2,000 for High, 1,000 for Normal, 500 for Low]

For an example, Our DRS cluster has 100 CPU cores (see the blue section) and thus the math is: $100 * 2,000 = 200,000$ for High and $100 * 500 = 50,000$ for Low.

That's why PROD-VM1 : CPU shares

High. https://www.vmware.com/files/pdf/drs_performance_best_practices_wp.pdf

Question 130

What command line utility can be used to upgrade an ESXi host?

- ☒ esxcli
- ☐ esxupdate
- ☐ vihostupdate
- ☐ esxcfg

Explanation:

Explanation:

ESXi Command-line interface - esxcli Using the vSphere Command Line Interface (CLI), you can upgrade ESXi 5.x hosts to version 6.0.

For more information about upgrading ESXi 5.x hosts using esxcli commands, see the Upgrading Hosts by Using esxcli Commands

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2109711

Question 131

Which condition would cause a vCenter Server installation to fail when installing on a Windows virtual machine?

- ☐ The virtual machine does not have at least four vCPUs.
- ☒ The virtual machine is running Windows Server 2008.
- ☐ The virtual machine has an E1000 network device.
- ☐ The virtual machine does not have 16GB of RAM.

Explanation:

Explanation:

The vCenter Server installation or uninstallation process might fail or stop responding on Windows Server 2008 R1 SP2

The Windows Installer on Windows Server 2008 R1 SP2 has issues with handling multiple packages. Because of these issues the vCenter Server installation might stop responding, or if you attempt to install and uninstall vCenter Server a few times, the process might fail.

Workaround: Apply the patch from [Microsoft KB 981669](#) which addresses the Windows Installer issue.

Reference: <https://www.vmware.com/support/vsphere6/doc/vsphere-esxi-vcenter-server-60-release-notes.html#compatibility>

Question 132

An administrator suspects that the MTU value for a vSphere Standard Switch is misconfigured

Which two commands can determine the value? (Choose two.)

- ☒ esxcfg-vswitch -l
- ☒ esxcli network vswitch standard list
- ☐ esxcfg-vss -l
- ☐ esxcli network standard vswitch list

Explanation:

Explanation: To display a list of vSwitches on the host, and to check that the configuration of the vSwitch is correct:

Run this command for ESX 3.5 and ESXi/ESX 4.x:

```
# esxcfg-vswitch -l
```

Run this command for ESXi 5.0:

```
# esxcli network vswitch standard list
```

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1007654

Question 133

An administrator notices that 8 out of 10 virtual machines have memory ballooning and swapping. However, virtual machine 9 is not ballooning or swapping and virtual machine 10 is not ballooning but is swapping.

◆

Which two statements explain the behavior of virtual machine 9 and virtual machine 10? (Choose two.)

- ☒ Virtual machine 9 has a 100% memory reservation.
- ☐ Virtual machine 10 has a memory limit configured.
- ☐ Virtual machine 9 has memory shares set to HIGH.
- ☒ Virtual machine 10 does not have VMware Tools enabled or installed.

Explanation:

Explanation:

If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot handle the demand for memory. This leads to memory reclamation, which might degrade performance.

If the host has enough free memory, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

D-) Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. If VMware tools are not installed/outdated showing on summary tab Ballooning will not work hence swapping will work because it has nothing to do with ballooning.

Fact: If memory is not available the memory can be mapped to the .vswp file on a vmfs or nfs datastore. The virtual machine has no knowledge if the memory is mapped to physical memory or to a disk. This is called hypervisor swapping, and this is the last resort for the vmkernel to use this mechanism. (Internal VMware KB article but for external you may check VMBlog or external links: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1003470)

Question 134

An administrator plans to use VMware Converter Standalone to convert a virtual machine to a vSphere environment. The source virtual machine has these properties:

Running Windows 2008 R2

Contains one NTFS formatted volume

◆

During conversion, how many virtual disks can the administrator add to the destination virtual machine?

- ☒ 0
- ☐ 1
- ☐ 2
- ☐ 3

Explanation:

Explanation:

After adding virtual disk on Windows 2008 and Windows 2008 R2, you cannot make disk online

Windows 2008 and Windows 2008 R2 virtual machine is configured with Virtual Hardware version 7

You see this error within Windows under Disk Management:

The disk is offline because of policy set by an administrator so the virtual disk will show 0.

Make disk online how follow the KB:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2000767

Question 135

Refer to the Exhibit.

This host currently has no management network redundancy

An administrator is changing the settings on a vSphere Distributed Switch (vDS). During this process, the ESXi Management IP address is set to an address which can no longer communicate with the vCenter Server.

◆◆

What is the most likely outcome of this action?

- ☐ The host will disconnect from the vCenter Server and remain disconnected.
- ☒ The host will automatically detect the communication issue and revert the change.
- ☐ The host will stay connected with the change, but show an alert.
- ☐ The host will disconnect and migrate the vDS portgroup to a standard switch.

Explanation:

Explanation: Procedure

1. Connect to the DCUI of the host.
2. From the Network Restore Options menu, select Restore vDS.
3. Configure the uplinks and optionally the VLAN for the management network.
4. Apply the configuration.

The DCUI creates a local ephemeral port and applies the values you provided for the VLAN and uplinks. The DCUI moves the VMkernel adapter for the management network to the new local port to restore connectivity to vCenter Server.

Next Steps:

After the connection of the host to vCenter Server is restored, correct the configuration of the distributed port group and re-add the VMkernel adapter to the group.

For information about networking rollback, recovery, and restore, see the [vSphere Networking documentation](#): Wait until vCenter Server applies the settings within the next 24 hours.

Question 136

An administrator wants to clone a virtual machine using the vSphere Client.

◆

Which explains why the Clone option is missing?

- ☒ The vSphere Client is directly connected to the ESXi host.◆
- ☐ The virtual machine is configured with a thin-provisioned virtual disk.
- ☐ The virtual machine is configured with outdated Virtual Hardware.
- ☐ Cloning can only be performed with vRealize Orchestrator

Explanation:

Explanation:

Clone a Virtual Machine in the vSphere Client

Cloning a virtual machine creates a duplicate of the virtual machine with the same configuration and installed software as the original.

Optionally, you can customize the guest operating system of the clone to change the virtual machine name, network settings, and other properties. This prevents conflicts that can occur if a virtual machine and a clone with identical guest operating system settings are deployed simultaneously.

Prerequisites

You must be connected to vCenter Server in order to clone a virtual machine. You cannot clone virtual machines if you connect directly to an ESXi host.

Reference:

https://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vsphere.vm_admin.doc_50%2FGUID-5C504B67-CDB3-42FC-8B3B-737201A725DD.html

Question 137

What are two use cases for Fibre Channel Zoning in a vSphere environment? (Choose two.)

- ☐ Increases the number of targets presented to an ESXi host.
- ☒ Controls and isolates paths in a fabric.
- ☐ Controls and isolates paths to an NFS share.
- ☒ Can be used to separate different environments.

Explanation:

Explanation : Using Zoning with Fibre Channel SANs

Zoning provides access control in the SAN topology. Zoning defines which HBAs can connect to which targets. When you configure a SAN by using zoning, the devices outside a zone are not visible to the devices inside the zone.

Zoning has the following effects:

Reduces the number of targets and LUNs presented to a host.

Controls and isolates paths in a fabric.

Can prevent non-ESXi systems from accessing a particular storage system, and from possibly destroying VMFS data.

Can be used to separate different environments, for example, a test from a production environment.

Reference: <https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-storage-guide.pdf>

Question 138

An ESXi 6.x host in the vCenter Server inventory has disconnected due to an All Paths Down (APD) situation. An administrator has corrected the APD issue on the host, but it still remains disconnected.

◆◆

What action should the administrator take next?

- ☒ Select Restart Management Agents from the DCUI.
- ☐ Execute esxcli system settings advanced set -d /Scsi/FailVMIOonAPD.
- ☐ Modify the advanced parameter /Disk/ApdTokenRetryCount
- ☐ Enable the advanced parameter /Misc/APDHandlingEnable.

Explanation:

Explanation:

All-Paths-Down (APD)

A datastore is shown as unavailable in the Storage view

A storage adapter indicates the Operational State of the device as Dead or Error

All paths to the device are marked as Dead

You are unable to connect directly to the ESXi host using the vSphere Client

The ESXi host shows as Disconnected in vCenter Server

In the `/var/log/vmkernel.log` file, you see entries similar to:

```
cpu1:2049)WARNING: NMP: nmp_IssueCommandToDevice:2954:I/O could not be issued to device
"naa.60a98000572d54724a34642d71325763" due to Not found
```

```
cpu1:2049)WARNING: NMP: nmp_DeviceRetryCommand:133:Device
"naa.60a98000572d54724a34642d71325763": awaiting fast path state update for failover with I/O
blocked. No prior reservation exists on the device.
```

```
cpu1:2049)WARNING: NMP: nmp_DeviceStartLoop:721:NMP Device
"naa.60a98000572d54724a34642d71325763" is blocked. Not starting I/O from device.
```

```
cpu1:2642)WARNING: NMP: nmpDeviceAttemptFailover:599:Retry world failover device
"naa.60a98000572d54724a34642d71325763" - issuing command 0x4124007ba7c0
```

```
cpu1:2642)WARNING: NMP: nmpDeviceAttemptFailover:658:Retry world failover device
"naa.60a98000572d54724a34642d71325763" - failed to issue command due to Not found (APD), try
again...
```

```
cpu1:2642)WARNING: NMP: nmpDeviceAttemptFailover:708:Logical device
"naa.60a98000572d54724a34642d71325763": awaiting fast path state update...
```

```
cpu0:2642)WARNING: NMP: nmpDeviceAttemptFailover:599:Retry world failover device
"naa.60a98000572d54724a34642d71325763" - issuing command 0x4124007ba7c0
```

cpu0:2642)WARNING: NMP: nmpDeviceAttemptFailover:658:Retry world failover device
 "naa.60a98000572d54724a34642d71325763" - failed to issue command due to Not found (APD), try
 again...

cpu0:2642)WARNING: NMP: nmpDeviceAttemptFailover:708:Logical device
 "naa.60a98000572d54724a34642d71325763": awaiting fast path state update...

A restart of the management agents may show these errors:

Not all VMFS volumes were updated; the error encountered was 'No connection'.

Errors:

Rescan complete, however some dead paths were not removed because they were in use by the
 system. Please use the 'storage core device world list' command to see the VMkernel worlds still
 using these paths.

Error while scanning interfaces, unable to continue. Error was Not all VMFS volumes were updated;
 the error encountered was 'No connection'.

You may also see that the device is no longer listed:

cpu17:10107)WARNING: Vol3: 1717: Failed to refresh FS 4beb089b-68037158-2ecc-00215eda1af6
 descriptor: Device is permanently unavailable

cpu17:10107)ScsiDeviceIO: 2316: Cmd(0x412442939bc0) 0x28, CmdSN 0x367bb6 from world 10107
 to dev "eui.00173800084f0005" failed H:0x1 D:0x0 P:0x0 Possible sense data: 0x0 0x0 0x0.

cpu17:10107)Vol3: 1767: Error refreshing PB resMeta: Device is permanently unavailable

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2004684

Question 139

An administrator is troubleshooting a CPU issue for a virtual machine. The following is seen in
 esxtop:

CPU0 is at 100% usage

The remaining logical CPUs are close to 0%

%RDY value is consistently above 10%



What is likely the cause of the CPU issue?

- ☐ The virtual machine has a CPU limit configured.
- ☐ The virtual machine's guest operating system is configured for SMP
- ☒ The virtual machine has CPU affinity configured.
- ☐ The virtual machine is configured with a CPU reservation.

Explanation:

Explanation:

Using CPU Affinity

By specifying a CPU affinity setting for each virtual machine, you can restrict the assignment of virtual machines to a subset of the available processors in multiprocessor systems. By using this feature, you can assign each virtual machine to processors in the specified affinity set. CPU affinity specifies virtual machine-to-processor placement constraints and is different from the relationship created by a VM-VM or VM-Host affinity rule, which specifies virtual machine-to-virtual machine host placement constraints.

Reference:

https://pubs.vmware.com/vsphere-4-esx-vcenter/index.jsp?topic=/com.vmware.vsphere.resource management.doc_41/managing_cpu_resources/c_using_cpu_affinity.html

Question 140

An administrator enables High Availability (HA) on a Virtual SAN cluster

There are four vmkernel port groups with the following IP addresses assigned:

Management: 192.168.12.10

vMotion: 192.168.13.10

Virtual SAN: 192.168.14.10

Fault Tolerance: 192.168.15.10

Which IP address will HA use for traffic?

- ☐ 192.168.12.10
- ☐ 192.168.13.10
- ☒ 192.168.14.10
- ☐ 192.168.15.10

Explanation:

Explanation:

Give the VMKernel port a label (e.g. iSCSI - if it will purely be used for iSCSI).

Enter the VLAN ID if the pNICs connected to the vSwitch are trunk ports and you have a specific VLAN for this traffic (e.g. a dedicated storage VLAN).

Now if the port will only be used for storage services (iSCSI and NFS).

If the port will be used for all VMkernel services (iSCSI, NFS, vMotion and FT), tick the additional boxes to allow this port to be used also for those services,

Assigned the FT (IP Address)

Clustering is configured to give you, the administrator of an environment, a form of fault tolerance, and VMware has taken this concept to a whole other level. Although VMware does not call FT clustering, it functions the same in that FT will failover the primary virtual machine to a secondary virtual machine. VMware Fault Tolerance (FT) is based on vLockstep technology and provides zero downtime, zero data loss, and continuous availability for your applications.

Next,

- Always assign a reasonably high relative share for the Fault Tolerance resource pool because FT is a very latency-sensitive traffic type.

Question 141

A task fails while creating a VMFS5 datastore on a disk with these characteristics:❖

Was previously used by a Linux server

Was not erased

Is visible with the vSphere Web Client

Which action can be performed to resolve the issue?

- ☒ Delete the partitions on the disk manually with partedUtil first.
- ☐ Create a VMFS3 file system first, then upgrade it.
- ☐ Create the VMFS5 file system manually using vmkfstools.
- ☐ Delete the data with the vmkfstools command.

Explanation:

Explanation: Please refer to link to understand more about it:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2046610

Question 142

An administrator is deploying ESXi❖❖6.x❖❖hosts using Auto Deploy and wants the image profile to be available, even after closing and opening a new PowerCLI window.

❖❖

Which command can be used to ensure that image profiles are preserved across PowerCLI sessions?

- ☐ Set-EsxImageProfile
- ☒ Export-EsxImageProfile
- ☐ Save-EsxImageProfile
- ☐ Preserve-EsxImageProfile

Explanation:

Explanation: Detailed Description

Exports an Image Profile object as either an ESXi ISO image that can be booted up and used as an ESXi installer, or an offline depot ZIP file that contains metadata plus the VIB packages. In both cases, downloads the VIB binaries and validates the VIB signatures.💎

You can perform the following tasks with the offline depot ZIP file:💎

- *Import the ZIP into VMware Update Manager for patch remediation💎
- *Download the ZIP to an ESXi host and used with esxcli for installation💎
- *Re-import the ZIP into ImageBuilder itself using Add-EsxSoftwareDepot💎

You can specify either -ExportToIso or -ExportToBundle but not both.

NAME

ImageProfile

TYPE

ImageProfile

DESCRIPTION

(pipeline input, prompt) Specifies the image profile to export.

Takes one of the following forms:

- *Name of an image profile, as displayed in the Name column of Get-EsxImageProfile, or the Name property of any ImageProfile object
- *ImageProfile object

REQUIRED?

True

PIPELINE INPUT

true (ByValue, ByPropertyName)

Reference: <https://www.vmware.com/support/developer/PowerCLI/PowerCLI501/html/Export-EsxImageProfile.html>

Question 143

When attempting to place a datastore in Maintenance Mode, the task remains at 1%.

What are two potential causes for this? (Choose two.)

- ☒ Storage DRS is disabled on one or more virtual machine(s) disk(s).
- ☒ Storage DRS rules prevent migration recommendations for the disk
- ☐ Storage DRS datastores have insufficient space to accommodate failover.
- ☐ Storage DRS provisioning network has been limited with I/O control policies.

Explanation:

Explanation:

Datastore Cannot Enter Maintenance Mode

You place a datastore in maintenance mode when you must take it out of usage to service it. A datastore enters or leaves maintenance mode only as a result of a user request.

Problem

A datastore in a datastore cluster cannot enter maintenance mode. The Entering Maintenance Mode status remains at 1%.

Cause

One or more disks on the datastore cannot be migrated with Storage vMotion. This condition can occur in the following instances.

Storage DRS is disabled on the disk.

Storage DRS rules prevent Storage DRS from making migration recommendations for the disk.

Solution

If Storage DRS is disabled, enable it or determine why it is disabled. See Storage DRS is Disabled on a Virtual Disk for reasons why Storage DRS might be disabled.

If Storage DRS rules are preventing Storage DRS from making migration recommendations, you can remove or disable particular rules.

- a In the vSphere Client inventory, right-click the datastore cluster and select Edit Settings.
- b Select Rules and click the rule.
- c Click Remove.
- d Click OK.

Alternatively, if Storage DRS rules are preventing Storage DRS from making migration recommendations, you can set the Storage DRS advanced option IgnoreAffinityRulesForMaintenance to 1.

- a In the vSphere Client inventory, right-click the datastore cluster and select Edit Settings.
- b Select SDRS Automation and click Advanced Options.
- c Enter 1 in the Value column for IgnoreAffinityRulesForMaintenance.
- d Click OK.

Reference:

<https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.troubleshooting.doc%2FGUID-91D4FBBC-0677-44DA-B0F2-43A3504F8CA0.html>

Question 144

An administrator is unable to see performance statistics for only the Past Week performance data. The vCenter Server is using a Microsoft SQL Server Database.



What are three likely causes contributing to this issue? (Choose three.)

- ☒ Performance statistics are turned off.
- ☒ The Past Day rollup job is not present.
- ☒ The stats_rollup_1_proc is not present.
- ☐ The VMware Performance Charts Service is stopped.
- ☐ The JDBC URL is incorrect in the vcdb.properties file.

Explanation:

Explanation: Please refer to this link to troubleshoot this issue: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1004382

Question 145

Which statement is true for the Path Selection Plug-In VMW_PSP_MRU?

- ☐ VMW_PSP_MRU is default for a majority of active-active and active-passive arrays.
- ☐ VMW_PSP_MRU will remain on the selected path even if the state were to change.
- ☐ VMW_PSP_MRU is recommended for Virtual SAN.
- ☒ VMW_PSP_MRU will have no preferred path setting for the Plug-In.

Explanation:

Explanation:

VMW_PSP_MRU

The host selects the path that it used most recently. When the path becomes unavailable, the host selects an alternative path. The host does not revert back to the original path when that path

becomes available again. There is no preferred path setting with the MRU policy. MRU is the default policy for most active-passive storage devices.

Reference: <https://pubs.vmware.com/vsphere-55/index.jsp?topic=%2Fcom.vmware.vsphere.storage.doc%2FGUID-B7AD0CA0-CBE2-4DB4-A22C-AD323226A257.html>

Question 146

An administrator attempts to place a Storage DRS enabled Datastore into Maintenance Mode. The task never completes, and the Entering Maintenance Mode status remains at 1%.



Which two actions should the administrator take to resolve this problem? (Choose two.)

- ☒ Set the Storage DRS advanced option IgnoreAffinityRulesForMaintenance = 1.
- ☐ Set the Storage DRS advanced option IgnoreAffinityRulesForMaintenance = 0.
- ☒ Disable Storage DRS affinity rules associated with this datastore cluster.
- ☐ Enable Storage DRS affinity rules associated with this datastore cluster.

Explanation:

Explanation: Problem

A datastore in a datastore cluster cannot enter maintenance mode. The Entering Maintenance Mode status remains at 1%.

Cause

One or more disks on the datastore cannot be migrated with Storage vMotion. This condition can occur in the following instances.

- ? Storage DRS is disabled on the disk.
- ? Storage DRS rules prevent Storage DRS from making migration recommendations for the disk.

Solution

? If Storage DRS is disabled, enable it or determine why it is disabled. See Storage DRS is Disabled on a Virtual Disk for reasons why Storage DRS might be disabled.

? If Storage DRS rules are preventing Storage DRS from making migration recommendations, you can remove or disable particular rules.

1.
 - In the vSphere Client inventory, right-click the datastore cluster and select Edit Settings.
 - Select Rules and click the rule.
 - Click Remove.
 - Click OK.

? Alternatively, if Storage DRS rules are preventing Storage DRS from making migration recommendations, you can set the Storage DRS advanced option IgnoreAffinityRulesForMaintenance to 1.

- In the vSphere Client inventory, right-click the datastore cluster and select Edit Settings.
- Select SDRS Automation and click Advanced Options.
- Enter 1 in the Value column for IgnoreAffinityRulesForMaintenance.
- Click OK.

Reference: <https://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.troubleshooting.doc%2FGUID-91D4FBBC-0677-44DA-B0F2-43A3504F8CA0.html>

Question 147

An administrator is writing a kickstart script to upgrade an ESXi 6.x host.

◆◆

In which three locations can the script reside? (Choose three.)

- ☒ NFS
- ☒ USB
- ☒ HTTP
- ☐ TFTP
- ☐ PXE

Explanation:

Explanation:

Scripted ESXi installation

Running a script is an efficient way to deploy multiple ESXi hosts with an unattended installation.

The installation script contains the configuration data for the ESXi hosts that you want to install.◆ You can apply the script to all hosts that you want to have a similar configuration.

Notes:

The script must be created using supported commands.

Ensure you consider all disks that are connected to the host at the time of installation to avoid◆ inadvertently overwriting data.◆ Disk names can vary from machine to machine, one of the

settings that you may want to configure and double check in a script is the disk selection for the ESXi installation.

The host installation script must be stored in a location that the host can access by FTP, HTTP, HTTPS, NFS, USB flash drive, or CD-ROM drive. The installation media can be mounted on the host locally, remotely, or through PXE boot.

Reference: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2109708#Script