


3V0-22.21 Braindumps PDF, VMware 3V0-22.21 Exam Cram [Q10-Q28]



3V0-22.21 Braindumps PDF, VMware 3V0-22.21 Exam Cram
New 2022 3V0-22.21 Sample Questions Reliable 3V0-22.21 Test Engine

VMware 3V0-22.21 Exam Syllabus Topics:

Topic 1- Configure and manage advanced cluster configurations (Quickstart, Admission Control)- Configure the VMware Tools and vSphere® Client repository (productLocker)
Topic 2- Troubleshoot resource pool configuration issues (e.g. Limits- Sharing, hierarchy)- Troubleshoot vSphere® configuration issues
Topic 3- Troubleshoot resource issues of the VMware vCenter Server®- Manage and remediate clusters using vSphere® Lifecycle Manager?
Topic 4- Configure CPU scheduler for security options- Set up a cluster solution for VMware vSAN?
Topic 5- Configure and manage VMware vCenter Server® backup and restore options- Set up log collection and VMware vCenter Server® backup
Topic 6- Identify and implement vSphere® resource optimization opportunities- Performance-tuning, Optimization, and Upgrades
Topic 7- Configure and manage High Availability options- Set up a clustering solution for VMware vSphere® High Availability (HA)
Topic 8- Manage advanced virtual machine configurations (per VM EVC, latency sensitivity)- Architectures and Technologies- Products and Solutions- Planning and Designing
Topic 9- Create and manage host profiles (e.g. edit and apply host profiles)- Implement VMware vSphere® virtual machine encryption (VM encryption)
Topic 10- Deploy and manage VMware vSphere® Distributed Switch (VDS)- Deploy and manage VMware vSphere® Standard Switch (vSS)

NO.10 The IT Team recently implemented a Unified Communication Service to provide conferencing for the organization. The UCS Application server is running from a virtual machine named VM1-UCS. Due to the nature of the UCS. the traffic for VM1-UCS has a higher Class of Service (CoS) demand.

In order to meet demand bandwidth for VM1 hosting the UCS. you been have given the task to improve performance by isolating the traffic and configuring CoS for this VM.

Your first task is to connect VM1-UCS to the distributed port group. UCS-Portgroup.

Your second task is to tag and mark the traffic for VM1-UCS with the following parameters:

* Name: UCS Network Traffic

* CoS Value: 4

* DSCP Value: 34

* Traffic Direction: Ingress/Egress

* Traffic Qualifier: System Traffic – Virtual Machine

The IT Team noticed that some of the parameter were not available on the distributed switch. Troubleshoot why these settings are missing and fix it.

* Send us your suggestions.

NO.11 The security team has decided to follow the VMware-recommended best practices in the vSphere hardening guide.

esxi02b:

Your first task is to create a local user in esxi02b:

* Name: SpecialUser

* Role: Administrator

Your second task is to ensure that SpecialUser is the ONLY user who is able to SSH into esxi02b via Putty.

Your final task is to enforce a strict lockdown on esxi02b.

Your second task is to ensure that SpecialUser is the ONLY user who is able to SSH into esxi02b via Putty.

Your final task is to enforce a strict lockdown on esxi02b.

Authentication and authorization govern access. vCenter Single Sign-On supports authentication, which means it determines whether a user can access vSphere components at all. Each user must also be authorized to view or manipulate vSphere objects.

vSphere supports several different authorization mechanisms, discussed in Understanding Authorization in vSphere. The focus of the information in this section is how the vCenter Server permission model works and how to perform user management tasks.

vCenter Server allows fine-grained control over authorization with permissions and roles. When you assign a permission to an object in the vCenter Server object hierarchy, you specify which user or group has which privileges on that object. To specify the privileges, you use roles, which are sets of privileges.

Initially, only the administrator user for the vCenter Single Sign-On domain, administrator@vsphere.local by default, is authorized to log in to the vCenter Server system. That user can then proceed as follows:

Add an identity source in which users and groups are defined to vCenter Single Sign-On. See the Platform Services Controller Administration documentation.

Give privileges to a user or group by selecting an object such as a virtual machine or a vCenter Server system and assigning a role on that object for the user or group.

NO.12 Your storage administrator is concerned about a new application being deployed on virtual machine (SIOCVM) in your vSphere 7.x environment.

You've been asked to create and apply a storage policy to make sure that the SIOCVM virtual machine does not exceed 500 IOPS.

Note: Name the Storage Policy 500IOPSLimit
Storage I/O Control v2

Storage I/O Control (SIOC) was initially introduced in vSphere 4.1 to provide I/O prioritization of virtual machines running on a cluster of ESXi hosts that had access to shared storage. It extended the familiar constructs of shares and limits, which existed for CPU and memory, to address storage utilization through a dynamic allocation of I/O queue slots across a cluster of ESXi servers. The purpose of SIOC is to address the 'noisy neighbor' problem, i.e. a low priority virtual machine impacting other higher priority virtual machines due to the nature of the application and its I/O running in that low priority VM.

vSphere 5.0 extended SIOC to provide cluster-wide I/O shares and limits for NFS datastores. This means that no single virtual machine should be able to create a bottleneck in any environment regardless of the type of shared storage used. SIOC automatically throttles a virtual machine which is consuming a disparate amount of I/O bandwidth when the configured latency threshold has been exceeded. To allow other virtual machines receive their fair share of I/O bandwidth on the same datastore, a share based fairness mechanism has been created which now is supported on both NFS and VMFS.

vSphere 5.1 introduced a new SIOC feature called Stats Only Mode. When enabled, it doesn't enforce throttling but gathers statistics to assist Storage DRS. Storage DRS now has statistics in advance for new datastores being added to the datastore cluster & can get up to speed on the datastores profile/capabilities much quicker than before.

Another 5.1 feature was Automatic Threshold Computation. The default latency threshold for SIOC is 30ms. Not all storage devices are created equal so this default was chosen as a sort of 'catch-all'. There are certain devices which will hit their natural contention point much earlier than others, for example All Flash Arrays, in which case the threshold should be lowered by the user. However, manually determining the correct latency can be difficult for users. This gave rise to the need for the latency threshold to get automatically determined at a correct level for each device. Using the I/O injector modeling of SIOC, peak throughput and corresponding latency of a datastore is measured. The latency threshold value at which Storage I/O Control will kick in is then set to 90% of this peak value (by default). vSphere administrators can change this 90% to another percentage value or they can still input a millisecond value if they so wish.

The default latency threshold for SIOC can be reduced to as low as 5ms.

SIOC V1 Overview

SIOC V1 is disabled by default. It needs to be enabled on a per datastore level, and it is only utilized when a specific level of latency has been reached. By default, the latency threshold for a datastore is set to 30ms, as mentioned earlier. If SIOC is triggered, disk shares (aggregated from all VMDKs using the datastore) are used to assign I/O queue slots on a per host basis to that datastore. In other words, SIOC limits the number of IOs that a host can issue. The more VMs/VMDKs that run on a particular host, the higher the number of shares, and thus the higher the number of IOs that that particular host can issue. The throttling is done by modifying the device queue depth of the various hosts sharing the datastore. When the period of contention passes, and latency returns to normal values, the device queue depths are allowed to return to default values on each host.

SIOC V2 Introduction

Before describing SIOC V2, it should be highlighted that SIOC V1 and SIOC V2 can co-exist on vSphere 6.5. This makes it much

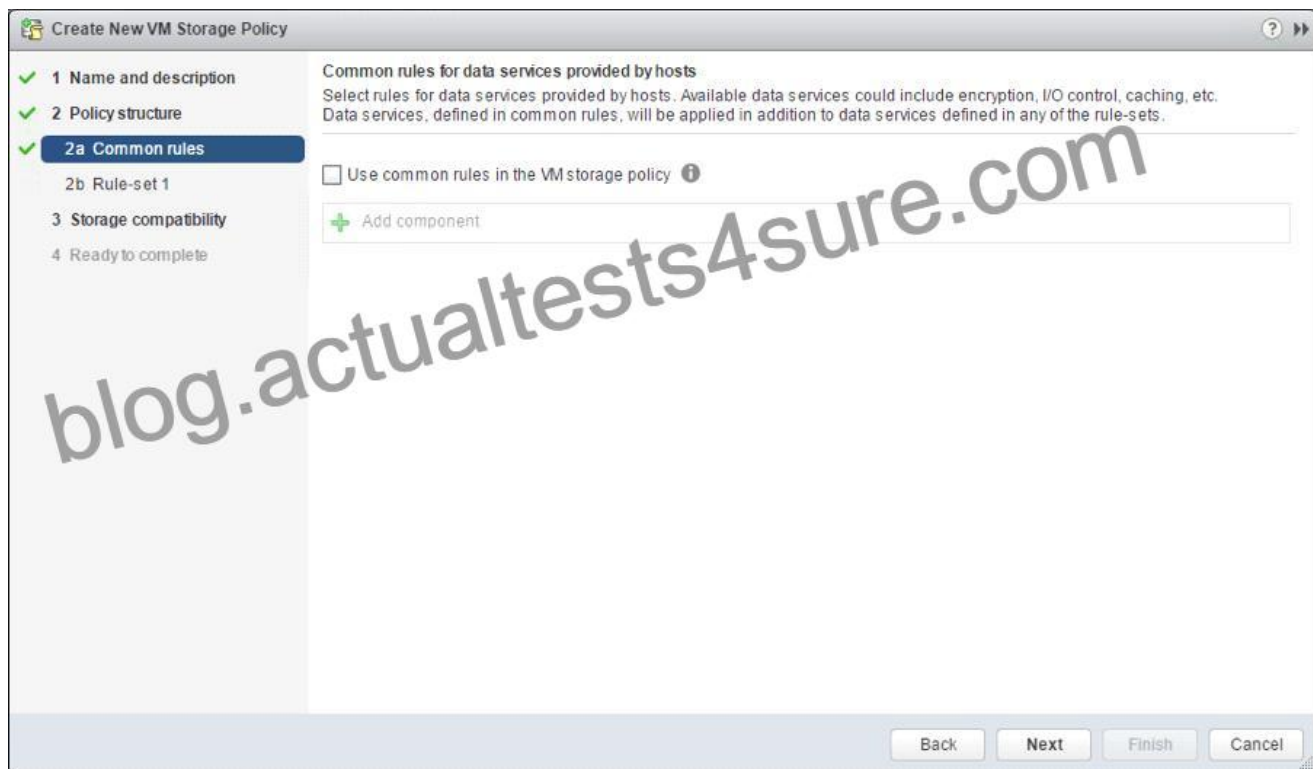
simpler when considering upgrades, or migrations between versions. With that in mind, SIOC V2 is considerably different from a user experience perspective when compared to V1. SIOCV2 is implemented using IO Filter framework Storage IO Control category. SIOC V2 can be managed using SPBM Policies. What this means is that you create a policy which contains your SIOC specifications, and these policies are then attached to virtual machines.

Creating an SIOC policy based

Creating an SIOC policy is done in exactly the same way as building a storage policy for VSAN or Virtual Volumes. Select the VM Storage Policy from the vSphere client home page, and from there select the option to create a new VM Storage Policy. VM Storage Policies in vSphere 6.5 has a new option called '“Common Rules”'. These are used for configuring data services provided by hosts, such as Storage I/O Control and Encryption.

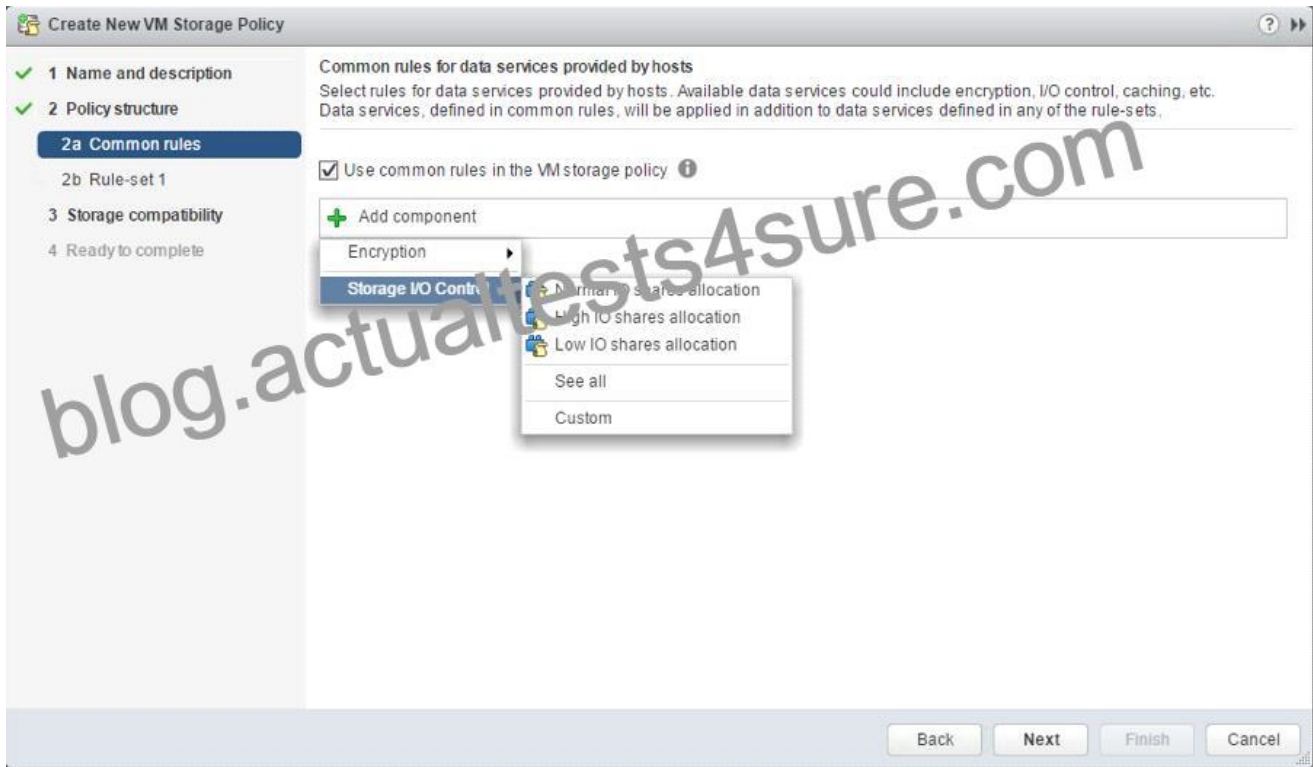
Use common rules in the VM storage policy

The first step is to click on the check box to enable common rules. This will then allow you to add components, such as SIOC, to the policy.



Add Component '– Storage I/O Control

In vSphere 6.5, there are two components available for common rules, Encryption and Storage I/O Control. Select Storage I/O Control in this case. Now you can select Normal, High, Low or Custom shares allocation.



This table describes the different Limits, Shares and Reservations associated with each setting:

HIGH

NORMAL

LOW

Limits

100,000

10,000

1,000

Reservation

100

50

10

Shares

2,000

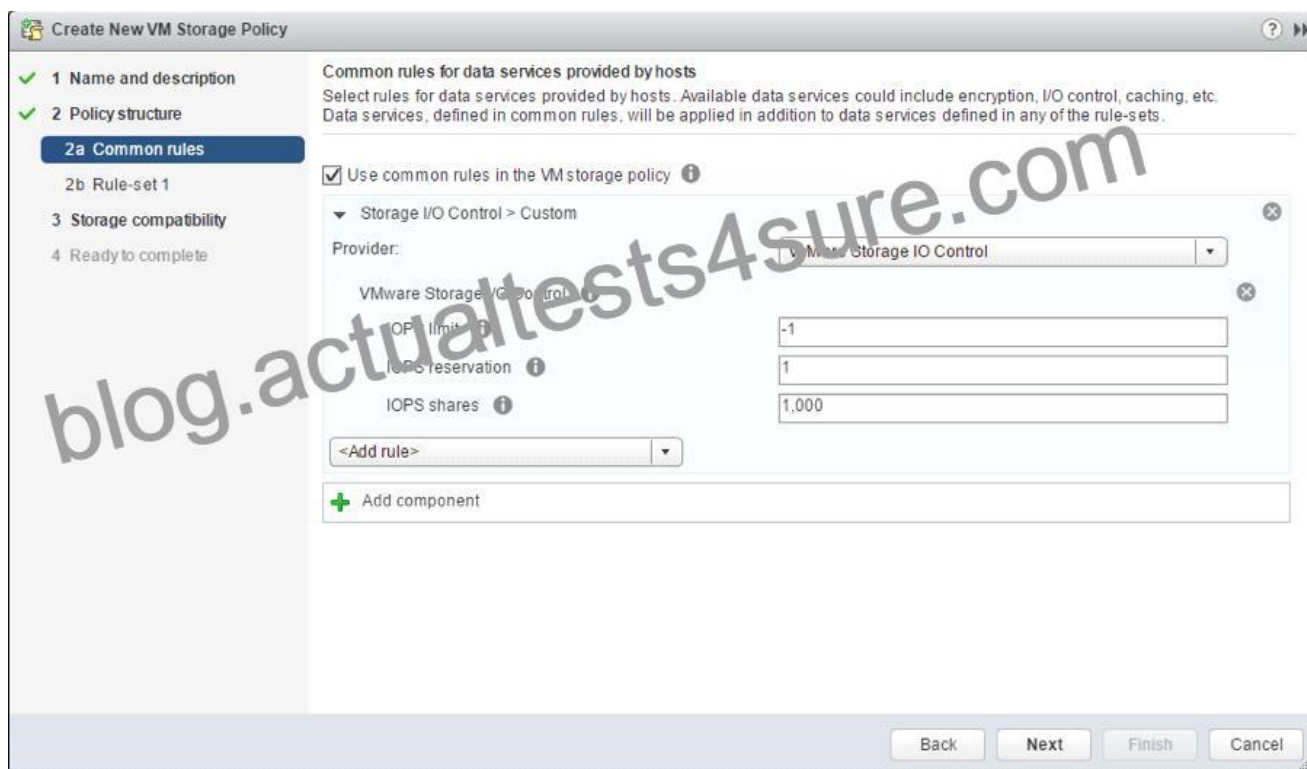
1,000

500

When the policy has been created, it may be assigned to newly deployed VMs during provisioning, or to already existing VMs by assigning this new policy to the whole VM (or just an individual VMDK) by editing its settings. One thing to note is that IO Filter based IOPS does not look at the size of the IO. For example, there is no normalization so that a 64K IOP is not equal to 2 x 32K IOPS. It is a fixed value of IOPS irrespective of the size of the IO.

Custom Allocation

If neither of the values in the Normal, High, Low allocations is appropriate, there is the ability to create custom settings for these values. In a custom setting, IOPS limit and IOPS reservation are both set to -1, implying unlimited. These may be modified as required.



Advanced Options

SchedCostUnit

This is an advanced parameter that was created for SIOC V1 only. SIOC V2 does not have SchedCostUnit implemented. For V1, SchedCostUnit determines the unit size (normalized size) of an IO operation for scheduling, and it is currently a constant value of 32K. This constant value, however, may not satisfy different requirements from different customers. Some customers may want to set this unit size to 4K. Other customers may want to set it up to 256K.

To satisfy these different requirements, SchedCostUnit is now configurable. It defaults to an IO size value of 32K, and allowable values range between 4K to 256K.

The SchedCostUnit dictates how requests are counted. A request with size \leq SchedCostUnit counts as a single I/O. Anything greater than SchedCostUnit will be counted as 2 or more requests.

For example, by changing the SchedCostUnit from 32K to 64K, the number of IOPS observed will halve. The size of the IO can be set using the:

```
&#8220;esxcli system settings advanced set -o /Disk/SchedCostUnit -i 65536&#8221;
```

and verified by using the”

```
&#8220;esxcli system settings advanced list -o /Disk/SchedCostUnit&#8221;
```

command. SIOC V2 counts guest IO directly. IOPS will be counted based on IO count, regardless of the IO size.

SchedReservationBurst

When limits are set on VMDKs, requests could have high average latency because the limit was enforced at a high (per request) granularity. This was due to the strict enforcement on a VM getting its share of IOs in interval of 1 second/L, where L is the user specified limit. The issue is more visible in fast storage, such as flash arrays. It was noted that SIOC V2 did not perform well when presented with a “bursty” workload on fast storage.

This SchedReservationBurst setting relaxes that constraint so a VM get its share of IOs at any time during a 1 second window, rather than enforce strict placement of IOs in intervals of 1/L. BURST option is turned-on by default.

SIOC V2 Limitations

In this initial release of SIOC V2 in vSphere 6.5, there is no support for vSAN or Virtual Volumes. SIOC v2 is only supported with VMs that run on VMFS and NFS datastores.

NO.13 Your team is experiencing intermittent issues with esxi0la and you have been asked to configure the host to export its syslog data to a preconfigured syslog collector.

To complete this task, you must:

- * Configure esxi0la.vciass.local to send syslog events to an external syslog collector on 172.20.10.10.
- * Ensure that the ESXi host security policies allow the syslog traffic to pass.
- * Send us your suggestions.

NO.14 The current vSphere environment will be adding new ESXi hosts that will be used to create a QA compute cluster. This cluster should have HA properties specific to the workloads that will be running in it.

In preparation of adding the new hosts, create the new cluster. QA-Cluster in. Datacenter-PROD on vcsa0la.vclass.local with the following HA requirements:

- * The cluster should not contain any ESXi hosts or VMs

- * Hosts should be monitored.
- * VMs should be restarted in the event of a host failure.
- * VMs should be restarted if guest heartbeats are not detected.
- * In the case of a host becoming isolated, shutdown and restart VMs.
- * If there is an All Paths Down event, any affected VMs must be moved to another host.
- * Reserve 10% of memory and CPU for failover capacity.

Part 2

You have been given a requirement for a virtual machine to have no downtime when an ESXi host failure occurs. Configure Fault Tolerance on VM1-FT in the PROD-B cluster. Use any compatible secondary host and datastore. Configure the following advanced cluster settings. Use SAN01 as the storage during configuration.

das.isolationaddress0 172.20.10.11

dass.igoreRedundantNetWaring true

Note: ignore any related host, customer, or bandwidth warnings as long as fault tolerance is configured and VM1-FT is running.

- * Send us your suggestions.

NO.15 The company's IT strategy is to adopt innovative and emerging technologies such as software-defined storage solution. The IT team has decided to run their business-critical workloads on an all-flash Virtual SAN (vSAN) as it provides excellent performance.

The IT team has purchased servers that are compatible with vSAN. However, all the solid-state drives (SSD) in the servers are shown incorrectly as hard-disk drives (HDD) instead.

In addition, some of the solid-state drives (SSD) will be used for other purposes instead of vSAN and should not be part of the vSAN cluster. These are the requirements for the vSAN cluster:

- * In each server, use the 3GB SSD as the cache tier and the 11GB SSD as the capacity tier
- * As a result the vSAN cluster will use a total of six SSDs (three SSDs for caching and three SSDs for capacity)
- * Ensure all the disks that will be used for vSAN are shown correctly as SSDs
- * Provide storage savings by using deduplication and compression.

Next, the IT team wants to improve the performance and availability of the business-critical workloads on the vSAN-datastore.

Ensure the following configurations will be applied on existing and new workloads located on vSAN-datastore:

Number of disk stripes per object: 2

Primary level of failures to tolerate: 2

Failure tolerance method: RAID-1 (Mirroring)

Force provisioning; Yes

The new configurations should be applied by default.

You may create new storage policy but do not edit the default vSAN storage policy as it may be used by other vSAN clusters in the future. Name the policy `“New vSAN Default’`.

Note: All tasks should be executed in PROD-A host cluster.

VMware vSphere ESXi can use locally attached SSDs (Solid State Disk) and flash devices in multiple ways. Since SSDs offer much higher throughput and much lower latency than traditional magnetic hard disks the benefits are clear. While offering lower throughput and higher latency, flash devices such as USB or SATADOM can also be appropriate for some use cases. The potential drawback to using SSDs and flash device storage is that the endurance can be significantly less than traditional magnetic disks and it can vary based on the workload type as well as factors such as the drive capacity, underlying flash technology, etc.

This KB outlines the minimum SSD and flash device recommendations based on different technologies and use case scenarios.

SSD and Flash Device Use Cases

A non-exhaustive survey of various usage models in vSphere environment are listed below.

Host swap cache

This usage model has been supported since vSphere 5.1 for SATA and SCSI connected SSDs. USB and low end SATA or SCSI flash devices are not supported.

The workload is heavily influenced by the degree of host memory over commitment.

Regular datastore

A (local) SSD is used instead of a hard disk drive.

This usage model has been supported since vSphere 7.0 for SATA and SCSI connected SSDs.

There is currently no support for USB connected SSDs or for low end flash devices regardless of connection type.

vSphere Flash Read Cache (aka Virtual Flash)

This usage model has been supported since vSphere 5.5 for SATA and SCSI connected SSDs.

There is no support for USB connected SSDs or for low end flash devices.

vSAN

This usage model has been supported since vSphere 5.5 for SATA and SCSI SSDs. For more information, see the vSAN Hardware Quick Reference Guide.

vSphere ESXi Boot Disk

A USB flash drive or SATADOM or local SSD can be chosen as the install image for ESXi, the vSphere hypervisor, which then boots from the flash device.

This usage model has been supported since vSphere 3.5 for USB flash devices and vSphere 4.0 for SCSI/SATA connected devices.

Installation to SATA and SCSI connected SSD, SATADOM and flash devices creates a full install image which includes a logging partition (see below) whereas installation to a USB device creates a boot disk image without a logging partition.

vSphere ESXi Coredump device

The default size for the coredump partition is 2.5 GiB which is about 2.7 GB and the installer creates a coredump partition on the boot device device for vSphere 5.5 and above. After installation the partition can be resized if necessary using partedUtil. For more information, see the vSphere documentation.

Any SATADOM or SATA/SCSI SSD may be configured with a coredump partition.

This usage model has been supported from vSphere 3.5 for boot USB flash devices and since vSphere 4.0 for any SATA or SCSI connected SSD that is local.

This usage model also applies to Autodeploy hosts which have no boot disk.

vSphere ESXi Logging device

A SATADOM or local SATA/SCSI SSD is chosen as the location for the vSphere logging partition (/scratch partition). This partition may be but need not be on the boot disk and this applies to Autodeploy hosts which lack a boot disk.

This usage model has been supported since vSphere 7.0 for any SATA or SCSI connected SSD that is local. SATADOMs that meet the requirement set forth in Table 1 are also supported.

This usage model can be supported in a future release of vSphere for USB flash devices that meet the requirement set forth in Table 1.

SSD Endurance Criteria

The flash industry often uses Tera Bytes Written (TBW) as a benchmark for SSD endurance. TBW is the number of terabytes that can be written to the device over its useful life. Most devices have distinct TBW ratings for sequential and random IO workloads, with the latter being much lower due to Write Amplification Factor (WAF) (defined below). Other measures of endurance commonly used are DWPD (Drive Writes Per Day) and P/E (Program/Erase) cycles.

Conversion formulas are provided here:

Converting DWPD (Drive Writes Per Day) to TBW (Terabytes Written):

$TBW = DWPD * Warranty \text{ (in Years)} * 365 * Capacity \text{ (in GB)} / 1,000 \text{ (GB per TB)}$ Converting Flash P/E Cycles per Cell to TBW (Terabytes Written):

$TBW = Capacity \text{ (in GB)} * (P/E \text{ Cycles per Cell}) / (1,000 \text{ (GB per TB)} * WAF)$ WAF is a measure of the induced writes caused by inherent properties of flash technology. Due to the difference between the storage block size (512 bytes), the flash cell size (typically 4KiB or 8KiB bytes) and the minimum flash erase size of many cells one write can force a number of induced writes due to copies,

garbage collection, etc. For sequential workloads typical WAFs fall in the range of single digits while for random workloads WAFs can approach or even exceed 100. Table 1 contains workload characterization for the various workloads excepting the Datastore and vSphere Flash Read Cache workloads which depend on the characteristics of the Virtual Machines workloads being run and thus cannot be characterized here. A WAF from the table can be used with the above P/E to TBW formula.

NO.16 A vSphere administrator has deployed a new server. The VM will have a workload which is prodApp1 to the following specifications:

- * The VM should never have any memory contention while powered on. even if the host that it resides
- * Configure the virtual machine for high latency sensitivity.
- * Send us your suggestions.

Feel VMware 3V0-22.21 Dumps PDF Will likely be The best Option:

<https://www.actualtests4sure.com/3V0-22.21-test-questions.html>