

DUMPS ARENA

Advanced Deploy VMware vSphere 7.x

VMware 3V0-22.21

Version Demo

Total Demo Questions: 5

Total Premium Questions: 17

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1 - (SIMULATION)

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

ANSWER: Seethesolutionbelow:**Explanation:**

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

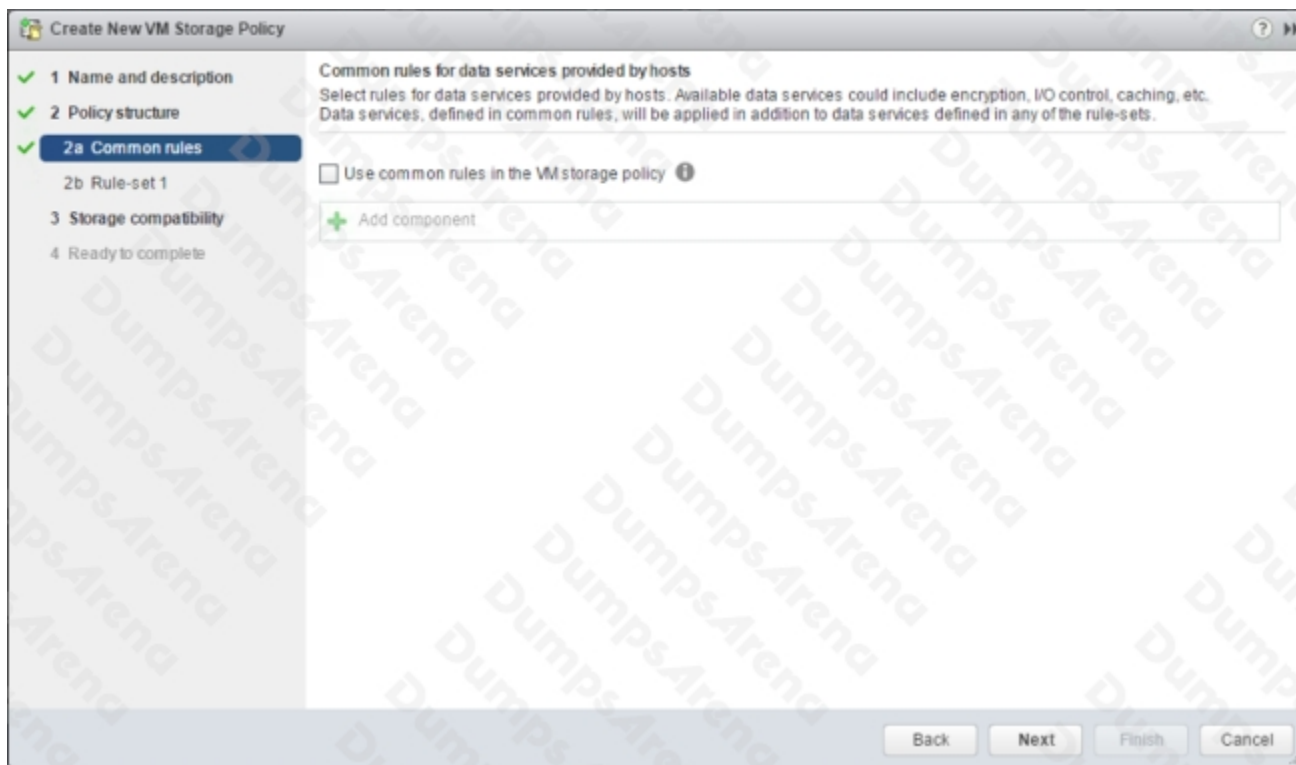
Before describing SIOC V2, it should be highlighted that SIOC V1 and SIOC V2 can co-exist on vSphere 6.5. This makes it much simpler when considering upgrades, or migrations between versions. With that in mind, SIOC V2 is considerably different from a user experience perspective when compared to V1. 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.

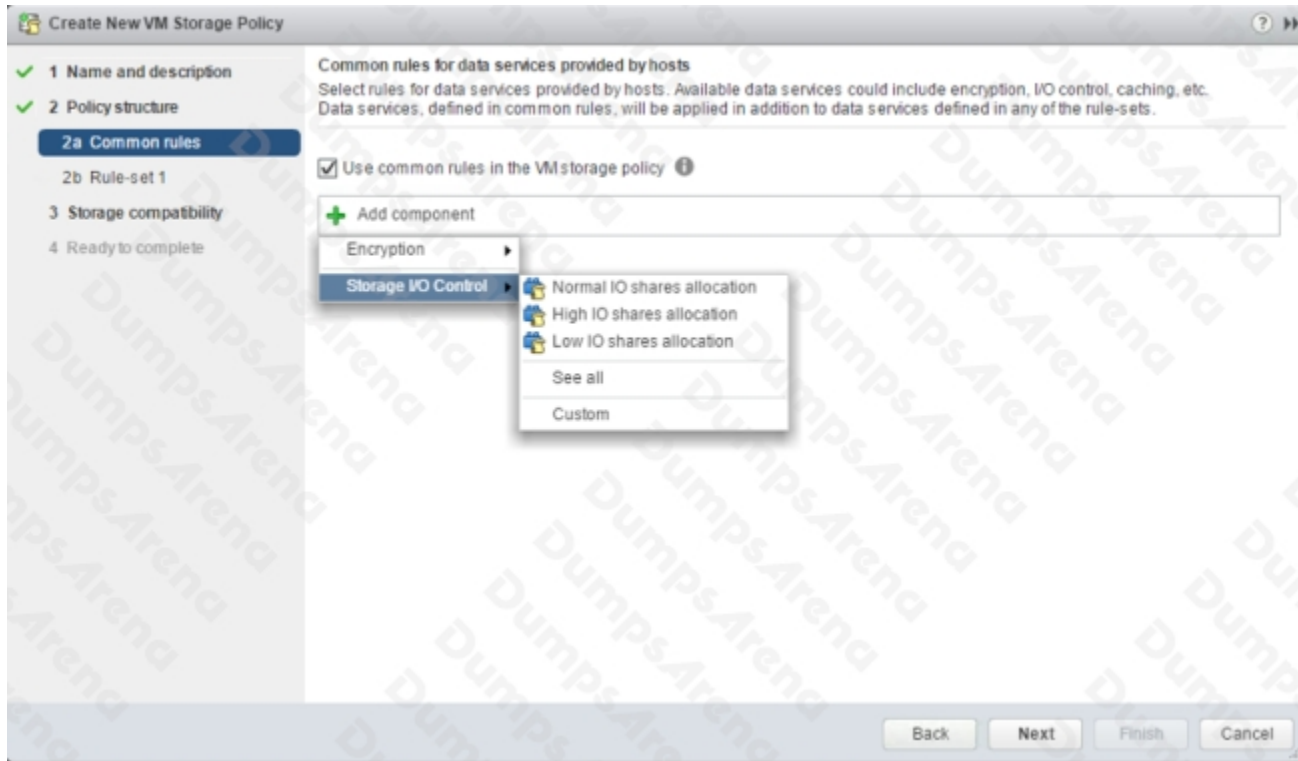
Use common rules in the VM storage policy

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

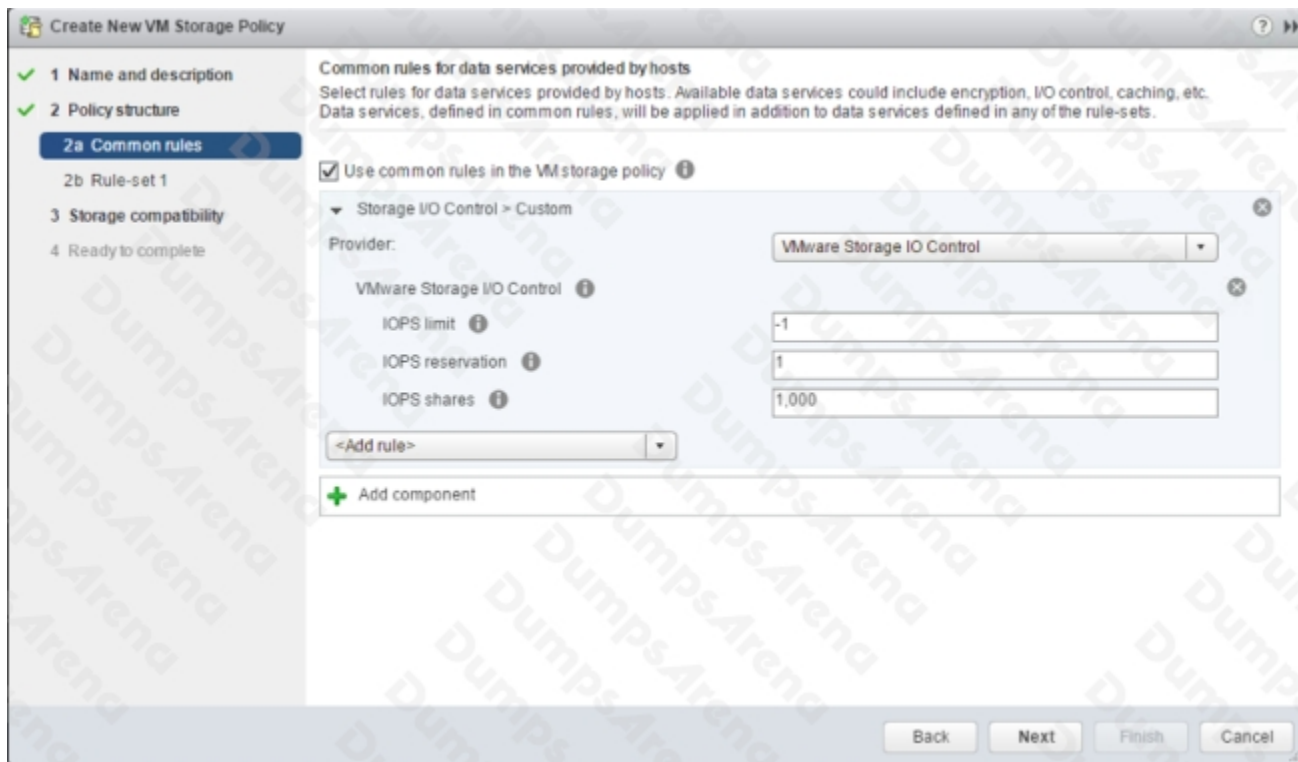
1,000

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:

```
"esxcli system settings advanced set -o /Disk/SchedCostUnit -i 65536"
```

and verified by using the"

"esxcli system settings advanced list -o /Disk/SchedCostUnit"

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.

QUESTION NO: 2 - (SIMULATION)

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:

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.

ANSWER: seethesolutionbelow:

Explanation:

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.

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:

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.

QUESTION NO: 3 - (SIMULATION)

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:

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".

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

ANSWER: seethesolutionbelow:

Explanation:

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.

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:

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.

QUESTION NO: 4 - (SIMULATION)

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 where not available on the distributed switch. Troubleshoot why these settings are missing and fix it.

ANSWER: Sendyoursuggestions.

QUESTION NO: 5 - (SIMULATION)

You have been asked to create a new datastore for the Production cluster; however, the ESXi host esxi02a is not able to add the datastore. Your storage team has determined that there is nothing incorrect with the storage presentation.

Troubleshoot why host esxi02a is not able to create the datastore. Add the datastore to the host once the proper configuration has been applied, and then reboot the host!

Use the following information to complete this task;

- ESXi host: esxi02a
- Datastore Name: new_datastore
- Datastore target: iSCSI Target
- Datastore LUN: 4
- Reboot the host

ANSWER: Seethesolutionbelow:

Explanation:

Cause

This issue may occur if:

Resolution

To prevent this issue: