Get 2024 Free Network Appliance NS0-593 Exam Practice Materials Collection [Q11-Q30

Get 2024 Free Network Appliance NS0-593 Exam Practice Materials Collection Get Latest and 100% Accurate NS0-593 Exam Questions Q11. After users start reporting the Inability to create new flies in a CIFS share, you find EMS events for

wafl.dir.size.max logged for the volume of the SVM to which the share points.

In this scenario, which action should you take to solve this issue?
* Delete unneeded files from the directory.
* Move the volume to a different aggregate.
* Increase the maximum numberof files for the volume.
* Increase the maximum directory size for the volume.

Q12. Your customer installed the shelf firmware for their NS224 shelf over a week ago, and the firmware has not upgraded on shelf 1 module B. The customer wants to know what the next steps would be to get the firmware upgraded after verifying that the shelf firmware is indeed loaded onto the system.

Which step would you perform to complete the firmware upgrade?
* Reseat the NSM100 module.

* Reseat the disk in Bay 0.
* Power cycle the shelf.
* Reseat the PSU of the shelf.

The question refers to a scenario where the shelf firmware for an NS224 shelf has not been upgraded on one of the NVMe shelf modules (NSM) after a week of installation.

The NSM is responsible for managing the communication between the drives and the I/O modules (IOM) in the shelf1.

The shelf firmware for the NSM is automatically updated when the NSM is inserted into the shelf or when the system is rebooted2.

If the automatic update does not work, the manual update process involves reseating the NSM, which means removing it from the shelf and inserting it back3.

Reseating the NSM triggers the firmware update and also resets the NSM’s state3.

The other options are not correct, because:

B) Reseating the disk in Bay 0 will not affect the NSM firmware update, as the disk is not connected to the NSM1.

C) Power cycling the shelf will disrupt the I/O operations and may cause data loss or corruption4.

D) Reseating the PSU of the shelf will not affect the NSM firmware update, as the PSU is not connected to the NSM1. Reference:

NS224 NVMe drive shelf overview – NetApp

Shelf firmware update process – NetApp

Module firmware upgrade stuck on NS224 shelf – NetApp Knowledge Base

Power cycle a disk shelf – NetApp

Q13. You receive the panic message shown in the exhibit.

```
Uncorrectable Machine Check Error at CPU0. MC0 Error: STATUS<0xb200000430000800>
(Val,UnCor,Enable,PCC,ErrCode(Src,NTO,Gen,Mem,L0)). MC5 Error: STATUS<0xf2000010c4300e0f>
(Val,OverF,UnCor,Enable,PCC,ErrCode(Gen,NTO,Gen,Gen,Gen)); Uncorrectable error at DIMM-1,
Channel 0, Serial: BA-00-1131-00098398!69002460-I01-NTA-T17!, FERR(0x400), NERR(0x402), MERR
M10Err, Rank 3, Bank 6, CAS 0x1e8, RAS 0x1bcf Uncorrectable error at DIMM-1, Channel 0, Serial:
BA-00-1131-00098398!69002460-I01-NTA-T17!, MERR M10Err, Rank 3, Bank 6, CAS 0x1e8, RAS 0x1bc.
```

In this scenario, which component should you troubleshoot first?

* the PCI card in slot 3
* the MetroCluster FC-Vl card in slot 6
* the memory module in slot 1
* the CPU

Q14. A user mentions that their home drive, that Is an export within a volume, is no longer allowing them to save

files. The drive reports that it Is full, even though It shows that minimal data is written to it.

Which statement would explain this behavior?
* The mount is stale and uses a cached version of the volume.
* Other users wrote to this user’s home drive.
* Other files within the volume are also owned by the user, exceeding the user quota.
* The client system needs to remount the export to show the proper space.

Q15. Your customer mentions that they have accidentally destroyed both root aggregates in their two-node cluster.

In this scenario, what are two actions that must be performed? (Choose two.)
* Rejoin the second node to the re-created cluster.
* Re-create the cluster from the local backup.
* Install ONTAP from a USB device.
* Re-create the cluster from the remote backup.

If both root aggregates are destroyed in a two-node cluster, the cluster will be inoperable and the data will be inaccessible. To recover from this situation, you need to perform the following actions:

Install ONTAP from a USB device on one of the nodes. This will create a new root aggregate and a new cluster on that node.

Rejoin the second node to the re-created cluster. This will also create a new root aggregate on the second node and synchronize it with the first node.

Restore the cluster configuration and data from a backup, if available. Reference = ONTAP 9 Documentation Center Storage System Recovery Troubleshooting Recovering from a root aggregate failure

Q16. After a motherboard replacement on a NetApp AFF A300 in a SAN environment, the customer states that

ports 0e and 0f are unable toconnect to the fabric. The ports report “offline”.

What would you examine first to troubleshoot the issue?
* vserver fcp wwpn-alias show command output
* system node hardware unified-connect show command output
* storage port show command output
* vserver fcp interface show command output

Q17. An administrator receives the following error message:

```
Mon Dec 23 00:20:36 EST [nodeA: wafl_exempt08: wafl.cp.toolong:error]: Aggregate
fas_01_DATA_AGGR experienced a long CP.
```

What are two causes for this error? (Choose two.)
* There Is excessive SSD load causing the wear leveling to become unbalanced.
* A disk is falling.
* There is excessive SATA HDD load.
* An SSD disk is performing garbage collection to create a dense data layout.

The error message “wafl.cp.toolong:error” indicates that a WAFL consistency point (CP) took longer than 30 seconds to complete. A CP is a process that flushes the data from the NVRAM buffer to the disk. A long CP can cause latency and performance issues for the system1 One possible cause for a long CP is excessive SSD load causing the wear leveling to become unbalanced. Wear leveling is a technique that distributes the write operations evenly across the SSD cells to extend the lifespan of the SSD. If some SSD cells are written more frequently than others, the wear leveling will become unbalanced and the SSD performance will degrade2 Another possible cause for a long CP is an SSD disk performing garbage collection to create a dense data layout. Garbage collection is a process that reclaims the space occupied by invalid or deleted data on the SSD. Garbage collection

can improve the write performance and storage efficiency of the SSD, but it can also consume CPU and disk resources and cause long CPs3 A disk failing or being failed is not a likely cause for a long CP, because the system will automatically mark the disk as failed and remove it from the aggregate. The system will also initiate a disk reconstruction or a RAID scrub to restore the data protection and redundancy4 There is no evidence that the system has SATA HDDs, so there is no reason to assume that there is excessive SATA HDD load. Moreover, SATA HDDs are usually used for secondary or backup storage, not for primary or performance-sensitive workloads5 Reference:

1: Are long Consistency Points (wafl.cp.toolong) normal? – NetApp Knowledge Base 2: How to troubleshoot SSD performance issues – NetApp Knowledge Base 3: How to troubleshoot SSD garbage collection issues – NetApp Knowledge Base 4: How to troubleshoot disk failures and replacements – NetApp Knowledge Base 5: ONTAP 9 – Hardware Universe – The Open Group

Q18. When you review performance data for a NetApp ONTAP cluster node, there are back-to-back (B2B) type

consistency points (CPs) found occurring onthe loot aggregate.

In this scenario, how will performance of the client operations on the data aggregates be affected?
* During B2B processing, clients will be unable to write data.
* Data aggregates will not be affected by B2B processing on another aggregate.
* During B2B processing, all I/O to the node is stopped.
* During B2B processing, clients will be unable to read data.

Q19. You have a customer complaining of long build times from their NetApp ONTAP-based datastores. They

provided you packet traces from the controller and client. Analysis of these traces shows an average service

response time of 1 ms. QoS outputconfirms the same. The client traces are reporting an average of 15 ms in

the same time period.

In this situation, what would be your next step?
* The cluster is responding slowly and requires further investigation using performance archives.
* The client that reports high latency should be investigated.
* The cluster interconnects should be investigated.
* A sync core should be triggered.

Q20. You have a customer complaining of long build times from their NetApp ONTAP-based datastores. They provided you packet traces from the controller and client. Analysis of these traces shows an average service response time of 1 ms. QoS output confirms the same. The client traces are reporting an average of 15 ms in the same time period.

In this situation, what would be your next step?
* The cluster is responding slowly and requires further investigation using performance archives.
* The client that reports high latency should be investigated.
* The cluster interconnects should be investigated.
* A sync core should be triggered.

The question describes a scenario where the controller and client have a significant difference in their reported latency for the same datastores.

The controller’s latency is 1 ms, which is within the normal range for ONTAP-based datastores1.

The client’s latency is 15 ms, which is much higher than the controller’s latency and could indicate a performance issue on the client side2.

Therefore, the next step is to investigate the client that reports high latency and identify the possible causes, such as network congestion, misconfiguration, resource contention, or application issues23.

The other options are not relevant or appropriate for this scenario, because:

A) The cluster is not responding slowly, as the controller’s latency is low and QoS output confirms the same.

C) The cluster interconnects are not likely to be the cause of the latency difference, as they are used for communication between nodes within the cluster, not between the controller and the client4.

D) A sync core is a diagnostic tool that captures the state of the system at a given point in time, and is not a troubleshooting step for performance issues5. Reference:

ONTAP 9 Performance – Resolution Guide – NetApp Knowledge Base

Performance troubleshooting – NetApp

How to troubleshoot performance issues in Data ONTAP 8 7-mode

Cluster interconnect network – NetApp

How to generate a sync core on a node – NetApp

Q21. Which two automation methods does NetApp ONTAP Select support? (Choose two.)
* REST
* Anslble
* Docker
* PHP

Q22. You are attempting to connect a NetApp ONTAP cluster to a very complex network that requires LIFs to fail over across subnets.

How would you accomplish this task?
* Configure an equal number of UFs on each subnet.
* Configure VIP LIFs using OSPF.
* Configure VIP LIFs using BGP.
* Configure a I IF failover policy for each subnet inside a single broadcast domain.

A LIF (Logical Interface) is a logical entity that represents a network connection point on a node1.

A VIP LIF (Virtual IP LIF) is a LIF that can fail over across subnets within an IPspace2.

BGP (Border Gateway Protocol) is a routing protocol that enables VIP LIFs to advertise their IP addresses to external routers and to update the routing tables when a failover occurs3.

To connect a NetApp ONTAP cluster to a complex network that requires LIFs to fail over across subnets, you need to configure VIP LIFs using BGP on the cluster and on the external routers3.

This way, you can ensure that the network traffic is routed to the optimal node and port for each VIP LIF, and that the network connectivity is maintained in the event of a node or port failure3. Reference:

1: Logical Interfaces, ONTAP 9 Documentation Center

2: VIP LIFs, ONTAP 9 Documentation Center

3: Configuring BGP on a cluster, ONTAP 9 Documentation Center

Q23. A storage administrator reports that a monitoring toot is reporting that the storage controller reads between 90% to 93% CPU use. You run the sysstat -m command against the node in question.

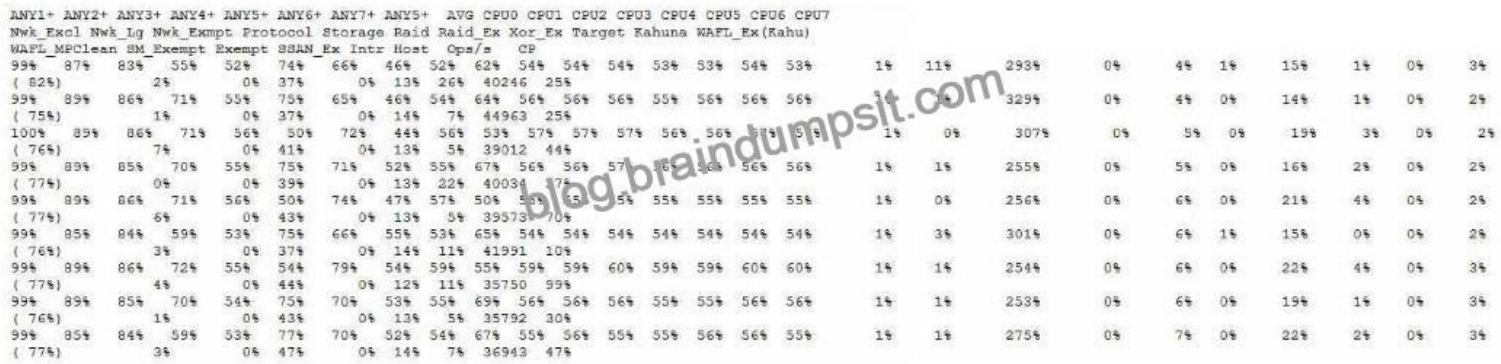

```
ANY1+ ANY2+ ANY3+ ANY4+ ANY5+ ANY6+ ANY7+ ANY5+  AVG CPU0 CPU1 CPU2 CPU3 CPU4 CPU5 CPU6 CPU7
Nwk_Excl Nwk_Ig Nwk_Exmpt Protocol Storage Raid Raid_Ex Xor_Ex Target Kahuna WAFL_Ex(Kahu)
WAFL_MPClean SM_Exempt Exempt SSAN_Ex Intr Host  Ops/s   CP
99%   87%   83%   55%   52%   74%   66%   46%  52%  62%  54%  54%  54%  53%  53%  54%  53%      1%   11%      293%      0%      4%   1%     15%     1%    0%    3%
( 82%)          2%          0%  37%       0%  13%  26%  40246  25%
99%   89%   86%   71%   55%   75%   65%   46%  54%  64%  56%  56%  56%  55%  56%  56%  56%      1%   0%       329%      0%      4%   0%     14%     1%    0%    2%
( 75%)          1%          0%  37%       0%  14%   7%  44963  25%
100%   89%   86%   71%   56%   50%   72%   44%  56%  53%  57%  57%  57%  56%  56%  57%  56%      1%   0%       307%      0%      5%   0%     19%     3%    0%    2%
( 76%)          7%          0%  41%       0%  13%   5%  39012  44%
99%   89%   85%   70%   55%   75%   71%   52%  55%  67%  56%  56%  57%  56%  56%  56%  56%      1%   1%       255%      0%      5%   0%     16%     2%    0%    2%
( 77%)          0%          0%  39%       0%  13%  22%  40036  17%
99%   89%   86%   71%   56%   50%   74%   47%  57%  50%  55%  55%  55%  55%  55%  55%  55%      1%   0%       256%      0%      6%   0%     21%     4%    0%    2%
( 77%)          6%          0%  43%       0%  13%   5%  39573  70%
99%   85%   84%   59%   53%   75%   66%   55%  53%  65%  54%  54%  54%  54%  54%  54%  54%      1%   3%       301%      0%      6%   1%     15%     0%    0%    2%
( 76%)          3%          0%  37%       0%  14%  11%  41991  10%
99%   89%   86%   72%   55%   54%   79%   54%  59%  55%  59%  59%  60%  59%  59%  60%  60%      1%   1%       254%      0%      6%   0%     22%     4%    0%    3%
( 77%)          4%          0%  44%       0%  12%  11%  35750  99%
99%   89%   85%   70%   54%   75%   70%   53%  55%  69%  56%  56%  56%  55%  55%  56%  56%      1%   1%       253%      0%      6%   0%     19%     1%    0%    3%
( 76%)          1%          0%  43%       0%  13%   5%  35792  30%
99%   85%   84%   59%   53%   77%   70%   52%  54%  67%  55%  56%  55%  55%  56%  56%  55%      1%   1%       275%      0%      7%   0%     22%     2%    0%    3%
( 77%)          3%          0%  47%       0%  14%   7%  36943  47%
```


Referring to the exhibit, which statement is correct?
* The customer should be advised to exclude certain workflows to reduce use.
* High network exempt use could be a problem.
* You should immediately investigate further by gathering perfstat data and opening a support case.
* The CPU Is not a first-order monitoring metric for ONTAP.

= CPU utilization in ONTAP is not a linear measure of the system load, nor can it be used alone as a measure of the overall system utilization. ONTAP uses a Coarse Symmetric Multiprocessing (CSMP) design which partitions system functions into logical processing domains, each with its own scheduling rules and resource availability. Therefore, a high CPU utilization does not necessarily indicate a performance problem, unless it is accompanied by other contributing factors such as high latency, low throughput, or high queue depth. ONTAP has several mechanisms to optimize CPU usage and balance the workload across the cores, such as WAFL parallelization, exempt processing, and CPU pinning. The CPU utilization reported by the sysstat command is an average across all cores and domains, and does not reflect the actual CPU activity or availability for each domain. Therefore, the CPU is not a first-order monitoring metric for ONTAP, and other metrics such as latency, throughput, and queue depth should be considered first. Reference = What is CPU utilization in Data ONTAP: Scheduling and Monitoring?, How to measure CPU utilization, What are CPU as a compute resource and the CPU domains in ONTAP 9?, Monitoring CPU utilization before ONTAP upgrade

Q24. Your customer has deployed a two petabyte NetApp ONTAP FlexGroup volume across their 4-node ONIAP 9.8 cluster. They plan to store over three billion files. They want to prevent file 10 conflicts since files are placed into the FlexGroup volume.

In this scenario, which two NFS SVM parameters should be enabled? (Choose two.)
* -v3-64bit-identifiers
* -v4-fsid-change
* -v3-fsid-change
* -v4-64bit-idantifiers

= To prevent file ID conflicts in a FlexGroup volume, you need to enable 64-bit NFSv3 and NFSv4 identifiers on the SVM that hosts the FlexGroup volume. This allows the SVM to use 64-bit file system IDs (FSIDs) and file IDs, which are unique across the cluster and can accommodate a large number of files. The -v3-64bit-identifiers and -v4-64bit-identifiers parameters enable this feature for NFSv3 and NFSv4 protocols respectively. Reference = Editing FlexGroup volumes, Enabling 64-bit NFSv3 identifiers on an SVM, NetApp ONTAP FlexGroup volumes – Best practices and implementation guide

Q25. After a normal power down of both nodes for building maintenance, Node01 of a 2-node cluster cannot be powered back up; however, all disk shelves are powered.

Which action should be performed to bring the cluster online and allow Node02 to serve data?
* Recreate the cluster with the system configuration recovery cluster recreate -from node command.
* Reboot the node With the system node reboot -node Node02 -bypass-optimization true command.
* Perform a takeover with the storage failover takeover -ofnode Node01 -option force command.
* Reinitialize the cluster with option 4a from the boot menu.

= The correct action to bring the cluster online and allow Node02 to serve data is to perform a takeover with the storage failover takeover -ofnode Node01 -option force command. This command will force Node02 to take over the resources of Node01 and serve the data from both nodes. This is necessary because Node01 is not responding and cannot initiate a graceful takeover. The other options are not correct because they will either destroy the existing cluster configuration (A and D) or reboot the node without taking over the resources of the other node (B). Reference = 1 Halt or reboot a node without initiating takeover in a two-node cluster – NetApp Documentation 2 Solved: Graceful shut down – NetApp Community

Q26. You notice poor performance on your FlexGroup and execute the system node run -node * flexgroup show

command for more Information. You notice the “Urge” column has non-zero values.

In this scenario, which statement is true?
* The aggregate is completely full.
* The constituent volumes are out ofInodes.
* The data placement is uneven.
* The constituent volumes are completely full.

Q27. You are trying to deploy a Connector in the AWScloud from NetApp Cloud Manager. The deployment fails

and shows the message ‘Insufficient permissions to deploy Cloud Connector”. You have verified the AWS

access key and the AWS secret key.

In this scenario, what Is the reason that the deployment failed?
* No AWS Marketplace subscription is associated with Cloud Manager.
* The required Identity and Access Management (IAM) policies were not Installed.
* The user lacks the permission to deploy within Cloud Manager.
* The Connector can be deployed only inAWS GovCloud (US).

Q28. After users start reporting the Inability to create new flies in a CIFS share, you find EMS events for wafl.dir.size.max logged for the volume of the SVM to which the share points.

In this scenario, which action should you take to solve this issue?
* Delete unneeded files from the directory.
* Move the volume to a different aggregate.
* Increase the maximum number of files for the volume.

* Increase the maximum directory size for the volume.

The wafl.dir.size.max event occurs when a directory has reached its maximum directory size (maxdirsize) limit, which prevents new file creation in that directory. The maxdirsize is a volume-level option that can be modified using the volume modify command. To solve this issue, you should increase the maximum directory size for the volume that contains the CIFS share, as long as it does not exceed 3% of the physical memory. Alternatively, you can reorganize the directory structure to avoid having too many files in one directory. Reference = Max Directory size error – NetApp Knowledge Base How to identify the target directory of wafl.dir.size.warning – NetApp Knowledge Base Post about wafl.dir.size.max | Source.kohlerville.com WAFL Max dirsize | uadmin

Q29. A user reports that a colleague saved a file called Test.txt from a UNIX system to a multiprotocol volume.

When opening the file later from a Windows system, it was not the file that they wanted. The file that they

wanted was named TEST~1.TXT.

Which statement explains this behavior?

* UNIX name mapping updated the filename.
* A Snapshot copy preserved two versions of the file.
* Windows Volume Shadow Copy Service stored an older version of the file.
* Case Insensitivity of SMB clients caused the file to be displayed with a different name.

Q30. When an administrator tries to create a share for an existing volume named voll, the process fails with an error.

```
cluster1::> vserver cifs share create -vserver svm1 -share-name voll -
path /vol1                                                                          Error:
command failed: The specified path "/vol1" does not exist in the namespace belonging to Vserver
"svm1".

cluster1::> vserver cifs share show
Vserver        Share         Path              Properties Comment  ACL
-------------- ------------- ----------------- ---------- -------- -----------
svm1           admin$        /                 browsable  -        -
svm1           c$            /                 oplocks    -        BUILTIN\Administrators /
Full Control
                                               browsable
                                               changenotify
                                               show-previous-versions
svm1           ipc$          /                 browsable  -        -

3 entries were displayed.

cluster1::> vserver cifs show
            Server          Status    Domain/Workgroup Authentication
Vserver     Name            Admin     Name             Style
----------- --------------- --------- ---------------- --------------
svm1        SVM1            up        DEMO             domain

cluster1::*> volume show -vservser svm1 -volume vol1

                                   Vserver Name: svm1
                                    Volume Name: vol1
                                 Aggregate Name: cluster1_01_SSD_1
    List of Aggregates for FlexGroup Constituents: cluster1_01_SSD_1
                                Encryption Type: none
               List of Nodes Hosting the Volume: cluster1-01
                                    Volume Size: 20MB
                                   Name Ordinal: base
                            Volume Data Set ID: 1028
                     Volume Master Data Set ID: 2162375168
                                   Volume State: online
                                   Volume Style: flex
                          Extended Volume Style: flexvol
                        FlexCache Endpoint Type: none
                         Is Cluster-Mode Volume: true
                         Is Constituent Volume: false
                                  Export Policy: default
                                        User ID: -
                                       Group ID: -
                                 Security Style: -
                               UNIX Permissions: ------------
                                  Junction Path: -
                           Junction Path Source: -
                                Junction Active: -
                         Junction Parent Volume: -
                                Junction Active: true
                         Junction Parent Volume: svm1_root
                            Vserver Root Volume: false
                                        Comment:
                                 Available Size: 18.76MB
                                Filesystem Size: 20MB
                        Total User-Visible Size: 19MB
                                      Used Size: 244KB
                                Used Percentage: 1%
                                            ...
                          Volume Tiering Policy: none
           Volume Tiering Minimum Cooling Days: -
            Performance Tier Inactive User Data: -
    Performance Tier Inactive User Data Percent: -
```

Referring to the exhibit, what Is the reason for the error?
* The volume must have a type of DP.
* The volume has not been mounted.
* The CIFS service is not authenticating properly with the domain controller.
* The CIFS service is not in workgroup mode.

The error message indicates that the specified path “/vol1” does not exist in the namespace belonging to Vserver “svm1”. This means that the volume “vol1” has not been mounted to the Vserver’s namespace, which is required for creating a share. The volume type, the CIFS service status, and the CIFS service mode are not relevant to the error. Reference = https://www.netapp.com/support-and-training/netapp-learning-services/certifications/support-engineer/

https://mysupport.netapp.com/site/docs-and-kb

The NS0-593 certification exam is a challenging exam that requires candidates to demonstrate their knowledge and skills in a variety of areas. To pass the exam, candidates must score a minimum of 70% on the exam. However, NetApp recommends that candidates aim for a higher score to demonstrate their mastery of the subject matter.

The NS0-593 exam is a specialist level certification and requires candidates to have a solid understanding of NetApp ONTAP storage systems. NS0-593 exam covers a wide range of topics including installation and configuration, data management, high availability, and troubleshooting. NS0-593 exam is designed to test the candidate's ability to apply their knowledge to real-world scenarios and solve complex problems.

Maximum Grades By Making ready With NS0-593 Dumps: https://www.braindumpsit.com/NS0-593_real-exam.html]