

# Troubleshoot Performance Issues in Hyperflex Clusters

## Contents

[Introduction](#)

[Identify](#)

[Cluster Types](#)

[Performance Chart Explanation](#)

[IOPS](#)

[Throughput](#)

[Latency](#)

[Frequency](#)

[External Factors](#)

[Identify Performance Issues at the Guest VM Level](#)

[Windows](#)

[ESXi](#)

[PVSCSI Check](#)

[Network Adapter Check](#)

[RSS Check](#)

## Introduction

This document describes the performance impact in a hyperflex environment, from the perspective of a Guest Virtual Machine (VM), ESXi host, and (SCVM)

## Identify

In order to troubleshoot the performance in a Hyperflex environment it is important to identify the type of cluster, the operation where the performance is degraded, the frequency of the performance degradation, and the level of performance impact that causes performance degradation.

There are multiple levels of impact in a hyperflex cluster, at the guest VM, the ESXI host level, and the Storage controller VM level.

## Cluster Types

â—◆ Hybrid nodes: Uses Solid State Drives (SSD) drives for caching and HDDs for the capacity layer.

â—◆ All-flash nodes: Uses SSD drives or Non-Volatile Memory Express (NVMe) storage for caching, and SSD drives for the capacity layer.

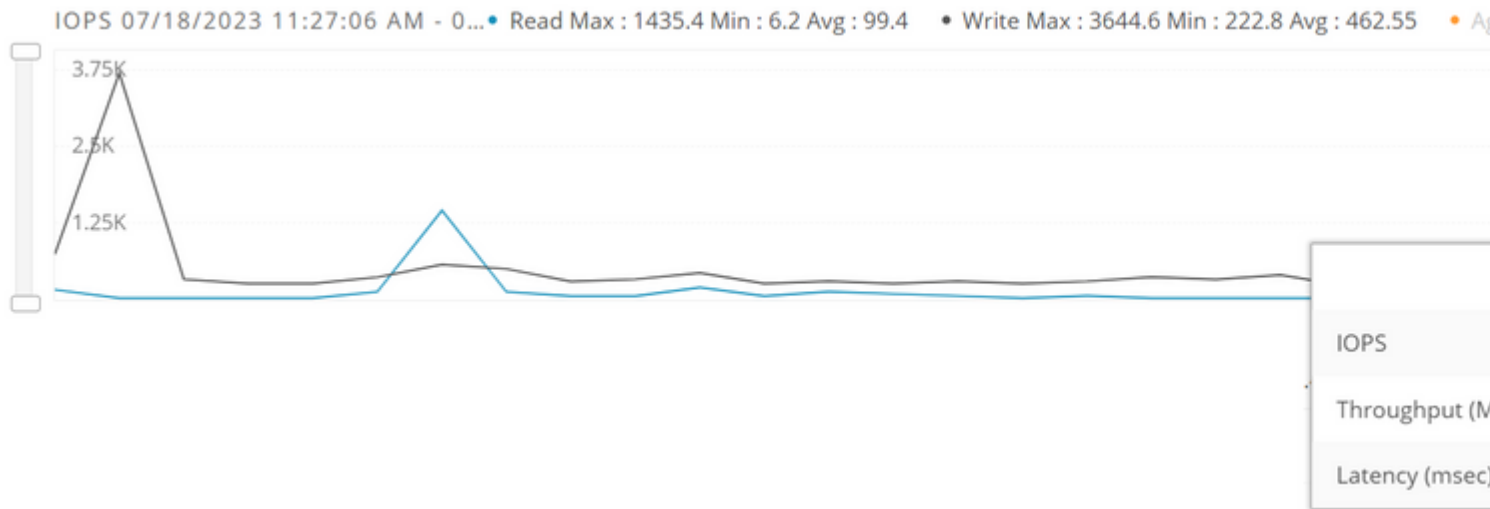
â—◆ All-NVMe nodes: Uses NVMe storage for both caching and the capacity layer all-NVMe nodes deliver the highest performance for the most demanding workloads with caching

## Performance Chart Explanation

The hyperflex systems have a feature to monitor performance, the charts display the read and write performance of the storage cluster.

## **IOPS**

Input/output operations per second (IOPS) is a common performance metric used to measure computer storage devices, including HDDs. This metric is used to evaluate performance for random I/O workloads.



IOPS performance chart.

## Throughput

The image shows the rate of data transfer in the storage cluster measured in Mbps.

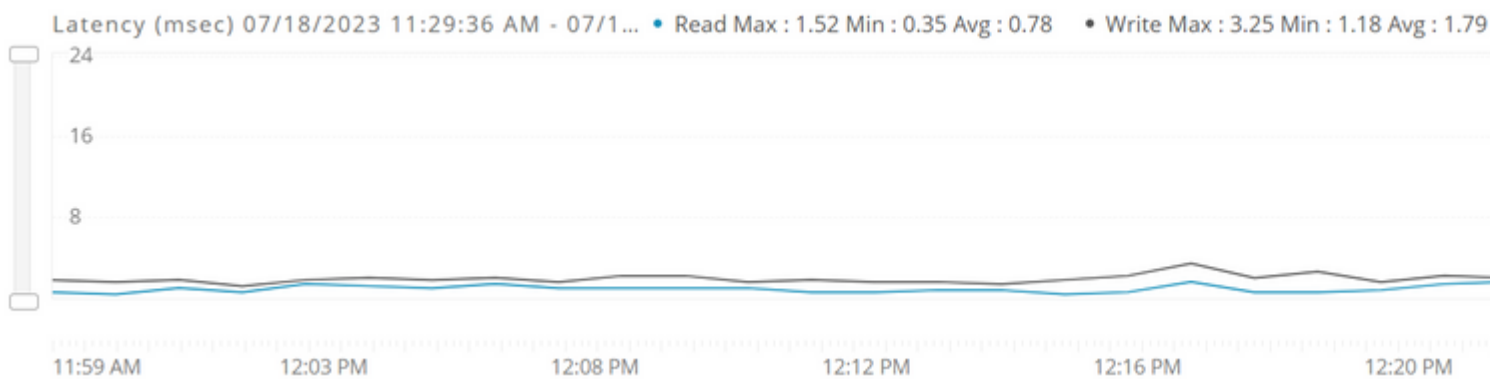


Throughput performance chart.

â€f

## Latency

Latency is a measure of how long it takes for a single I/O request to complete. It is the duration between issuing a request and receiving a response, is measured in milliseconds.



Latency performance chart.

â€f

## Frequency

It is important to define the frequency and duration of the performance impact to review the possible impact

[Performance Monitor](#)

[Resource Monitor](#)

## ESXi

After identifying the performance impact and reviewing the possible causes of the performance degradation, there are some performance checks to improve the performance.

- Check for overprovision (the total number of vCPUs that are assigned to all the VMs must not exceed the total number of physical cores available on the ESXi host machine).
- Even if the guest OS does not use some of its vCPUs, configuring VMs with those vCPUs still imposes some small resource requirements on ESXi that translate to real CPU consumption on the host.
- Over-allocating memory also unnecessarily increases the VM memory overhead and can lead to memory contention, especially if reservations are used.
- Verify that the Balloon driver is not retaining a hold on the memory, for more information, refer to this [link](#).

Review to [Troubleshooting ESX/ESXi virtual machine performance issues](#).

## PVSCSI Check

Paravirtual SCSI (PVSCSI) adapters are high-performance storage adapters that can result in greater throughput and lower CPU utilization for virtual machines with high disk IO requirements, it is recommended to use PVSCSI adapters. PVSCSI controller is a virtualization-aware, high-performance SCSI adapter that allows the lowest possible latency and highest throughput with the lowest CPU overhead.

> CPU	4	▼	ⓘ
> Memory	4	▼	GB ▼
> Hard disk 1	24	GB ▼	
▼ SCSI controller 0 *	LSI Logic SAS		
Change Type	LSI Logic SAS	▼	
SCSI Bus Sharing	None	▼	

*PVSCSI adapter.*

## Network Adapter Check

VMXNET 3 is a paravirtualized NIC designed for performance and provides high-performance features commonly used on modern networks, such as jumbo frames, multi-queue support (also known as Receive Side Scaling in Windows), IPv6 offloads, and MSI/MSI-X interrupt delivery and hardware offloads.

Ensure the adapter type is VMXNET3.

Network adapter 1	Storage Controller Managemen
Status	<input checked="" type="checkbox"/> Connect At Power On
Adapter Type	VMXNET 3
DirectPath I/O	<input checked="" type="checkbox"/> Enable

Network adapter.

## RSS Check

---

**Note:** This check only applies to the guest virtual machines that are running a Windows operating system.

---

Receive side scaling (RSS) is a network driver technology that enables the efficient distribution of network receive processing across multiple CPUs in multiprocessor systems.

Windows servers have a driver configuration that enables the distribution of the kernel-mode network processing load across multiple CPUs.

Verify if it is enabled run this command on the Windows PowerShell:

```
netsh interface tcp set global rss=enabled
```

In order to enable RSS review this [link](#)

## CPU Hot-Plug Check

CPU hotplug is a feature that enables the VM administrator to add CPUs to the VM without having to power it off. This allows adding CPU resources on the fly with no disruption to service. When CPU hotplug is enabled on a VM, the vNUMA capability is disabled.

# Edit Settings | centos8

Virtual Hardware | VM Options

ADD NEW D

▼ CPU	1 ▼	
Cores per Socket	1 ▼	Sockets: 1
CPU Hot Plug	<input type="checkbox"/> Enable CPU Hot Add	
Reservation	0	MHz ▼
Limit	Unlimited	MHz ▼
Shares	Normal ▼	1000 ▼
Hardware virtualization	<input type="checkbox"/> Expose hardware assisted virtualization to the guest OS	
Performance Counters	<input type="checkbox"/> Enable virtualized CPU performance counters	
I/O MMU	<input type="checkbox"/> Enabled	
> Memory	2	GB ▼
> Hard disk 1	16	GB ▼
> SCSI controller 0	VMware Paravirtual	

CANCEL

*CPU hot plug disabled.*

Review the best practices for common Operating systems and applications:

Windows.

[Performance Tuning Guidelines for Windows Server 2022.](#)

Red Hat.

[3 tips for Linux process performance improvement with priority and affinity.](#)

SQL Server.

[Architecting Microsoft SQL Server on VMware.](#)

RedHat.

[Performance Tuning Guide.](#)

---

: CRC errors and MTU mismatch especially in the storage network generate latency issues. Storage Traffic must use Jumbo Frames.

---

## Storage I/O Control and Queue Depth Check

Storage I/O Control (SIOC) is used to control the I/O usage of a virtual machine and to gradually enforce the predefined I/O share levels, is necessary to have this feature disabled in Hyperflex Clusters.

Queue depth is the number of pending input/output ([I/O](#)) requests that a storage resource can handle at any one time.

You can use these steps to verify SIOC is disabled and that the Queue Depth configuration.

### Confirm that SIOC is running on ESXi and the Queue Depth Configuration

Step 1. SSH to an HX ESXi host and issue the command to list the datastores.

```
[root@] vsish -e ls /vmkModules/nfsclient/mnt
encrypted_app/
Prod/           <----- Datastore name
Dev/
App/
```

Step 2. Use the Data Store name and issue the command.

```
vsish -e get /vmkModules/nfsclient/mnt/<datastore name>/properties
```

```
[root@] vsish -e get /vmkModules/nfsclient/mnt/Prod/properties
mount point information {
  volume name:Prod
  server name:7938514614702552636-8713662604223381594
  server IP:127.0.0.1
  server volume:172.16.3.2:Prod
  UUID:63dee313-dfecdf62
  client src port:641
  busy:0
  socketSendSize:1048576
  socketReceiveSize:1048576
  maxReadTransferSize:65536
  maxWriteTransferSize:65536
  reads:0
  readsFailed:0
  writes:285
  writesFailed:0
  readBytes:0
  writeBytes:10705
  readTime:0
  writeTime:4778777
  readSplitsIssued:0
  writeSplitsIssued:285
```

```

readIssueTime:0
writeIssueTime:4766494
cancels:0
totalReqsQueued:0
metadataReqsQueued(non IO):0
reqsInFlight:0
readOnly:0
hidden:0
isPE:0
isMounted:1
isAccessible:1
unstableWrites:0
unstableNoCommit:0
maxQDepth:1024      <----- Max Qdepth configuration
iormState:0         <----- I/O control disabled
latencyThreshold:30
shares:52000
podID:0
iormInfo:0
NFS operational state: 0 -> Up
enableDnlc:1
closeToOpenCache:0
highToAvgLatRatio:10
latMovingAvgSmoothingLevel:2
activeWorlds:55
inPreUnmount:0
}

```

Step 3. In the output look for the line

```
iormState:0 0= disabled 2= enabled
```

The line maxQDepth must be 1024

Step 4. The same steps must be repeated for the rest of the Datastores

## Disable SIOC

In order to disable the SIOC run these steps.

Step 1. Login to vsphere using the HTML client.

Step 2. From the drop-down menu, select Storage and then select the applicable HX Datastore in the left pane.

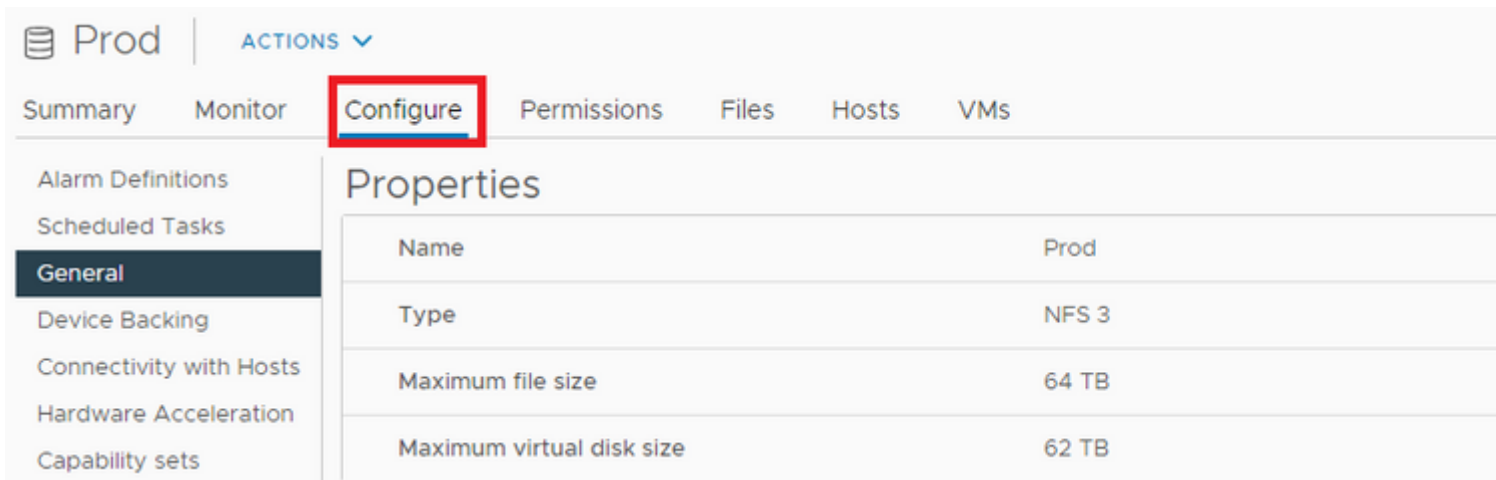




Select datastore.

â€f

Step 3. In the right pane top section of the Datastore, select the configure tab.



Configure tab.

â€f

Step 4. In the right pane middle section Under More, select General, and on the right side scroll down to DataStore Capabilities and click Edit



```
vsish -e set /vmkModules/nfsclient/mnt/<yourdatastorename>/properties maxQDepth 1024
```

## Check for Rx\_no\_Buff

Hyperflex servers with heavy network traffic or network traffic with microbursts can lead to packet loss seen in the form of rx\_no\_bufs.

To identify this Issue run these commands in the ESXi host to check the rx\_no\_buf counters.

```
/usr/lib/vmware/vm-support/bin/nicinfo.sh | egrep "^NIC:|rx_no_buf"  
NIC: vmnic0  
rx_no_bufs: 1  
NIC: vmnic1  
rx_no_bufs: 2  
NIC: vmnic2  
rx_no_bufs: 2  
NIC: vmnic3  
rx_no_bufs: 71128211 <-----Very high rx_no_bufs counter  
NIC: vmnic4  
rx_no_bufs: 1730  
NIC: vmnic5  
rx_no_bufs: 897  
NIC: vmnic6  
rx_no_bufs: 24952  
NIC: vmnic7  
rx_no_bufs: 2
```

Wait a few minutes and run the command again and check if the rx\_no\_bufs counters are not increasing.

- If these counters are low (< 1,000) then with little packet loss due to the default queue configuration and likely no tuning needs to be done.
- If these counters are high (> 10,000) then there is some impact due to this queue configuration and tuning could help a little.
- If these counters are very high (> 1,000,000) then there is a more significant impact, increasing queues is highly recommended.
- If rx\_no\_bufs are actively incrementing, this means the packet made it all the way across the network and arrives at the virtualized layer then the packet is dropped.

If you see the counter on these values please contact Cisco TAC to tune the vNIC configuration for better performance.

Review the best practices and additional checks at ESXI Level.

[Performance Best Practices for VMware vSphere 7.0.](#)

# Identify Performance Issues Storage Controller Virtual Machine (SCVM) Level

## Cluster Health

Verify if the cluster is healthy.

```
hxshell:~$ sysmtool --ns cluster --cmd healthdetail
Cluster Health Detail:
-----:
State: ONLINE                               <----- State of the cluster
HealthState: HEALTHY                       <----- Health of the cluster
Policy Compliance: COMPLIANT
Creation Time: Tue May 30 04:48:45 2023
Uptime: 7 weeks, 19 hours, 45 mins, 51 secs
Cluster Resiliency Detail:
-----:
Health State Reason: Storage cluster is healthy.
# of nodes failure tolerable for cluster to be fully available: 1
# of node failures before cluster goes into readonly: NA
# of node failures before cluster goes to be critical and partially available: 3
# of node failures before cluster goes to enospace warn trying to move the existing data: NA
# of persistent devices failures tolerable for cluster to be fully available: 2
# of persistent devices failures before cluster goes into readonly: NA
# of persistent devices failures before cluster goes to be critical and partially available: 3
# of caching devices failures tolerable for cluster to be fully available: 2
# of caching failures before cluster goes into readonly: NA
# of caching failures before cluster goes to be critical and partially available: 3
Current ensemble size: 3
Minimum data copies available for some user data: 3
Minimum cache copies remaining: 3
Minimum metadata copies available for cluster metadata: 3
Current healing status:
Time remaining before current healing operation finishes:
# of unavailable nodes: 0

hxshell:~$
```

This output shows an unhealthy cluster due to an unavailable node.

```
hxshell:~$ sysmtool --ns cluster --cmd healthdetail
Cluster Health Detail:
-----:
State: ONLINE                               <-----State of the cluster
HealthState: UNHEALTHY                     <-----Health of the cluster
Policy Compliance: NON-COMPLIANT
Creation Time: Tue May 30 04:48:45 2023
Uptime: 7 weeks, 19 hours, 55 mins, 9 secs
Cluster Resiliency Detail:
-----:
Health State Reason: Storage cluster is unhealthy.Storage node 172.16.3.9 is unavailable.
# of nodes failure tolerable for cluster to be fully available: 0
# of node failures before cluster goes into readonly: NA
```

```

# of node failures before cluster goes to be critical and partially available: 2
# of node failures before cluster goes to enospace warn trying to move the existing data: NA
# of persistent devices failures tolerable for cluster to be fully available: 1
# of persistent devices failures before cluster goes into readonly: NA
# of persistent devices failures before cluster goes to be critical and partially available: 2
# of caching devices failures tolerable for cluster to be fully available: 1
# of caching failures before cluster goes into readonly: NA
# of caching failures before cluster goes to be critical and partially available: 2
Current ensemble size: 3
Minimum data copies available for some user data: 2
Minimum cache copies remaining: 2
Minimum metadata copies available for cluster metadata: 2
Current healing status: Rebuilding/Healing is needed, but not in progress yet. Warning: Insufficient nodes
Time remaining before current healing operation finishes:
# of unavailable nodes: 1

hxshell:~$

```

This output shows an unhealthy cluster due to the rebuilding.

```

Cluster Health Detail:
-----:
State: ONLINE
HealthState: UNHEALTHY
Policy Compliance: NON-COMPLIANT
Creation Time: Tue May 30 04:48:45 2023
Uptime: 7 weeks, 20 hours, 2 mins, 4 secs
Cluster Resiliency Detail:
-----:
Health State Reason: Storage cluster is unhealthy.
# of nodes failure tolerable for cluster to be fully available: 1
# of node failures before cluster goes into readonly: NA
# of node failures before cluster goes to be critical and partially available: 2
# of node failures before cluster goes to enospace warn trying to move the existing data: NA
# of persistent devices failures tolerable for cluster to be fully available: 1
# of persistent devices failures before cluster goes into readonly: NA
# of persistent devices failures before cluster goes to be critical and partially available: 2
# of caching devices failures tolerable for cluster to be fully available: 1
# of caching failures before cluster goes into readonly: NA
# of caching failures before cluster goes to be critical and partially available: 2
Current ensemble size: 3
Minimum data copies available for some user data: 3
Minimum cache copies remaining: 2
Minimum metadata copies available for cluster metadata: 2
Current healing status: Rebuilding is in progress, 58% completed.
Time remaining before current healing operation finishes: 18 hr(s), 10 min(s), and 53 sec(s)
# of unavailable nodes: 0

```

These commands show an overall summary of the health of the cluster and let you know if there is something is affecting the operation of the cluster, for instance, if there is a blacklisted disk, an offline node, or if the cluster is healing.

## Nodes Participating In I/O

The performance can be impacted by a node not participating in the input and output operations, to check



```
bash-4.2# stcli cleaner start
WARNING: This command should be executed ONLY by Cisco TAC support as it may have very severe consequences
bash-4.2# stcli cleaner info
{ 'type': 'node', 'id': '1f82077d-6702-214d-8814-e776ffc0f53c', 'name': '172.16.3.7' }: ONLINE
{ 'type': 'node', 'id': 'c4a24480-e935-6942-93ee-987dc8e9b5d9', 'name': '172.16.3.8' }: ONLINE
{ 'type': 'node', 'id': '50a5dc5d-c419-9c48-8914-d91a98d43fe7', 'name': '172.16.3.9' }: ONLINE
bash-4.2#
```

---

**Caution:** This command must be executed with Cisco TAC approval.

---

## Rebalance

The storage cluster is rebalanced on a regular schedule. It is used to realign the distribution of stored data across changes in available storage and to restore storage cluster health.

Rebalance runs in clusters for different reasons:

- A physical resource (node/disk) is down and HX is relocating those Vnodes to a different physical resource in the cluster.
- The individual drives across the cluster are not all comparably utilized and therefore some hotspots have been created in terms of data availability within the HX cluster (data collocation).
- Rebalance can also run if Zone compliance is not there even if the cluster is healthy.
- When a new node is added to the existing cluster, the added node(s) take on new writes as soon as it joins the existing cluster.

Verify that the cluster has rebalance enabled.

```
hxshell:~$ stcli rebalance status
rebalanceStatus:
  percentComplete: 0
  rebalanceState: cluster_rebalance_not_running
rebalanceEnabled: True    <-----Rebalance should be enabled
hxshell:~$
```

---

**Caution:** Any operation related to Rebalance must be done with Cisco TAC approval.

---

## Disk Failure

For proper operation, the cluster must not have any blacklisted disks or offline resources.

You need to check if there is any blacklisted disk on the cluster in the HX Connect interface.

## Hyperconverged Nodes

Node	Hypervisor	HyperFlex Controller	Disk Overview ( 25 in use   1 empty slots )
HX240C-M5SX	<span style="color: green;">✔</span> Online	<span style="color: green;">✔</span> Online	

**Slot Number 4** Persistence

Disk State ❌ BLACKLISTED

Locator LED

Used / Total Capacity 1013 GB of 2.2 TB

Serial Number

Storage Usage  45%

Version CK04

Disk Drive Interface SAS

*Blacklisted disk.*

Check on the CLI for any offline resources on each Converge Node.

```

sysmtool --ns cluster --cmd offlineresources
UUID                               Type      State      InUse      Last modified
----                               -
000cca0b019b4a80:0000000000000000  DISK     DELETED    YES        <----- Offline disk
5002538c405e0bd1:0000000000000000  DISK     BLOCKLISTED NO        <----- Blacklisted disk
5002538c405e299e:0000000000000000  DISK     DELETED    NO
Total offline resources: 3, Nodes: 0, Disks: 3
  
```

Verify if there are any blacklisted resources.

```

hxshell:~$ sysmtool --ns disk --cmd list | grep -i blacklist
Blacklist Count: 0
Blacklist Count: 0
Blacklist Count: 0
Blacklist Count: 0
State: BLACKLISTED
Blacklist Count: 5
Blacklist Count: 0
Blacklist Count: 0
  
```

```

admin:~$ cat /var/log/springpath/diskslotmap-v2.txt
0.0.1:5002538e000d59a3:Samsung:SAMSUNG_MZ7LH3T8HMLT-00003:S4F3NY0M302248:HXT76F3Q:SATA:SSD:3662830:Inactive
1.0.2:5002538c40be79ac:Samsung:SAMSUNG_MZ7LM240HMHQ-00003:S4EGNX0KC04551:GXT51F3Q:SATA:SSD:228936:Active
1.0.3:5002538e000d599e:Samsung:SAMSUNG_MZ7LH3T8HMLT-00003:S4F3NY0M302243:HXT76F3Q:SATA:SSD:3662830:Active
1.0.4:5002538e000d59a0:Samsung:SAMSUNG_MZ7LH3T8HMLT-00003:S4F3NY0M302245:HXT76F3Q:SATA:SSD:3662830:Active
1.0.5:5002538e000eb00b:Samsung:SAMSUNG_MZ7LH3T8HMLT-00003:S4F3NY0M302480:HXT76F3Q:SATA:SSD:3662830:Active
1.0.6:5002538e000d599b:Samsung:SAMSUNG_MZ7LH3T8HMLT-00003:S4F3NY0M302240:HXT76F3Q:SATA:SSD:3662830:Active
1.0.7:5002538e000d57f6:Samsung:SAMSUNG_MZ7LH3T8HMLT-00003:S4F3NY0M301819:HXT76F3Q:SATA:SSD:3662830:Active
1.0.8:5002538e000d59ab:Samsung:SAMSUNG_MZ7LH3T8HMLT-00003:S4F3NY0M302256:HXT76F3Q:SATA:SSD:3662830:Active
1.0.9:5002538e000d59a1:Samsung:SAMSUNG_MZ7LH3T8HMLT-00003:S4F3NY0M302246:HXT76F3Q:SATA:SSD:3662830:Active
1.0.10:5002538e0008c68f:Samsung:SAMSUNG_MZ7LH3T8HMLT-00003:S4F3NY0M200500:HXT76F3Q:SATA:SSD:3662830:Active
0.1.192:000cca0b01c83180:HGST:UCSC-NVMEHW-H1600:SDM000026904:KNCCD111:NVMe:SSD:1526185:Active:/dev/nvme0n1
admin:~$

```

Example of a Node without any disk failure.

```

hxshell:~$ sysmtool --ns cluster --cmd offlineresources
No offline resources found <----- No offline resources

hxshell:~$ sysmtool --ns disk --cmd list | grep -i blacklist
hxshell:~$ <----- No blacklisted disks
hxshell:~$ cat /var/log/springpath/diskslotmap-v2.txt
1.14.1:55cd2e404c234bf9:Intel:INTEL_SSDSC2BX016T4K:BTHC618505B51P6PGN:G201CS01:SATA:SSD:1526185:Active:/dev/sda
1.14.2:5000c5008547c543:SEAGATE:ST1200MM0088:Z4009D7Y0000R637KMU7:N0A4:SAS:10500:1144641:Active:/dev/sdb
1.14.3:5000c5008547be1b:SEAGATE:ST1200MM0088:Z4009G0B0000R635L4D3:N0A4:SAS:10500:1144641:Active:/dev/sdc
1.14.4:5000c5008547ca6b:SEAGATE:ST1200MM0088:Z4009F9N0000R637JZRF:N0A4:SAS:10500:1144641:Active:/dev/sdd
1.14.5:5000c5008547b373:SEAGATE:ST1200MM0088:Z4009GPM0000R634ZJHB:N0A4:SAS:10500:1144641:Active:/dev/sde
1.14.6:5000c500854310fb:SEAGATE:ST1200MM0088:Z4008XFJ0000R6374ZE8:N0A4:SAS:10500:1144641:Active:/dev/sdf
1.14.7:5000c50085424b53:SEAGATE:ST1200MM0088:Z4008D2S0000R635M4VF:N0A4:SAS:10500:1144641:Active:/dev/sdg
1.14.8:5000c5008547bcfb:SEAGATE:ST1200MM0088:Z4009G3W0000R637K1R8:N0A4:SAS:10500:1144641:Active:/dev/sdh
1.14.9:5000c50085479abf:SEAGATE:ST1200MM0088:Z4009J510000R637KL1V:N0A4:SAS:10500:1144641:Active:/dev/sdi
1.14.11:5000c5008547c2c7:SEAGATE:ST1200MM0088:Z4009FR0000R637JPEQ:N0A4:SAS:10500:1144641:Active:/dev/sdj
1.14.13:5000c5008547ba93:SEAGATE:ST1200MM0088:Z4009G8V0000R634ZKLX:N0A4:SAS:10500:1144641:Active:/dev/sdk
1.14.14:5000c5008547b69f:SEAGATE:ST1200MM0088:Z4009GG80000R637KM30:N0A4:SAS:10500:1144641:Active:/dev/sdl
1.14.15:5000c5008547b753:SEAGATE:ST1200MM0088:Z4009GH90000R635L5F6:N0A4:SAS:10500:1144641:Active:/dev/sdm
1.14.16:5000c5008547ab7b:SEAGATE:ST1200MM0088:Z4009H3P0000R634ZK8T:N0A4:SAS:10500:1144641:Active:/dev/sdn
hxshell:~$

```

## Free Memory

Check the free memory with this command, the free memory must be more than 2048 MB (free +cache).

```

hxshell:~$ free -m

```



	total	used	free	shared	buff/cache	available
Mem:	74225624	32194300	38893712	1672	3137612	41304336
Swap:	0	0	0			

hxshell:~\$

if the free + cache memory is less than 2048, is necessary to identify the process that is generating the Out Of Memory condition.

---

**Note:** You can use the **top** command to identify processes that consume a lot of memory, however, any changes have to be done with TAC approval, contact Cisco TAC to troubleshoot OOM conditions.

---

## End Of Space Condition

The best practice of storage cluster space utilization is to not go beyond 76 percent at the HX Connect capacity view. Beyond 76 percent, usage at the HX Connect capacity view results in performance degradation.

If the storage cluster is experiencing an ENOSPC condition, the cleaner automatically runs at high priority, which can create performance issues in the cluster, the priority is determined by cluster space usage.

If the storage cluster reaches an ENOSPC WARN condition, the cleaner increases its intensity by increasing the number of I/O to collect garbage with an ENOSPC set condition, it runs at the highest priority.

You can check the ENOSPCINFO status on the cluster with this command.

```

hxshell:~$ sysmtool --ns cluster --cmd enospcinfo
Cluster Space Details:
-----:
Cluster state: ONLINE
Health state: HEALTHY
Raw capacity: 42.57T
Usable capacity: 13.06T
Used capacity: 163.08G
Free capacity: 12.90T
Enospc state: ENOSPACE_CLEAR <-----End of space status
Space reclaimable: 0.00
Minimum free capacity
required to resume operation: 687.12G
Space required to clear
ENOSPC warning: 2.80T <-----Free space until the end of space warning appears
Rebalance In Progress: NO
Flusher in progress: NO
Cleaner in progress: YES
Disk Enospace: NO

hxshell:~$

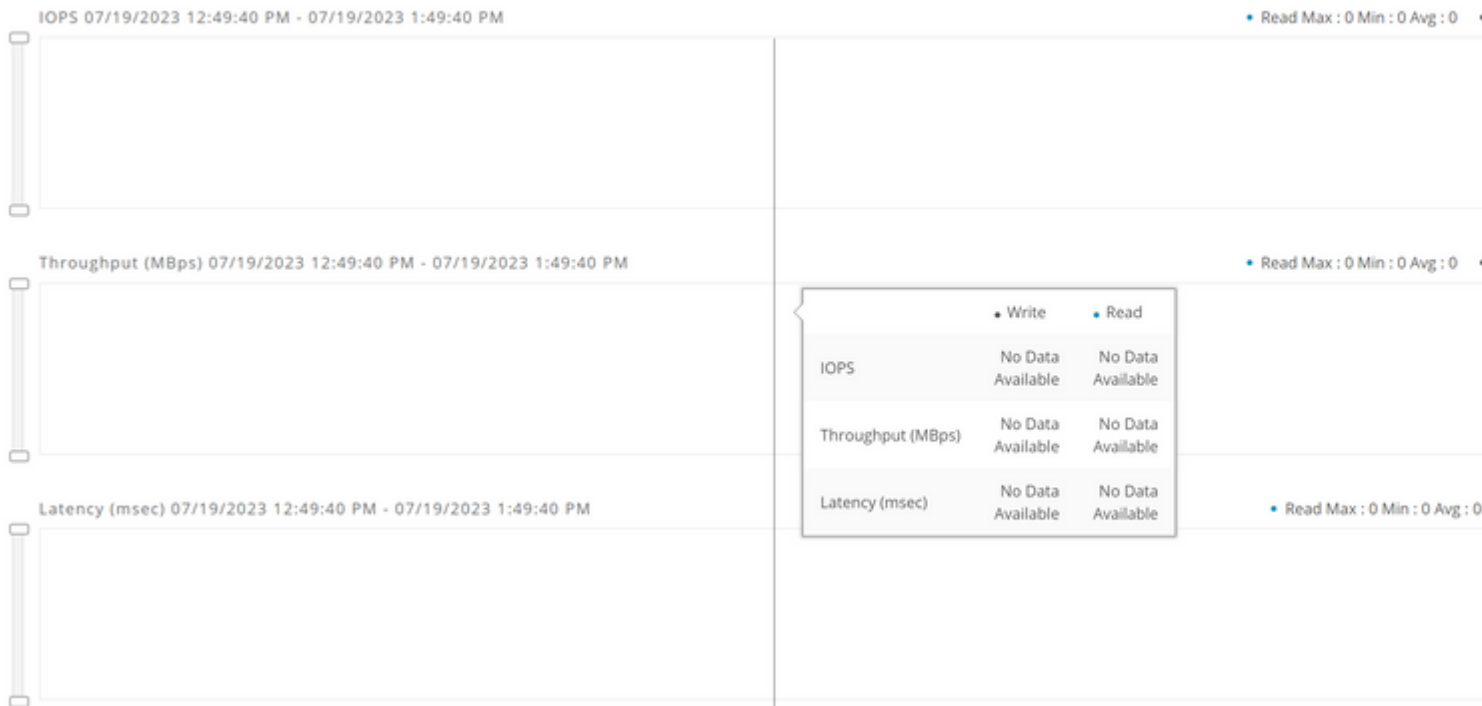
```

Review the [Capacity Management in Cisco HyperFlex White Paper](#) to identify the best practice to manage

the space on your Hyperflex Cluster.

## **Troubleshooting Performance Charts**

Sometimes the hyperflex performance charts are not displaying information.



Hyperflex performance charts.

â€f

If you face this behavior you need to review if the stats services are running in the cluster.

```

hxshell:~$ priv service carbon-cache status
carbon-cache stop/waiting

hxshell:~$ priv service carbon-aggregator status
carbon-aggregator stop/waiting

hxshell:~$ priv service statsd status
statsd stop/waiting
  
```

If the processes are not running, manually start the services.

```

hxshell:~$ priv service carbon-cache start
carbon-cache start/running, process 15750

hxshell:~$ priv service carbon-aggregator start
carbon-aggregator start/running, process 15799

hxshell:~$ priv service statsd start
statsd start/running, process 15855
  
```

## Related Information

- [Capacity Management in Cisco HyperFlex White Paper](#)
- [Mission-Critical Workload Performance Testing of Different Hyperconverged Approaches on th](#)