



## ESXi

### SAN hosts and cloud clients

NetApp  
February 26, 2024

# Table of Contents

- ESXi ..... 1
  - Use VMware vSphere 8.x with ONTAP ..... 1
  - Use VMware vSphere 7.x with ONTAP ..... 10
  - Use VMware vSphere 6.5 and 6.7 with ONTAP ..... 18

# ESXi

## Use VMware vSphere 8.x with ONTAP

You can configure ONTAP SAN host settings for the VMware vSphere 8.x release with FC, FCoE, and iSCSI protocols.

### Hypervisor SAN booting

#### What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the [NetApp Interoperability Matrix Tool](#) to verify that your OS, HBA, HBA firmware and the HBA boot BIOS, and ONTAP version are supported.

#### Steps

1. Map the SAN boot LUN to the host.
2. Verify that multiple paths are available.



Multiple paths become available after the host OS is up and running on the paths.

3. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

For information on how to enable the HBA BIOS, see your vendor-specific documentation.

4. Reboot the host to verify that the boot is successful.

### Multipathing

ESXi provides an extensible multipathing module called Native Multipathing Plug-In (NMP) that manages the sub-plugins, Storage Array Type Plugins (SATPs) and Path Selection Plugins (PSPs). By default, these SATP rules are available in ESXi.

For NetApp ONTAP storage, `VMW_SATP_ALUA` plugin is used by default with `VMW_PSP_RR` as a path selection policy (PSP). You can confirm by using the following command:

```
`esxcli storage nmp satp rule list -s VMW_SATP_ALUA`
```

Example output:

Name	Device	Vendor	Model	Driver	Transport	Options
-----						
VMW_SATP_ALUA		LSI	INF-01-00			
reset_on_attempted_reserve		system				
VMW_SATP_ALUA		NETAPP				
reset_on_attempted_reserve		system				
Rule Group	Claim Options	Default PSP	PSP Options	Description		
-----						
tpgs_on	VMW_PSP_MRU			NetApp E-Series arrays with		
ALUA support						
tpgs_on	VMW_PSP_RR			NetApp arrays with ALUA		
support						

## Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

### Example

The following example displays the correct output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

```
# esxcli storage nmp device list -d naa.600a0980383148693724545244395855
```

Example output:

```

naa.600a0980383148693724545244395855
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a0980383148693724545244395855)
  Storage Array Type: VMW_SATP_ALUA
  Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off;
{TPG_id=1000,TPG_state=ANO}{TPG_id=1001,TPG_state=AO}}
  Path Selection Policy: VMW_PSP_RR
  Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0; lastPathIndex=1:
NumIOsPending=0,numBytesPending=0}
  Path Selection Policy Device Custom Config:
  Working Paths: vmhba4:C0:T0:L11, vmhba3:C0:T0:L11
  Is USB: false

```

```
# esxcli storage nmp path list -d naa.600a0980383148693724545244395855
```

#### Example output:

```

fc.20000024ff7f4a51:21000024ff7f4a51-fc.2009d039ea3ab21f:2003d039ea3ab21f-
naa.600a0980383148693724545244395855
  Runtime Name: vmhba4:C0:T0:L11
  Device: naa.600a0980383148693724545244395855
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a0980383148693724545244395855)
  Group State: active
  Array Priority: 0
  Storage Array Type Path Config: {TPG_id=1001,
TPG_state=AO,RTP_id=4,RTP_health=UP}
  Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff7f4a50:21000024ff7f4a50-fc.2009d039ea3ab21f:2002d039ea3ab21f-
naa.600a0980383148693724545244395855
  Runtime Name: vmhba3:C0:T0:L11
  Device: naa.600a0980383148693724545244395855
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a0980383148693724545244395855)
  Group State: active
  Array Priority: 0
  Storage Array Type Path Config: {TPG_id=1001,
TPG_state=AO,RTP_id=3,RTP_health=UP}

```

```
Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path configuration.
```

```
fc.20000024ff7f4a51:21000024ff7f4a51-fc.2009d039ea3ab21f:2001d039ea3ab21f-naa.600a0980383148693724545244395855
```

```
Runtime Name: vmhba4:C0:T3:L11
```

```
Device: naa.600a0980383148693724545244395855
```

```
Device Display Name: NETAPP Fibre Channel Disk  
(naa.600a0980383148693724545244395855)
```

```
Group State: active unoptimized
```

```
Array Priority: 0
```

```
Storage Array Type Path Config: {TPG_id=1000,  
TPG_state=ANO,RTP_id=2,RTP_health=UP}
```

```
Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path configuration.
```

```
fc.20000024ff7f4a50:21000024ff7f4a50-fc.2009d039ea3ab21f:2000d039ea3ab21f-naa.600a0980383148693724545244395855
```

```
Runtime Name: vmhba3:C0:T3:L11
```

```
Device: naa.600a0980383148693724545244395855
```

```
Device Display Name: NETAPP Fibre Channel Disk  
(naa.600a0980383148693724545244395855)
```

```
Group State: active unoptimized
```

```
Array Priority: 0
```

```
Storage Array Type Path Config: {TPG_id=1000,  
TPG_state=ANO,RTP_id=1,RTP_health=UP}
```

```
Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path configuration.
```

## All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

### Example

The following example displays the correct output for an ONTAP LUN:

```
esxcli storage nmp device list -d naa.600a098038304759563f4e7837574453
```

Example output:

```

naa.600a098038314962485d543078486c7a
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038314962485d543078486c7a)
  Storage Array Type: VMW_SATP_ALUA
  Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off;
{TPG_id=1001,TPG_state=AO}{TPG_id=1000,TPG_state=AO}}
  Path Selection Policy: VMW_PSP_RR
  Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0; lastPathIndex=3:
NumIOsPending=0,numBytesPending=0}
  Path Selection Policy Device Custom Config:
  Working Paths: vmhba4:C0:T0:L14, vmhba4:C0:T1:L14, vmhba3:C0:T0:L14,
vmhba3:C0:T1:L14
  Is USB: false

```

```
# esxcli storage nmp path list -d naa.600a098038314962485d543078486c7a
```

#### Example output:

```

fc.200034800d756a75:210034800d756a75-fc.2018d039ea936319:2015d039ea936319-
naa.600a098038314962485d543078486c7a
  Runtime Name: vmhba4:C0:T0:L14
  Device: naa.600a098038314962485d543078486c7a
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038314962485d543078486c7a)
  Group State: active
  Array Priority: 0
  Storage Array Type Path Config: {TPG_id=1000,
TPG_state=AO,RTP_id=2,RTP_health=UP}
  Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.200034800d756a75:210034800d756a75-fc.2018d039ea936319:2017d039ea936319-
naa.600a098038314962485d543078486c7a
  Runtime Name: vmhba4:C0:T1:L14
  Device: naa.600a098038314962485d543078486c7a
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038314962485d543078486c7a)
  Group State: active
  Array Priority: 0
  Storage Array Type Path Config: {TPG_id=1001,

```

```
TPG_state=AO,RTP_id=4,RTP_health=UP}
```

```
Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path configuration.
```

```
fc.200034800d756a74:210034800d756a74-fc.2018d039ea936319:2014d039ea936319-naa.600a098038314962485d543078486c7a
```

```
Runtime Name: vmhba3:C0:T0:L14
```

```
Device: naa.600a098038314962485d543078486c7a
```

```
Device Display Name: NETAPP Fibre Channel Disk  
(naa.600a098038314962485d543078486c7a)
```

```
Group State: active
```

```
Array Priority: 0
```

```
Storage Array Type Path Config: {TPG_id=1000,
```

```
TPG_state=AO,RTP_id=1,RTP_health=UP}
```

```
Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path configuration.
```

```
fc.200034800d756a74:210034800d756a74-fc.2018d039ea936319:2016d039ea936319-naa.600a098038314962485d543078486c7a
```

```
Runtime Name: vmhba3:C0:T1:L14
```

```
Device: naa.600a098038314962485d543078486c7a
```

```
Device Display Name: NETAPP Fibre Channel Disk  
(naa.600a098038314962485d543078486c7a)
```

```
Group State: active
```

```
Array Priority: 0
```

```
Storage Array Type Path Config: {TPG_id=1001,
```

```
TPG_state=AO,RTP_id=3,RTP_health=UP}
```

```
Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path configuration.
```

## vVol

Virtual Volumes (vVols) are a VMware object type that corresponds to a Virtual Machine (VM) disk, its snapshots, and fast clones.

ONTAP tools for VMware vSphere includes the VASA Provider for ONTAP, which provides the integration point for a VMware vCenter to leverage vVols based storage. When you deploy the ONTAP tools Open Virtualization Appliance (OVA), it is automatically registered with the vCenter server and enables the VASA Provider.

When you create a vVols datastore using the vCenter user interface, it guides you to create FlexVols as backup storage for the datastore. vVols within vVols datastores are accessed by ESXi hosts using a protocol endpoint (PE). In SAN environments, one 4MB LUN is created on each FlexVol in the datastore for use as a PE. A SAN PE is an administrative logical unit (ALU). vVols are subsidiary logical units (SLUs).

Standard requirements and best practices for SAN environments apply when using vVols, including (but not limited to) the following:

- Create at least one SAN LIF on each node per SVM you intend to use. The best practice is to create at least two per node, but no more than necessary.



- Eliminate any single point of failure. Use multiple VMkernel network interfaces on different network subnets that use NIC teaming when multiple virtual switches are used, or use multiple physical NICs connected to multiple physical switches to provide HA and increased throughput.
- Configure zoning, VLANs, or both as required for host connectivity.
- Verify that all required initiators are logged into the target LIFs on the desired SVM.



You must deploy ONTAP tools for VMware vSphere to enable the VASA Provider. The VASA Provider will manage all of your iGroup settings for you, therefore there is no need to create or manage iGroups in a vVols environment.

NetApp does not recommend changing any vVols settings from default at this time.

Refer to the [NetApp Interoperability Matrix Tool](#) for specific versions of ONTAP tools, or legacy VASA Provider for your specific versions of vSphere and ONTAP.

For detailed information on provisioning and managing vVols, refer to ONTAP tools for VMware vSphere documentation, [TR-4597](#), and [TR-4400](#).

## Recommended settings

### ATS locking

ATS locking is **mandatory** for VAAI compatible storage and upgraded VMFS5 and is required for proper interoperability and optimal VMFS shared storage I/O performance with ONTAP LUNs. Refer to VMware documentation for details on enabling ATS locking.

Settings	Default	ONTAP Recommended	Description
HardwareAcceleratedLocking	1	1	Helps enable the use of Atomic Test and Set (ATS) locking
Disk IOPs	1000	1	IOPS limit: The Round Robin PSP defaults to an IOPS limit of 1000. In this default case, a new path is used after 1000 I/O operations are issued.
Disk/QFullSampleSize	0	32	The count of QUEUE FULL or BUSY conditions it takes before ESXi starts throttling.



Enable `space-alloc` setting for all the LUNs mapped to VMware vSphere for UNMAP to work. For more details, refer to ONTAP Documentation.

### Guest OS timeouts

You can manually configure the virtual machines with the recommended guest OS tunings. After tuning updates, you must reboot the guest for the updates to take effect.

### GOS timeout values:

Guest OS Type	Timeouts
Linux variants	disk timeout = 60
Windows	disk timeout = 60
Solaris	disk timeout = 60 busy retry = 300 not ready retry = 300 reset retry = 30 max.throttle = 32 min.throttle = 8

### Validate the vSphere tunable

You can use the following command to verify the HardwareAcceleratedLocking setting.

```
esxcli system settings advanced list --option /VMFS3/HardwareAcceleratedLocking
```

```

Path: /VMFS3/HardwareAcceleratedLocking
Type: integer
Int Value: 1
Default Int Value: 1
Min Value: 0
Max Value: 1
String Value:
Default String Value:
Valid Characters:
Description: Enable hardware accelerated VMFS locking (requires
compliant hardware). Please see http://kb.vmware.com/kb/2094604 before
disabling this option.
```

### Validate the Disk IOPs setting

You can use the following command to verify the IOPs setting.

```
esxcli storage nmp device list -d naa.600a098038304731783f506670553355
```

```

naa.600a098038304731783f506670553355
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304731783f506670553355)
  Storage Array Type: VMW_SATP_ALUA
  Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off;
{TPG_id=1000,TPG_state=ANO}{TPG_id=1001,TPG_state=AO}}
  Path Selection Policy: VMW_PSP_RR
  Path Selection Policy Device Config: {policy=rr,
iops=1,bytes=10485760,useANO=0; lastPathIndex=0:
NumIOsPending=0,numBytesPending=0}
  Path Selection Policy Device Custom Config:
  Working Paths: vmhba4:C0:T0:L82, vmhba3:C0:T0:L82
  Is USB: false

```

## Validate the QFullSampleSize

You can use the following command to verify the QFullSampleSize.

```
esxcli system settings advanced list --option /Disk/QFullSampleSize
```

```

Path: /Disk/QFullSampleSize
Type: integer
Int Value: 32
Default Int Value: 0
Min Value: 0
Max Value: 64
String Value:
Default String Value:
Valid Characters:
Description: Default I/O samples to monitor for detecting non-transient
queue full condition. Should be nonzero to enable queue depth throttling.
Device specific QFull options will take precedence over this value if set.

```

## Known issues

The VMware vSphere 8.x with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description
1543660	I/O error occurs when Linux VMs using vNVMe adapters encounter a long all paths down (APD) window	Linux VMs running vSphere 8.x and later and using virtual NVMe (vNVME) adapters encounter an I/O error because the vNVMe retry operation is disabled by default. To avoid a disruption on Linux VMs running older kernels during an all paths down (APD) or a heavy I/O load, VMware has introduced a tunable "VSCSIDisableNvmeRetry" to disable the vNVMe retry operation.

#### Related information

- [TR-4597-VMware vSphere with ONTAP](#)
- [VMware vSphere 5.x, 6.x and 7.x support with NetApp MetroCluster \(2031038\)](#)
- [NetApp ONTAP with NetApp SnapMirror Business Continuity \(SM-BC\) with VMware vSphere Metro Storage Cluster \(vMSC\)](#)

## Use VMware vSphere 7.x with ONTAP

You can use ONTAP SAN host configuration settings for the vSphere 7.x release with FC, FCoE and iSCSI protocols.

### Hypervisor SAN Booting

#### What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the [NetApp Interoperability Matrix Tool](#) to verify that your OS, HBA, HBA firmware and the HBA boot BIOS, and ONTAP version are supported.

#### Steps

1. Map the SAN boot LUN to the host.
2. Verify that multiple paths are available.



Multiple paths become available after the host OS is up and running on the paths.

3. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

For information on how to enable the HBA BIOS, see your vendor-specific documentation.

4. Reboot the host to verify that the boot is successful.

### Multipathing

ESXi provides an extensible multipathing module called Native Multipathing Plug-In (NMP) that manages the sub-plugins Storage Array Type Plugins (SATPs), and Path Selection Plugins (PSPs). These SATP rules are available by default in ESXi.

For NetApp ONTAP storage, VMW\_SATP\_ALUA plugin is used by default with VMW\_PSP\_RR as a path selection policy (PSP). This can be confirmed by using the below command.

```
esxcli storage nmp satp rule list -s VMW_SATP_ALUA
```

Name	Device	Vendor	Model	Driver	Transport	Options
VMW_SATP_ALUA		NETAPP				reset_on_attempted_reserve

  

Rule Group	Claim Options	Default PSP	PSP Options	Description
system	tpgs_on	VMW_PSP_RR		NetApp arrays with ALUA support

## Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

### Example

The following example displays the correct output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

```
esxcli storage nmp device list -d naa.600a098038313530772b4d673979372f
```

```
naa.600a098038313530772b4d673979372f
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038313530772b4d673979372f)
  Storage Array Type: VMW_SATP_ALUA
  Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off;
{TPG_id=1000,TPG_state=AO}{TPG_id=1001,TPG_state=ANO}}
  Path Selection Policy: VMW_PSP_RR
  Path Selection Policy Device Config:
{policy=rr,iops=1,bytes=10485760,useANO=0; lastPathIndex=1:
NumIOsPending=0,numBytesPending=0}
  Path Selection Policy Device Custom Config:
Working Paths: vmhba3:C0:T3:L21, vmhba4:C0:T2:L21
Is USB: false
```

**esxcli storage nmp path list -d naa.600a098038313530772b4d673979372f**

fc.20000090fae0ec8e:10000090fae0ec8e-fc.201000a098dfe3d1:200b00a098dfe3d1-naa.600a098038313530772b4d673979372f

Runtime Name: vmhba3:C0:T2:L21

Device: naa.600a098038313530772b4d673979372f

Device Display Name: NETAPP Fibre Channel Disk  
(naa.600a098038313530772b4d673979372f)

Group State: active unoptimized

Array Priority: 0

Storage Array Type Path Config:

{TPG\_id=1001,TPG\_state=ANO,RTP\_id=29,RTP\_health=UP}

Path Selection Policy Path Config: PSP VMW\_PSP\_RR does not support path configuration.

fc.20000090fae0ec8e:10000090fae0ec8e-fc.201000a098dfe3d1:200700a098dfe3d1-naa.600a098038313530772b4d673979372f

Runtime Name: vmhba3:C0:T3:L21

Device: naa.600a098038313530772b4d673979372f

Device Display Name: NETAPP Fibre Channel Disk  
(naa.600a098038313530772b4d673979372f)

Group State: active

Array Priority: 0

Storage Array Type Path Config:

{TPG\_id=1000,TPG\_state=AO,RTP\_id=25,RTP\_health=UP}

Path Selection Policy Path Config: PSP VMW\_PSP\_RR does not support path configuration.

fc.20000090fae0ec8f:10000090fae0ec8f-fc.201000a098dfe3d1:200800a098dfe3d1-naa.600a098038313530772b4d673979372f

Runtime Name: vmhba4:C0:T2:L21

Device: naa.600a098038313530772b4d673979372f

Device Display Name: NETAPP Fibre Channel Disk  
(naa.600a098038313530772b4d673979372f)

Group State: active

Array Priority: 0

Storage Array Type Path Config:

{TPG\_id=1000,TPG\_state=AO,RTP\_id=26,RTP\_health=UP}

Path Selection Policy Path Config: PSP VMW\_PSP\_RR does not support path configuration.

fc.20000090fae0ec8f:10000090fae0ec8f-fc.201000a098dfe3d1:200c00a098dfe3d1-naa.600a098038313530772b4d673979372f

Runtime Name: vmhba4:C0:T3:L21

Device: naa.600a098038313530772b4d673979372f

Device Display Name: NETAPP Fibre Channel Disk

```
(naa.600a098038313530772b4d673979372f)
  Group State: active unoptimized
  Array Priority: 0
  Storage Array Type Path Config:
{TPG_id=1001,TPG_state=ANO,RTP_id=30,RTP_health=UP}
  Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.
```

## All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

### Example

The following example displays the correct output for an ONTAP LUN:

```
esxcli storage nmp device list -d naa.600a098038304759563f4e7837574453
```

```
naa.600a098038304759563f4e7837574453
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
  Storage Array Type: VMW_SATP_ALUA
  Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off;
{TPG_id=1001,TPG_state=AO}{TPG_id=1000,TPG_state=AO}}
  Path Selection Policy: VMW_PSP_RR
  Path Selection Policy Device Config:
{policy=rr,iops=1,bytes=10485760,useANO=0; lastPathIndex=2:
NumIOsPending=0,numBytesPending=0}
  Path Selection Policy Device Custom Config:
  Working Paths: vmhba4:C0:T0:L9, vmhba3:C0:T1:L9, vmhba3:C0:T0:L9,
vmhba4:C0:T1:L9
  Is USB: false
```

```
esxcli storage nmp device list -d naa.600a098038304759563f4e7837574453
```

```
fc.20000024ff171d37:21000024ff171d37-fc.202300a098ea5e27:204a00a098ea5e27-
naa.600a098038304759563f4e7837574453
  Runtime Name: vmhba4:C0:T0:L9
  Device: naa.600a098038304759563f4e7837574453
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
  Group State: active
  Array Priority: 0
```

```

Storage Array Type Path Config:
{TPG_id=1000,TPG_state=AO,RTP_id=6,RTP_health=UP}
Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff171d36:21000024ff171d36-fc.202300a098ea5e27:201d00a098ea5e27-
naa.600a098038304759563f4e7837574453
Runtime Name: vmhba3:C0:T1:L9
Device: naa.600a098038304759563f4e7837574453
Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
Group State: active
Array Priority: 0
Storage Array Type Path Config:
{TPG_id=1001,TPG_state=AO,RTP_id=3,RTP_health=UP}
Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff171d36:21000024ff171d36-fc.202300a098ea5e27:201b00a098ea5e27-
naa.600a098038304759563f4e7837574453
Runtime Name: vmhba3:C0:T0:L9
Device: naa.600a098038304759563f4e7837574453
Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
Group State: active
Array Priority: 0
Storage Array Type Path Config:
{TPG_id=1000,TPG_state=AO,RTP_id=1,RTP_health=UP}
Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff171d37:21000024ff171d37-fc.202300a098ea5e27:201e00a098ea5e27-
naa.600a098038304759563f4e7837574453
Runtime Name: vmhba4:C0:T1:L9
Device: naa.600a098038304759563f4e7837574453
Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
Group State: active
Array Priority: 0
Storage Array Type Path Config:
{TPG_id=1001,TPG_state=AO,RTP_id=4,RTP_health=UP}
Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

```



## vVol

Virtual Volumes (vVols) are a VMware object type that corresponds to a Virtual Machine (VM) disk, and its snapshots and fast-clones.

ONTAP tools for VMware vSphere includes the VASA Provider for ONTAP, which provides the integration point for a VMware vCenter to leverage vVols based storage. When you deploy the ONTAP tools OVA, it is automatically registered with the vCenter server and enables the VASA Provider.

When you create a vVols datastore using the vCenter user interface, it guides you to create FlexVols as backup storage for the datastore. vVols within a vVols datastores are access by ESXi hosts using a protocol endpoint (PE). In SAN environments, one 4MB LUN is created on each FlexVol in the datastore for use as a PE. A SAN PE is an administrative logical unit (ALU). vVols are subsidiary logical units (SLUs).

Standard requirements and best practices for SAN environments apply when using vVols, including (but not limited to) the following:

1. Create at least one SAN LIF on each node per SVM you intend to use. The best practice is to create at least two per node, but no more than necessary.
2. Eliminate any single point of failure. use multiple VMkernel network interfaces on different network subnets that use NIC teaming when multiple virtual switches are used. Or use multiple physical NICs connected to multiple physical switches to provide HA and increased throughput.
3. Configure zoning and/or VLANs as required for host connectivity.
4. Ensure all required initiators are logged into the target LIFs on the desired SVM.



You must deploy ONTAP tools for VMware vSphere to enabled the VASA Provider. The VASA Provider will manage all of your igroup settings for you, so there is no need to create or manage igroups in a vVols environment.

NetApp does not recommend changing any vVols settings from the default at this time.

Refer to the [NetApp Interoperability Matrix Tool](#) for specific versions of ONTAP tools, or legacy VASA Provider for your specific versions of vSphere and ONTAP.

For detailed information on provisioning and managing vVols, please refer to ONTAP tools for VMware vSphere documentation as well as [TR-4597-VMware vSphere with ONTAP](#) and [TR-4400](#).

## Recommended Settings

### ATS Locking

ATS locking is **mandatory** for VAAI compatible storage and upgraded VMFS5, and is required for proper interoperability and optimal VMFS shared storage I/O performance with ONTAP LUNs. Refer to VMware documentation for details on enabling ATS locking.

Settings	Default	ONTAP Recommended	Description
HardwareAcceleratedLocking	1	1	Helps enable the use of Atomic Test and Set (ATS) locking

Settings	Default	ONTAP Recommended	Description
Disk IOPs	1000	1	IOPS limit: The Round Robin PSP defaults to an IOPS limit of 1000. In this default case, a new path is used after 1000 I/O operations are issued.
Disk/QFullSampleSize	0	32	The count of QUEUE FULL or BUSY conditions it takes before ESXi starts throttling.



Enable Space-alloc setting for all the LUN's mapped to VMware vSphere for UNMAP to work. For More details, refer to ONTAP Documentation.

### Guest OS timeouts

You can manually configure the virtual machines with the recommended guest OS tunings. After tuning updates, you must reboot the guest for the updates to take effect.

#### GOS timeout values:

Guest OS Type	Timeouts
Linux variants	disk timeout = 60
Windows	disk timeout = 60
Solaris	disk timeout = 60 busy retry = 300 not ready retry = 300 reset retry = 30 max.throttle = 32 min.throttle = 8

### Validating the vSphere tunable

Use the following command to verify the HardwareAcceleratedLocking setting.

```
esxcli system settings advanced list --option /VMFS3/HardwareAcceleratedLocking
```

```
Path: /VMFS3/HardwareAcceleratedLocking
Type: integer
Int Value: 1
Default Int Value: 1
Min Value: 0
Max Value: 1
String Value:
Default String Value:
Valid Characters:
Description: Enable hardware accelerated VMFS locking (requires
compliant hardware). Please see http://kb.vmware.com/kb/2094604 before
disabling this option.
```

### Validating the Disk IOPs setting

Use the following command to verify the IOPs setting.

```
esxcli storage nmp device list -d naa.600a098038304731783f506670553355
```

```
naa.600a098038304731783f506670553355
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304731783f506670553355)
  Storage Array Type: VMW_SATP_ALUA
  Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off;
{TPG_id=1000,TPG_state=ANO}{TPG_id=1001,TPG_state=AO}}
  Path Selection Policy: VMW_PSP_RR
  Path Selection Policy Device Config:
{policy=rr,iops=1,bytes=10485760,useANO=0; lastPathIndex=0:
NumIOsPending=0,numBytesPending=0}
  Path Selection Policy Device Custom Config:
  Working Paths: vmhba4:C0:T0:L82, vmhba3:C0:T0:L82
  Is USB: false
```

### Validating the QFullSampleSize

Use the following command to verify the QFullSampleSize

```
esxcli system settings advanced list --option /Disk/QFullSampleSize
```

```
Path: /Disk/QFullSampleSize
Type: integer
Int Value: 32
Default Int Value: 0
Min Value: 0
Max Value: 64
String Value:
Default String Value:
Valid Characters:
Description: Default I/O samples to monitor for detecting non-transient
queue full condition. Should be nonzero to enable queue depth throttling.
Device specific QFull options will take precedence over this value if set.
```

## Known issues

There are no known issues for the VMware vSphere 7.x with ONTAP release.

### Related information

- [TR-4597-VMware vSphere with ONTAP](#)
- [VMware vSphere 5.x, 6.x and 7.x support with NetApp MetroCluster \(2031038\)](#)
- [NetApp ONTAP with NetApp SnapMirror Business Continuity \(SM-BC\) with VMware vSphere Metro Storage Cluster \(vMSC\)](#)

## Use VMware vSphere 6.5 and 6.7 with ONTAP

You can use ONTAP SAN host configuration settings for the vSphere 6.5.x and 6.7.x releases with FC, FCoE and iSCSI protocols.

## Hypervisor SAN Booting

### What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the [NetApp Interoperability Matrix Tool](#) to verify that your OS, HBA, HBA firmware and the HBA boot BIOS, and ONTAP version are supported.

### Steps

1. Map the SAN boot LUN to the host.
2. Verify that multiple paths are available.



Multiple paths become available after the host OS is up and running on the paths.

3. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

For information on how to enable the HBA BIOS, see your vendor-specific documentation.

4. Reboot the host to verify that the boot is successful.

## Multipathing

ESXi provides an extensible multipathing module called Native Multipathing Plug-In (NMP) that manages the sub-plugins Storage Array Type Plugins (SATPs), and Path Selection Plugins (PSPs). These SATP rules are available by default in ESXi.

For NetApp ONTAP storage, VMW\_SATP\_ALUA plugin is used by default with VMW\_PSP\_RR as a path selection policy (PSP). This can be confirmed by using the below command:

```
esxcli storage nmp satp rule list -s VMW_SATP_ALUA
```

Name	Device	Vendor	Model	Driver	Transport	Options
-----	-----	-----	-----	-----	-----	-----
VMW_SATP_ALUA		LSI	INF-01-00			
reset_on_attempted_reserve						
VMW_SATP_ALUA		NETAPP				
reset_on_attempted_reserve						

  

Rule Group	Claim Options	Default PSP	PSP Options	Description
-----	-----	-----	-----	-----
system	tpgs_on	VMW_PSP_MRU		NetApp E-Series arrays
with ALUA support				
system	tpgs_on	VMW_PSP_RR		NetApp arrays with ALUA
support				

## All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

### Example

The following example displays the correct output for an ONTAP LUN:

```
esxcli storage nmp device list -d naa.600a098038304759563f4e7837574453
```

```
fc.20000024ff171d37:21000024ff171d37-fc.202300a098ea5e27:204a00a098ea5e27-
naa.600a098038304759563f4e7837574453
  Runtime Name: vmhba4:C0:T0:L9
  Device: naa.600a098038304759563f4e7837574453
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
  Group State: active
  Array Priority: 0
  Storage Array Type Path Config:
{TPG_id=1000,TPG_state=AO,RTP_id=6,RTP_health=UP}
  Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.
```

```

fc.20000024ff171d36:21000024ff171d36-fc.202300a098ea5e27:201d00a098ea5e27-
naa.600a098038304759563f4e7837574453
  Runtime Name: vmhba3:C0:T1:L9
  Device: naa.600a098038304759563f4e7837574453
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
  Group State: active
  Array Priority: 0
  Storage Array Type Path Config:
{TPG_id=1001,TPG_state=AO,RTP_id=3,RTP_health=UP}
  Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff171d36:21000024ff171d36-fc.202300a098ea5e27:201b00a098ea5e27-
naa.600a098038304759563f4e7837574453
  Runtime Name: vmhba3:C0:T0:L9
  Device: naa.600a098038304759563f4e7837574453
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
  Group State: active
  Array Priority: 0
  Storage Array Type Path Config:
{TPG_id=1000,TPG_state=AO,RTP_id=1,RTP_health=UP}
  Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff171d37:21000024ff171d37-fc.202300a098ea5e27:201e00a098ea5e27-
naa.600a098038304759563f4e7837574453
  Runtime Name: vmhba4:C0:T1:L9
  Device: naa.600a098038304759563f4e7837574453
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
  Group State: active
  Array Priority: 0
  Storage Array Type Path Config:
{TPG_id=1001,TPG_state=AO,RTP_id=4,RTP_health=UP}
  Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

```

In the above example, LUN has been mapped from NetApp storage with 4 paths (4 active-optimized).

## Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a

different controller. The non-optimized paths are only used when no optimized paths are available.

### Example

The following example displays the correct output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

```
esxcli storage nmp path list -d naa.600a098038313530772b4d673979372f
```

```
fc.20000090fae0ec8e:10000090fae0ec8e-fc.201000a098dfe3d1:200b00a098dfe3d1-
naa.600a098038313530772b4d673979372f
  Runtime Name: vmhba3:C0:T2:L21
  Device: naa.600a098038313530772b4d673979372f
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038313530772b4d673979372f)
  Group State: active unoptimized
  Array Priority: 0
  Storage Array Type Path Config:
{TPG_id=1001,TPG_state=ANO,RTP_id=29,RTP_health=UP}
  Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000090fae0ec8e:10000090fae0ec8e-fc.201000a098dfe3d1:200700a098dfe3d1-
naa.600a098038313530772b4d673979372f
  Runtime Name: vmhba3:C0:T3:L21
  Device: naa.600a098038313530772b4d673979372f
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038313530772b4d673979372f)
  Group State: active
  Array Priority: 0
  Storage Array Type Path Config:
{TPG_id=1000,TPG_state=AO,RTP_id=25,RTP_health=UP}
  Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000090fae0ec8f:10000090fae0ec8f-fc.201000a098dfe3d1:200800a098dfe3d1-
naa.600a098038313530772b4d673979372f
  Runtime Name: vmhba4:C0:T2:L21
  Device: naa.600a098038313530772b4d673979372f
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038313530772b4d673979372f)
  Group State: active
  Array Priority: 0
  Storage Array Type Path Config:
{TPG_id=1000,TPG_state=AO,RTP_id=26,RTP_health=UP}
  Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.
```

```
fc.20000090fae0ec8f:10000090fae0ec8f-fc.201000a098dfe3d1:200c00a098dfe3d1-  
naa.600a098038313530772b4d673979372f  
  Runtime Name: vmhba4:C0:T3:L21  
  Device: naa.600a098038313530772b4d673979372f  
  Device Display Name: NETAPP Fibre Channel Disk  
(naa.600a098038313530772b4d673979372f)  
  Group State: active unoptimized  
  Array Priority: 0  
  Storage Array Type Path Config:  
{TPG_id=1001,TPG_state=ANO,RTP_id=30,RTP_health=UP}  
  Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path  
configuration.
```

In the above example, LUN has been mapped from NetApp storage with 4 paths (2 active-optimized and 2 active-unoptimized).

## vVol

Virtual Volumes (vVols) are a VMware object type that corresponds to a Virtual Machine (VM) disk, and its snapshots and fast clones.

ONTAP tools for VMware vSphere includes the VASA Provider for ONTAP, which provides the integration point for a VMware vCenter to leverage vVols based storage. When you deploy the ONTAP tools OVA, it is automatically registered with the vCenter server and enables the VASA Provider.

When you create a vVols datastore using the vCenter user interface, it guides you to create FlexVols as backup storage for the datastore. vVols within a vVols datastores are accessed by ESXi hosts using a protocol endpoint (PE). In SAN environments, one 4MB LUN is created on each FlexVol in the datastore for use as a PE. A SAN PE is an administrative logical unit (ALU); vVols are subsidiary logical units (SLUs).

Standard requirements and best practices for SAN environments apply when using vVols, including (but not limited to) the following:

1. Create at least one SAN LIF on each node per SVM you intend to use. The best practice is to create at least two per node, but no more than necessary.
2. Eliminate any single point of failure. Use multiple VMkernel network interfaces on different network subnets that use NIC teaming when multiple virtual switches are used or use multiple physical NICs connected to multiple physical switches to provide HA and increased throughput.
3. Configure zoning and/or VLANs as required for host connectivity.
4. Ensure all required initiators are logged into the target LIFs on the desired SVM.



You must deploy ONTAP tools for VMware vSphere to enable the VASA Provider. The VASA Provider will manage all of your igroup settings for you, so there is no need to create or manage iGroups in a vVols environment.

NetApp does not recommend changing any vVols settings from default at this time.

Refer to the [NetApp Interoperability Matrix Tool](#) for specific versions of ONTAP tools, or legacy VASA Provider for your specific versions of vSphere and ONTAP.



For detailed information on provisioning and managing vVols, please refer to ONTAP tools for VMware vSphere documentation as well as [TR-4597](#) and [TR-4400](#).

## Recommended Settings

### ATS Locking

ATS locking is **mandatory** for VAAI compatible storage and upgraded VMFS5 and is required for proper interoperability and optimal VMFS shared storage I/O performance with ONTAP LUNs. Refer to VMware documentation for details on enabling ATS locking.

Settings	Default	ONTAP Recommended	Description
HardwareAcceleratedLocking	1	1	Helps enable the use of Atomic Test and Set (ATS) locking
Disk IOPs	1000	1	IOPS limit: The Round Robin PSP defaults to an IOPS limit of 1000. In this default case, a new path is used after 1000 I/O operations are issued.
Disk/QFullSampleSize	0	32	The count of QUEUE FULL or BUSY conditions it takes before ESXi starts throttling.



Enable Space-alloc setting for all the LUN's mapped to VMware vSphere for UNMAP to work. For more details, refer to [ONTAP Documentation](#).

### Guest OS timeouts

You can manually configure the virtual machines with the recommended guest OS tunings. After tuning updates, you must reboot the guest for the updates to take effect.

#### GOS timeout values:

Guest OS Type	Timeouts
Linux variants	disk timeout = 60
Windows	disk timeout = 60
Solaris	disk timeout = 60 busy retry = 300 not ready retry = 300 reset retry = 30 max.throttle = 32 min.throttle = 8

### Validating the vSphere tunable

Use the following command to verify the `HardwareAcceleratedLocking` setting:

```
esxcli system settings advanced list --option /VMFS3/HardwareAcceleratedLocking
```

```
Path: /VMFS3/HardwareAcceleratedLocking
Type: integer
Int Value: 1
Default Int Value: 1
Min Value: 0
Max Value: 1
String Value:
Default String Value:
Valid Characters:
Description: Enable hardware accelerated VMFS locking (requires
compliant hardware). Please see http://kb.vmware.com/kb/2094604 before
disabling this option.
```

### Validating the Disk IOPs setting

Use the following command to verify the IOPs setting:

```
esxcli storage nmp device list -d naa.600a098038304731783f506670553355
```

```
naa.600a098038304731783f506670553355
  Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304731783f506670553355)
  Storage Array Type: VMW_SATP_ALUA
  Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off;
{TPG_id=1000,TPG_state=ANO}{TPG_id=1001,TPG_state=AO}}
  Path Selection Policy: VMW_PSP_RR
  Path Selection Policy Device Config:
{policy=rr,iops=1,bytes=10485760,useANO=0; lastPathIndex=0:
NumIOsPending=0,numBytesPending=0}
  Path Selection Policy Device Custom Config:
Working Paths: vmhba4:C0:T0:L82, vmhba3:C0:T0:L82
Is USB: false
```

### Validating the QFullSampleSize

Use the following command to verify the QFullSampleSize:

```
esxcli system settings advanced list --option /Disk/QFullSampleSize
```

Path: /Disk/QFullSampleSize

Type: integer

Int Value: 32

Default Int Value: 0

Min Value: 0

Max Value: 64

String Value:

Default String Value:

Valid Characters:

Description: Default I/O samples to monitor for detecting non-transient queue full condition. Should be nonzero to enable queue depth throttling. Device specific QFull options will take precedence over this value if set.

## Known issues

The VMware vSphere 6.5 and 6.7 with ONTAP release has the following known issues:

OS version	NetApp Bug ID	Title	Description
ESXi 6.5 and ESXi 6.7.x	1413424	WFC RDM luns fails during testing	Windows failover clustering raw device mapping between Windows Virtual Machines like Windows 2019, Windows 2016, and Windows 2012 across VMWare ESXi host failed during storage failover testing on all the 7-mode, C-mode cluster controllers.
ESXi 6.5.x and ESXi 6.7.x	1256473	PLOGI issue seen during testing on Emulex adapters	

## Related information

- [TR-4597-VMware vSphere with ONTAP](#)
- [VMware vSphere 5.x, 6.x and 7.x support with NetApp MetroCluster \(2031038\)](#)
- [NetApp ONTAP with NetApp SnapMirror Business Continuity \(SM-BC\) with VMware vSphere Metro Storage Cluster \(vMSC\)](#)

## Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

## Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.