ONTAP Select documentation

ONTAP Select

NetApp
November 08, 2026

This PDF was generated from https://docs.netapp.com/us-en/ontap-select-9191/index.html on November 08, 2026. Always check docs.netapp.com for the latest.

Table of Contents

ONTAP Select documentation . . . 1
Release Notes . . . 2
ONTAP Select Release Notes . . . 2
What's new for ONTAP Select . . . 2
ONTAP Select 9.19.1 . . . 2
ONTAP Select 9.18.1 . . . 2
ONTAP Select 9.17.1 . . . 3
ONTAP Select 9.16.1 . . . 5
ONTAP Select 9.15.1 . . . 5
ONTAP Select 9.14.1 . . . 6
ONTAP Select 9.13.1 . . . 6
ONTAP Select 9.12.1 . . . 6
ONTAP Select 9.11.1 . . . 7
ONTAP Select 9.10.1 . . . 7
ONTAP Select 9.9.1 . . . 8
ONTAP Select 9.8 . . . 8
Concepts . . . 9
Learn about ONTAP Select . . . 9
Software-defined storage . . . 9
Two software components . . . 9
Illustration of a typical deployment . . . 10
FlexCache write-back with ONTAP Select . . . 10
Compare ONTAP Select and ONTAP 9 . . . 11
Learn about ONTAP Select Deploy . . . 12
Core functionality . . . 12
Ways you can access the Deploy utility . . . 12
Business use cases . . . 13
ONTAP Select business needs and usage scenarios . . . 13
ONTAP Select use cases for remote and branch offices . . . 14
ONTAP Select support for private cloud and data centers . . . 15
ONTAP Select data protection and efficiency . . . 16
ONTAP Select terminology and key concepts . . . 18
Plan . . . 23
ONTAP Select installation and deployment workflow . . . 23
ONTAP Select . . . 23
ONTAP Select requirements and planning considerations . . . 23
ONTAP Select VMware hypervisor and hardware considerations . . . 27
ONTAP Select storage and RAID considerations . . . 29
External storage requirements . . . 33
ONTAP Select networking requirements . . . 35
ONTAP Select two-node clusters with HA . . . 38
ONTAP Select remote and branch office deployments . . . 39
ONTAP Select MetroCluster SDS deployment requirements . . . 39

ONTAP Select VMware vCenter server on ESXi 40
ONTAP Select Deploy 42
ONTAP Select Deploy general requirements and planning 42
ONTAP Select Deploy hypervisor host considerations 44
Summary of best practices for ONTAP Select deployment 47
Storage 47
Networking 48
HA 49
License 51
Options 51
Learn about evaluation licenses for ONTAP Select deployments 51
Production licenses 52
Purchase a license 60
License purchasing workflow for ONTAP Select 60
Capacity Tiers 62
Capacity Pools 64
ONTAP Select support for ONTAP features 65
ONTAP features licensed by default 65
ONTAP features that are separately licensed 66
Install 68
Pre-installation checklist 68
Host preparation checklist 68
Required information for ONTAP Select Deploy utility installation 77
Required information for ONTAP Select installation 78
Configure an ONTAP Select host to use NVMe drives 79
Install ONTAP Select Deploy 84
Step 1: Download the virtual machine image 84
Step 2: Verify the ONTAP Select Deploy OVA signature 85
Step 3: Deploy the virtual machine 86
Step 4: Sign in to the Deploy web interface 88
Deploy an ONTAP Select cluster using the web UI 88
Step 1: Prepare for the deployment 89
Step 2: Create a single-node or multi-node cluster 90
Step 3: Complete the deployment 92
Initial state of the ONTAP Select cluster after deployment 92
Administer 94
Learn about administering ONTAP Select 94
Administer ONTAP Select 94
Perform additional ONTAP configuration 94
Upgrade ONTAP Select nodes 95
General procedure 95
Revert an ONTAP Select node 96
Use the VMXNET3 network driver 96
Manage ONTAP Select diagnostics and support 96
Configure the Deploy system 96

Display the ONTAP Select Deploy event messages . . . 96
Enable AutoSupport . . . 97
Generate and download an AutoSupport package . . . 97
Secure an ONTAP Select deployment . . . 98
Change the Deploy administrator password . . . 98
Add a management server account . . . 98
Configure MFA . . . 99
ONTAP Select Deploy CLI MFA login using YubiKey PIV or FIDO2 authentication . . . 99
Configure the public key in ONTAP Select Deploy . . . 100
Log in to ONTAP Select Deploy using YubiKey PIV Authentication over SSH . . . 100
ONTAP Select Deploy CLI MFA login using ssh-keygen . . . 101
Confirm connectivity among the ONTAP Select nodes . . . 103
Administer the ONTAP Select Deploy mediator services . . . 104
View the status of the mediator service . . . 104
Clusters . . . 104
Manage ONTAP Select clusters . . . 104
Expand or contract an ONTAP Select cluster on ESXi or KVM hosts . . . 106
Nodes and hosts . . . 109
Access an ONTAP Select video console . . . 109
Resize ONTAP Select cluster nodes . . . 109
Replace failed software RAID drives for ONTAP Select . . . 110
Upgrade ONTAP Select node to VMFS6 using Storage vMotion . . . 117
Manage ONTAP Select licenses . . . 119
Manage the Capacity Tier licenses . . . 120
Manage the Capacity Pool licenses . . . 120
Reinstall a Capacity Pool license . . . 121
Convert an evaluation license to a production license . . . 122
Manage an expired Capacity Pool license . . . 122
Manage add-on licenses . . . 123
Deep dive . . . 124
Storage . . . 124
Learn about the general storage characteristics for ONTAP Select . . . 124
Hardware RAID services for ONTAP Select with local attached storage . . . 129
ONTAP Select software RAID configuration services for local attached storage . . . 136
ONTAP Select vSAN and external array configurations . . . 144
Considerations for increasing ONTAP Select storage capacity . . . 147
ONTAP Select storage efficiency support . . . 150
Networking . . . 153
ONTAP Select networking concepts and characteristics . . . 153
ONTAP Select single and multiple node network configurations . . . 155
ONTAP Select internal and external networks . . . 160
Supported ONTAP Select network configurations . . . 162
ONTAP Select VMware vSphere vSwitch configuration on ESXi . . . 163
ONTAP Select physical switch configuration . . . 173
ONTAP Select data and management traffic separation . . . 175

High availability architecture . . . 177
ONTAP Select high availability configurations . . . 177
ONTAP Select HA RSM and mirrored aggregates . . . 180
Data protection enhancements with ONTAP Select HA . . . 183
Learn about ONTAP Select performance . . . 185
Automate with REST . . . 187
Concepts . . . 187
Learn about REST web services for deploying and managing ONTAP Select clusters . . . 187
Learn about accessing the ONTAP Select Deploy API . . . 188
ONTAP Select Deploy API basic operational characteristics . . . 188
Request and response API transaction for ONTAP Select . . . 190
Asynchronous processing using the Job object for ONTAP Select . . . 193
Access with a browser . . . 194
Learn about accessing the ONTAP Select Deploy API with a browser . . . 194
Access the ONTAP Select Deploy documentation page . . . 195
Understand and execute an ONTAP Select Deploy API call . . . 195
Workflow processes . . . 196
Learn about using the ONTAP Select Deploy API workflows . . . 196
Create an ONTAP Select single-node evaluation cluster on ESXi . . . 196
Access with Python . . . 203
Learn about accessing the ONTAP Select Deploy API using Python . . . 203
Learn about the Python scripts for ONTAP Select Deploy . . . 203
Python code samples . . . 205
Script to create an ONTAP Select cluster . . . 205
JSON for script to create an ONTAP Select cluster . . . 212
Script to add an ONTAP Select node license . . . 217
Script to delete an ONTAP Select cluster . . . 220
Common support Python module for ONTAP Select . . . 222
Script to resize ONTAP Select cluster nodes . . . 226
Use the CLI . . . 231
Sign in to ONTAP Select Deploy using SSH . . . 231
Deploy an ONTAP Select cluster using the CLI . . . 231
Step 1: Prepare for the deployment . . . 231
Step 2: Upload and register a license file . . . 232
Step 3: Add hypervisor hosts . . . 232
Step 4: Create and configure an ONTAP Select cluster . . . 234
Step 5: Configure an ONTAP Select node . . . 235
Step 6: Attach storage to the ONTAP Select nodes . . . 237
Step 7: Deploy an ONTAP Select cluster . . . 239
Secure an ONTAP Select deployment . . . 239
Confirm network connectivity among ONTAP Select nodes . . . 240
Manage ONTAP Select clusters using the CLI . . . 241
Back up the ONTAP Select Deploy configuration data . . . 241
Delete an ONTAP Select cluster . . . 241
Nodes and hosts . . . 242

Upgrade to VMware ESXi 8.0 or later for ONTAP Select . . . 242
Modify a host management server for ONTAP Select Deploy . . . 246
Deploy utility . . . 247
Upgrade an ONTAP Select Deploy instance . . . 247
Migrate an ONTAP Select Deploy instance to a new virtual machine . . . 249
Add an ONTAP Select image to Deploy . . . 251
Remove an ONTAP Select image from Deploy . . . 252
Recover the ONTAP Select Deploy utility for a two-node cluster . . . 253
Try ONTAP Select . . . 258
Access ONTAP Select evaluation software . . . 258
Step 1: Register for an account . . . 258
Step 2: Download ONTAP Select evaluation software . . . 259
Deploy a 90-day evaluation instance of an ONTAP Select cluster . . . 259
Step 1: Prepare the ONTAP Select cluster host . . . 260
Step 2: Deploy a single-node ONTAP Select cluster using an OVF template . . . 260
FAQ for ONTAP Select . . . 262
General . . . 262
Licenses, installation, upgrades, and reverts . . . 263
Storage . . . 264
vCenter . . . 267
HA and clusters . . . 268
Mediator service . . . 270
Legal notices for ONTAP Select . . . 271
Copyright . . . 271
Trademarks . . . 271
Patents . . . 271
Privacy policy . . . 271
Open source . . . 271

ONTAP Select documentation

Release Notes

ONTAP Select Release Notes

The Release Notes for ONTAP Select provide release-specific information, including new features, supported configurations, upgrade notes, known issues, fixed issues, and known limitations.

You need an account to sign in to the NetApp Support Site to access the Release Notes.

Current version of ONTAP Select

You can access the ONTAP Select 9.19.1 Release Notes to view the details about the current version.

What's new for ONTAP Select

Learn about the new features and enhancements in supported versions of ONTAP Select.

ONTAP Select 9.19.1

Support for external key management (KMIP)

Beginning with ONTAP Select 9.19.1, ONTAP Select supports external key management using the Key Management Interoperability Protocol (KMIP). You can enable external key management after deploying an ONTAP Select 9.19.1 cluster or upgrading an existing cluster to ONTAP Select 9.19.1. For more information, refer to ONTAP features licensed by default.

ONTAP Select does not support cloud key managers.

Support for Advanced Micro Devices (AMD) processors

Beginning with ONTAP Select 9.19.1, ONTAP Select supports AMD processors on Kernel-Based Virtual Machine (KVM) and ESXi hypervisor hosts. For more information, refer to VMware hypervisor and hardware considerations.

ONTAP Select 9.18.1

Updates to the ONTAP Select product download options

The ONTAP Select product download options on the NetApp Support Site are changing with the ONTAP Select 9.18.1 release to **ONTAP Select Deploy** and **ONTAP Select Image**. These two new product download options replace the previous four options, as shown in the following table:

Previous product download option	New product download option	For more information, go to….
ONTAP Select	ONTAP Select Deploy	Install ONTAP Select Deploy
ONTAP Select Image Install	ONTAP Select Image	Add an ONTAP Select image to Deploy
ONTAP Select Deploy Upgrade	ONTAP Select Deploy	Upgrade an ONTAP Select Deploy instance
ONTAP Select Node Upgrade	ONTAP Select Image	Upgrade ONTAP Select nodes

Enhanced support for cluster expansion and contraction

Beginning with ONTAP Select 9.18.1, cluster expansion and contraction is supported between four-node and twelve-node ONTAP Select clusters on ESXi and KVM hypervisor hosts.

You can increase the cluster size for an existing cluster on an ESXi or KVM host in the following increments:

- From four nodes to six, eight, ten, or twelve nodes
- From six nodes to eight, ten, or twelve nodes
- From eight nodes to ten or twelve nodes
- From ten to twelve nodes

You can decrease the cluster size for an existing cluster on an ESXi or KVM host in the following increments:

- From twelve nodes to ten, eight, six, or four nodes
- From ten nodes to eight, six, or four nodes
- From eight to six or four nodes
- From six to four nodes

Learn about cluster expansion and contraction and the supported KVM and ESXi hypervisor versions.

Updated KVM hypervisor support

Beginning with ONTAP Select 9.18.1, KVM hypervisor is supported on Rocky Linux 10.1.

ONTAP Select 9.17.1

Software RAID support for locally attached NVMe disks on KVM hosts

ONTAP Select 9.17.1 introduces software RAID support for ONTAP Select locally attached NVMe disks on a KVM hypervisor host.

To use software RAID for locally attached NVMe disks, you need to configure PCI passthrough (DirectPath IO) when you are preparing the KVM host. This provides direct access for the KVM host to the locally attached NVMe disks, which is required for the following tasks:

- Configuring the KVM host to use NVMe drives
- Using software RAID after you deploy the cluster

Learn about the software RAID configuration services for local attached storage.

Support for cluster expansion and contraction on KVM hosts

Beginning with ONTAP Select 9.17.1, the cluster expansion and contraction feature is supported for both Kernel-Based Virtual Machine (KVM) and ESXi hypervisor hosts. For ONTAP Select 9.16.1 and 9.15.1, cluster expansion and contraction is only supported for ESXi hypervisor hosts.

You can use the cluster expansion and contraction feature on an existing ONTAP Select cluster to increase the size from a six-node to an eight-node cluster or decrease the size from an eight-node to a six-node cluster.

Learn about cluster expansion and contraction and the supported KVM and ESXi hypervisor versions.

Enhanced support for cluster expansion and contraction on ESXi hosts

Beginning with ONTAP Select 9.17.1, cluster expansion and contraction is supported between six-node and twelve-node clusters on ESXi hosts:

You can increase the cluster size for an existing ESXi cluster in the following increments:

- From six nodes to eight, ten, or twelve nodes
- From eight nodes to ten or twelve nodes
- From ten to twelve nodes

You can decrease the cluster size for an existing ESXi cluster in the following increments:

- From twelve nodes to ten, eight, or six nodes
- From ten nodes to eight or six nodes
- From eight to six nodes

Learn about cluster expansion and contraction and the supported KVM and ESXi hypervisor versions.

Support for SnapMirror cloud

Beginning with ONTAP Select 9.17.1, SnapMirror cloud is supported with ONTAP Select. SnapMirror cloud is not licensed by default for ONTAP Select. When you deploy an ONTAP Select 9.17.1 cluster or upgrade an existing ONTAP Select cluster to version 9.17.1, you need to request a SnapMirror cloud license. For more information, refer to ONTAP features that are separately licensed.

Support for SnapLock Select

Beginning with ONTAP Select 9.17.1, new deployments have SnapLock Select (including tamperproof snapshot locking) licensed automatically. For upgrades from ONTAP Select 9.16.1 and earlier, you can download the SnapLock Select license free of charge from the NetApp Support Site and apply it manually. For more information, refer to ONTAP features licensed by default.

Support for vSAN ESA on multi-node clusters

Beginning with ONTAP Select 9.17.1, support for vSAN Express Storage Architecture (ESA) is extended to deploy multi-node clusters. This enhancement enables you to deploy vSAN ESA configurations.

vSAN ESA is an ESXi specific configuration and is a new architecture for VMware vSAN introduced in vSphere 8. vSAN ESA is designed to offer enhanced efficiency, scalability, and performance, particularly when using NVMe-based TLC flash devices.

ONTAP Select Deploy and ONTAP Select for ESXi support the configuration of an ONTAP Select single-node and multi-node cluster, using either a vSAN or an external array type of datastore for its storage pool.



To perform Content Library operations, you need access to additional permission levels. This access is required for vSAN ESA deployments if you are working with restricted roles.

Learn about accessing permission levels for Content Library operations.

Updated NDA driver support

Beginning with ONTAP Select 9.17.1, the FreeBSD NDA (NVMe Direct Access) device driver replaces the NVD (NVMe express disk driver). The FreeBSD driver provides support for direct access devices by implementing the NVMe command protocol. When you deploy an ONTAP Select 9.17.1 cluster or upgrade an existing ONTAP Select cluster to version 9.17.1, the FreeBSD driver is started by default when you configure the NVMe devices.

For ONTAP Select 9.16.1 and earlier, the NVD continues to provide support for NVMe devices in your ONTAP Select deployment.

Updated KVM hypervisor support

Beginning with ONTAP Select 9.17.1, KVM hypervisor is supported on Red Hat Enterprise Linux (RHEL) 10.1, 10.0, 9.7, and 9.6 and Rocky Linux 10.0, 9.7, and 9.6.

There are software RAID workflow limitations for KVM hypervisors on RHEL 10.1 and 10.0 and Rocky Linux 10.0 hosts. For more information, see the following Knowledge Base articles:

- CDEPLOY-4020: ONTAP Select Deploy: Warning message while create the Cluster HWR using RHEL 10 and ROCKY 10
- CDEPLOY-4025: ONTAP Select DeployGUI: Storage pools and storage disks not visible for SWR in cluster creation page on hosts with RHEL10/Rocky 10

Enhanced VMware ESXi support

ONTAP Select 9.17.1 includes support for VMware ESXi 9.0.

ONTAP Select 9.16.1

Updated NetApp License File support

Beginning with ONTAP Select 9.16.1, the NetApp License File (NLF) support is updated. The new NLF format includes licenses for the ARP, ONTAP S3, and S3 SnapMirror features. Learn more.

For a new ONTAP Select 9.16.1 Deploy, the new NLF format is automatically applied. When you upgrade an existing ONTAP Select Deploy to 9.16.1, the new NLF format is not applied. To get the ARP, ONTAP S3, and S3 SnapMirror feature licenses, you must download the updated NLF after you upgrade. If you revert to ONTAP Select 9.15.1 or earlier, you must reinstall the feature licenses that you had before the upgrade.

Support for Autonomous Ransomware Protection

ONTAP Select 9.16.1 introduces support for Autonomous Ransomware Protection (ARP). ONTAP Select 9.16.1 only supports manual updates for ARP, automatic updates are not supported. The ARP feature license is included in the NLF for ONTAP Select 9.16.1. Learn more.

Enhanced VMware ESXi support

ONTAP Select 9.16.1 includes support for VMware ESXi 8.0 U3.

Updated KVM hypervisor support

Beginning with ONTAP Select 9.16.1, KVM hypervisor is supported on RHEL 9.5 and Rocky Linux 9.5.

ONTAP Select 9.15.1

Updated KVM hypervisor support

Beginning with ONTAP Select 9.15.1, Kernel-Based Virtual Machine (KVM) hypervisor is supported on RHEL 9.4 and Rocky Linux 9.4.

Support for cluster expansion and contraction

Beginning with ONTAP Select 9.15.1, cluster expansion and contraction is supported.

- Cluster expansion from six-node to eight-node clusters

You can increase the cluster size from a six-node cluster to an eight-node cluster with the cluster expansion feature. Cluster expansions from one, two, or four-node clusters to six or eight-node clusters are not currently supported. Learn more.

- Cluster contraction eight-node to six-node clusters

  You can decrease the cluster size from an eight-node cluster to six-node cluster with the cluster contraction feature. Cluster contractions from six or eight-node clusters to one, two, or four-node clusters are not currently supported. Learn more.

Support for cluster expansion and contraction is limited to ESXi clusters only.

ONTAP Select 9.14.1

Support for KVM hypervisor

Beginning with ONTAP Select 9.14.1, support for KVM hypervisor has been reinstated. Previously, support for deploying a new cluster on a KVM hypervisor was removed in ONTAP Select 9.10.1 and support for managing existing KVM clusters and hosts, except to take offline or delete, was removed in ONTAP Select 9.11.1.

Deploy VMware vCenter plug-in is no longer supported

Beginning with ONTAP Select 9.14.1, the Deploy VMware vCenter plug-in is no longer supported.

Updated ONTAP Select Deploy support

If you are running a version of ONTAP Select Deploy 9.14.1 lower than 9.14.1P2, you should upgrade to ONTAP Select Deploy 9.14.1P2 as soon as possible. For more information, refer to the ONTAP Select 9.14.1 Release Notes.

Enhanced VMware ESXi support

ONTAP Select 9.14.1 includes support for VMware ESXi 8.0 U2.

ONTAP Select 9.13.1

Support for NVMe over TCP

When you upgrade to ONTAP Select 9.13.1, you must have the new license to support NVMe over TCP. This license is automatically included when you deploy ONTAP Select for the first time from version 9.13.1.

Updated VMware ESXi support

Beginning with ONTAP 9.13.1, VMware ESXi 8.0.1 GA (build 20513097) is supported with hardware version 4 and later.

Updated ONTAP Select Deploy support

As of April 2024, ONTAP Select Deploy 9.13.1 is no longer available on the NetApp Support Site. If you are running ONTAP Select Deploy 9.13.1, you should upgrade to ONTAP Select Deploy 9.14.1P2 as soon as possible. For more information, refer to the ONTAP Select 9.14.1 Release Notes.

ONTAP Select 9.12.1

ONTAP Select 9.12.1 benefits from most of the new developments in the current release of the core ONTAP product. It does not include any new features or improvements specific to ONTAP Select.

As of April 2024, ONTAP Select Deploy 9.12.1 is no longer available on the NetApp Support Site. If you are

running ONTAP Select Deploy 9.12.1, you should upgrade to ONTAP Select Deploy 9.14.1P2 as soon as possible. For more information, refer to the ONTAP Select 9.14.1 Release Notes.

ONTAP Select 9.11.1

Enhanced VMware ESXi support

ONTAP Select 9.11.1 includes support for VMware ESXi 7.0 U3C.

Support for VMware NSX-T

ONTAP Select 9.10.1 and later releases have been qualified for VMware NSX-T version 3.1.2. There are no functional issues or deficiencies when using NSX-T with an ONTAP Select single-node cluster deployed with an OVA file and the ONTAP Select Deploy administration utility. However, when using NSX-T with an ONTAP Select multi-node cluster, you should note the following limitation for ONTAP Select 9.11.1:

- Network connectivity checker

  The network connectivity checker available through the Deploy CLI fails when it is run against an NSX-T based network.

KVM hypervisor is no longer supported

- Beginning with ONTAP Select 9.10.1, you can no longer deploy a new cluster on the KVM hypervisor.
- Beginning with ONTAP Select 9.11.1, all manageability functionality is no longer available for existing KVM clusters and hosts, except for the take offline and delete functions.

  NetApp strongly recommends that customers plan and execute a full data migration from ONTAP Select for KVM to any other ONTAP platform, including ONTAP Select for ESXi. For more information, refer to the EOA Notice

ONTAP Select 9.10.1

Support for VMware NSX-T

ONTAP Select 9.10.1 has been qualified for VMware NSX-T version 3.1.2. There are no functional issues or deficiencies when using NSX-T with an ONTAP Select single-node cluster deployed with an OVA file and the ONTAP Select Deploy administration utility. However, when using NSX-T with an ONTAP Select multi-node cluster, you should note the following requirements and limitations:

- Cluster MTU

  You must manually adjust the cluster MTU size to 8800 before deploying the cluster to account for the additional overhead. The VMware guidance is to allow for a 200-byte buffer when using NSX-T.

- Network 4x10Gb configuration

  For ONTAP Select deployments on a VMware ESXi host configured with four NICs, the Deploy utility will prompt you to follow the best practice of splitting internal traffic across two different port groups and external traffic across two different port groups. However, when using an overlay network this configuration does not work and you should disregard the recommendation. In this case, you should instead use only one internal port group and one external port group.

- Network connectivity checker

  The network connectivity checker available through the Deploy CLI fails when it is run against an NSX-T based network.

KVM hypervisor is no longer supported

Beginning with ONTAP Select 9.10.1 you can no longer deploy a new cluster on the KVM hypervisor. However, if you upgrade a cluster from a previous release to 9.10.1 you can still use the Deploy utility to administer the cluster.

ONTAP Select 9.9.1

Processor family support

Beginning with ONTAP Select 9.9.1, only CPU models from Intel Xeon Sandy Bridge or later are supported for ONTAP Select.

Updated VMware ESXi support

Support for VMware ESXi has been enhanced with ONTAP Select 9.9.1. The following releases are now supported:

- ESXi 7.0 U2
- ESXi 7.0 U1

ONTAP Select 9.8

High speed interface

The high speed interface feature enhances network connectivity by providing an option for both 25G (25GbE) and 40G (40GbE). To achieve the best performance when using these higher speeds, you should follow the best practices regarding port mapping configurations as described in the ONTAP Select documentation.

Updated VMware ESXi support

There are two changes for ONTAP Select 9.8 regarding the support for VMware ESXi.

- ESXi 7.0 is supported (GA build 15843807 and later)
- ESXi 6.0 is no longer supported

Concepts

Learn about ONTAP Select

ONTAP Select is a software-only version of ONTAP that you can deploy as a virtual machine on a hypervisor host. It complements the suite of mainstream AFF and FAS ONTAP offerings as well as other software-only options such as Cloud Volumes ONTAP.

ONTAP Select converts internal disk drives, NVMe, SSD or HDD, and external array storage, into a flexible storage system with many of the same benefits offered by dedicated ONTAP storage systems. You can also deploy ONTAP Select on new servers or on your existing server infrastructure. ONTAP Select is easy to manage and leverages the same management software as ONTAP based solutions, which means reduced operational overhead and training requirements.

ONTAP Select adapts to your capacity consumption needs with two complementary licensing models — Capacity Tiers and the completely flexible Capacity Pools. Both licensing models allow you to increase capacity in increments of as little as 1TB. For example, you can start with a few terabytes and dynamically add capacity as your project grows. If you are using Capacity Pools, you can redistribute the capacity as you see fit when your project is finished.

ONTAP Select is integrated into cloud management frameworks such as VMware vSphere and Kernel-based Virtual Machine (KVM). This helps accelerate deployment of new projects such as file services, home directories, software development environments, and application testing.

SnapMirror software lets you move data between ONTAP storage across the hybrid cloud so that you can easily access your data where you want it. For example, you can bring up an environment quickly to prove an idea. Later, you can move the project to a dedicated storage platform for production deployment or make it more accessible in the cloud as part of a development workflow.

Software-defined storage

The implementation and delivery of IT services through software allows administrators to rapidly provision resources with a speed and agility that was previously not possible. As modern data centers move to a software-defined infrastructure (SDI) architecture, you can separate the most valuable IT assets from the underlying physical infrastructure, providing flexibility, scalability, and programmability.

In a commodity world where data is fragmented across silos of direct-attached storage (DAS), data access and management have become more complex problems. Software-defined storage (SDS) has emerged as an important part of the SDI landscape to address these and other issues.

ONTAP Select is the NetApp solution for the SDS market. ONTAP Select brings enterprise-class storage management features to the software-defined data center and extends the NetApp Data Fabric architecture to the extreme edge use cases, including the Internet of Things (IoT) and tactical servers.

Two software components

ONTAP Select is composed of two major software components:

ONTAP Select node

An ONTAP Select cluster is composed of one, two, four, six, eight, ten, or twelve nodes. Each cluster node is deployed as a separate virtual machine and runs a specially-designed version of the ONTAP 9 software.

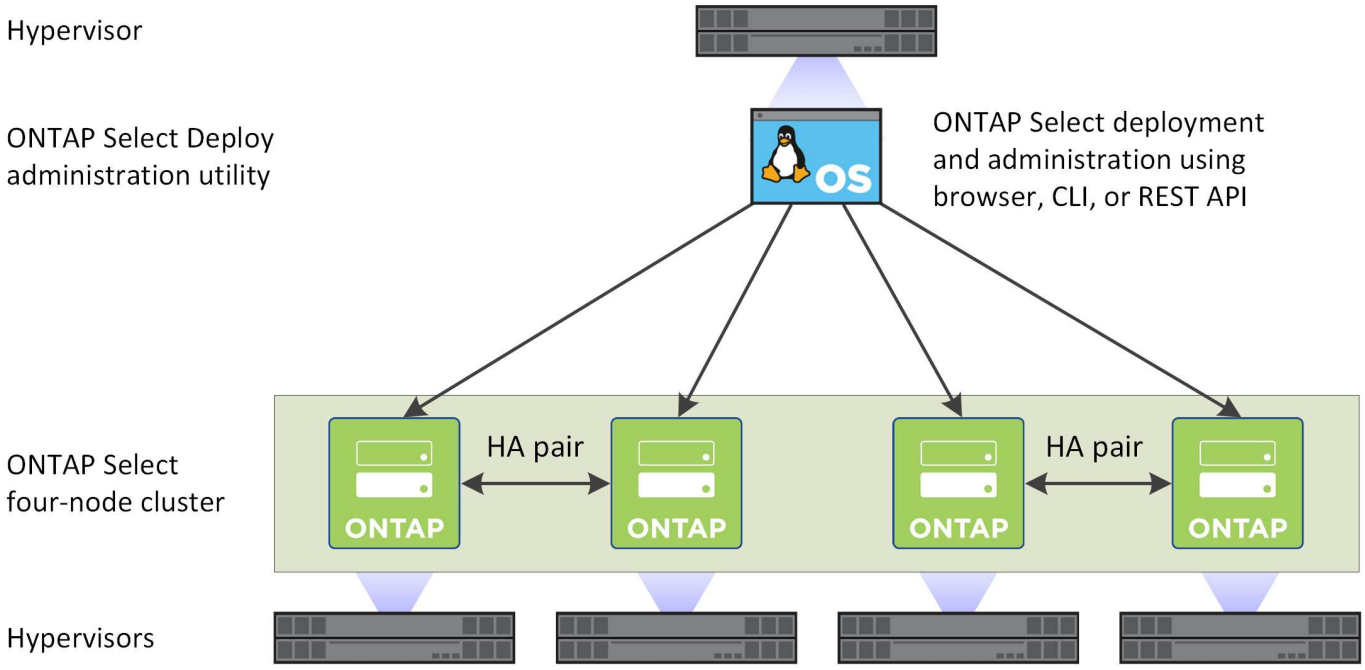
ONTAP Select Deploy administration utility

The Deploy administration utility is packaged and installed as a separate Linux virtual machine. You must use the utility to deploy ONTAP Select clusters in a production environment. A current version of the ONTAP Select node image is bundled with the Deploy utility.

> The Deploy administration utility is not assigned a separate version number. Deploy has the same version number as the associated ONTAP Select release. However, each update of the Deploy utility within a specific ONTAP Select release has a unique build number.

Illustration of a typical deployment

The following figure illustrates the NetApp ONTAP Select Deploy administration utility being used to deploy and support a four-node ONTAP Select cluster. The Deploy utility and ONTAP Select nodes run as separate virtual machines on dedicated hypervisor hosts.



FlexCache write-back with ONTAP Select

ONTAP Select supports FlexCache cache volumes with write-back. The overall ONTAP Select instance performance is comparable to an entry-level ONTAP system.

FlexCache cache volumes with write-back have the following requirements with ONTAP Select:

- Use ONTAP Select on a medium (8-core) or large (16-core) instance.
- Use NVMe-backed solid state drives. SAS-backed storage has a lower performance compared to NVMe-backed storage for ONTAP Select.
- Test ONTAP Select for FlexCache write-back before starting the deployment. FlexCache write-back is not recommended for a large number of small files.

ONTAP Select does not support FlexCache cache volumes with write-back on a small instance (4-core). Additionally, ONTAP Select does not support FlexCache origin volumes with write-back on small, medium, or large instances.

Learn about ONTAP FlexCache write-back

Compare ONTAP Select and ONTAP 9

Both hardware-based ONTAP and ONTAP Select provide enterprise class storage solutions. However, because they are designed and implemented differently, each can address different business requirements and usage scenarios. You should become familiar with the major differences between the platforms before planning an ONTAP Select deployment.

Different HA architecture

Depending on the number of nodes you define in a cluster, ONTAP Select provides an HA capability. For example, a four-node cluster consists of two HA pairs. The HA architecture used with ONTAP Select is based on a non-shared storage model. That is, one node in an HA pair cannot directly access the storage owned by the other node. This design can affect certain ONTAP Select operational characteristics.

Capacity licensing

ONTAP Select introduces a consumption-based licensing model. You must purchase a license with storage capacity for each node or shared Capacity Pool when deploying an ONTAP Select cluster in a production environment. Using the Deploy utility, you must apply the license files which establish the storage capacity for the cluster nodes.

ONTAP feature licensing

Each node in an ONTAP Select cluster is automatically licensed to use several ONTAP features. You do not need to manually install or apply these feature licenses.

ONTAP features not supported in ONTAP Select

Several ONTAP features are not supported with ONTAP Select. In most cases, these features require special hardware that is not available in the virtualized ONTAP Select environment.

- Automatic updates for Autonomous Ransomware Protection (ARP)

  Beginning with ONTAP Select 9.16.1, manual updates to ARP are supported and the ARP feature license is included in the NetApp License File (NLF).

- Cloud key managers
- Cluster IPspace

  Any modification to Cluster IPspace, including adding or removing ports, virtual LANs (VLANs), or link aggregation groups is not supported.

- Consistency groups in Asynchronous SnapMirror
- Fibre Channel

  Fibre Channel and Fibre Channel over Ethernet are not supported.

- Health monitors

  The traditional health monitoring used with a hardware-based ONTAP deployment is specific to the underlying hardware components. Due to the virtualized environment used with ONTAP Select, health

monitors are not active.

- Interface groups

  Interface groups are not supported.

- NIC offload support

  Due to the virtualized environment used with ONTAP Select, the NIC offload facility is not supported.

- NetApp storage encryption drives
- ONTAP port properties

  Modifying the properties of the ONTAP ports, including speed, duplex, and flow-control, is not supported.

- Service processors
- SVM-Migrate
- SnapLock Compliance
- SnapMirror active sync
- VMware HCX

Related information

- Learn about ONTAP features licensed by default
- Learn about ONTAP Select license options

Learn about ONTAP Select Deploy

ONTAP Select Deploy is the administration utility used to deploy and manage ONTAP Select clusters. Deploy is packaged as a Linux virtual machine which you must install before creating an ONTAP Select cluster.

Core functionality

The Deploy administration utility performs the following core functions:

- Record the details of each hypervisor host where ONTAP Select is deployed
- Configure the hosts and install the required licenses
- Deploy and manage the ONTAP Select clusters
- Maintain an inventory of ONTAP Select clusters and hosts
- Collect and send AutoSupport data to NetApp
- Maintain an internal set of ONTAP Select node images
- Support the hypervisor-specific command formats and protocols

Ways you can access the Deploy utility

You have several options available when accessing the Deploy administration utility. All the external interfaces are functionally equivalent. You should select the access option that best matches your specific deployment

goals and requirements. In all cases, you must sign in using the administrator account with a valid password.

Web graphical user interface

You can access the Deploy utility through a modern web browser. The web UI provides an intuitive and easy-to-use interface, and in most cases will be your primary interface when using the utility.

Command line interface

A text-based command line interface is available through the management shell. You can access the CLI management shell in the following ways:

- Secure shell (SSH)
- Virtual machine console

You typically use the virtual machine console as part of the installation and initial configuration process. However, in most situations SSH provides a more flexible and convenient option.

REST web services API

The REST web services API exposed to external clients provides another option when connecting to the Deploy utility. You can access the API using any mainstream programming language or tool that supports REST web services. Popular choices include:

- Python
- Java
- Curl

Using a programming or scripting language provides an opportunity to automate the deployment and management of the ONTAP Select clusters.

ONTAP Select online documentation web page

Displaying the online documentation web page at the Deploy utility is an alternative way to access the REST web services API. However, instead of using a programming language, you access the management API through the page using a browser. The following features are provided:

- A detailed description of every call in the REST web services API
- The ability to manually issue any of the API calls

You can access the online documentation page using the IP or domain name of the Deploy virtual machine. To display the page, enter a URL with the following format in your browser (substituting the appropriate IP address or domain name for your Deploy VM instance): `http://<ip_address>/api/ui`

Business use cases

ONTAP Select business needs and usage scenarios

ONTAP Select is suitable for several different types of applications based on the inherent flexibility provided through the hypervisor virtualization.

Deployment

From a high level, you can deploy ONTAP Select in two different ways regarding the workload on the hypervisor host servers.

Dedicated deployment

With the dedicated deployment model, a single instance of ONTAP Select runs on the host server. No other significant processing runs on the same hypervisor host.

Collocated deployment

With the collocated deployment model, ONTAP Select shares the host with other workloads. Specifically, there are additional virtual machines, each typically running computational applications. These compute workloads are local to the ONTAP Select cluster. This model supports specialized application and deployment requirements. As with the dedicated deployment model, each ONTAP Select virtual machine must run on a separate and dedicated hypervisor host.

Storage

ONTAP Select can be used as primary or secondary storage, depending on your business needs.

Primary storage

In certain cases, you may choose to deploy ONTAP Select as your primary storage platform. These types of implementations vary and depend on the workload characteristics of the applications as well as your business objectives.

Disaster recovery and secondary storage

You can use ONTAP Select to implement additional storage that augments your primary storage capabilities. The additional storage can be used to support your organization’s disaster recovery efforts and data backup plans.

Development and testing

As you deploy various applications within your organization, you can use ONTAP Select as an integral part of the overall application development and testing process. For example, you may need temporary storage to hold test input or output data. The length of these types of deployments can vary based on the application characteristics and requirements.

ONTAP Select use cases for remote and branch offices

Deploy ONTAP Select in remote office/branch office (ROBO) situations to support smaller offices while maintaining centralized administration and control.

The following ROBO configurations are supported:

- Two-node cluster with HA capability
- Single-node cluster

The ONTAP Select VM can be collocated with application VMs, making it an optimal solution for ROBOs.

Using ONTAP Select to provide enterprise-class file services while allowing bidirectional replication to other

ONTAP Select or FAS clusters enables resilient solutions to be built in low-touch or low-cost environments. ONTAP Select comes prepopulated with feature licenses for CIFS, NFS, and iSCSI protocol services as well as both SnapMirror and SnapVault replication technologies. Therefore, all of these features are available immediately upon deployment.

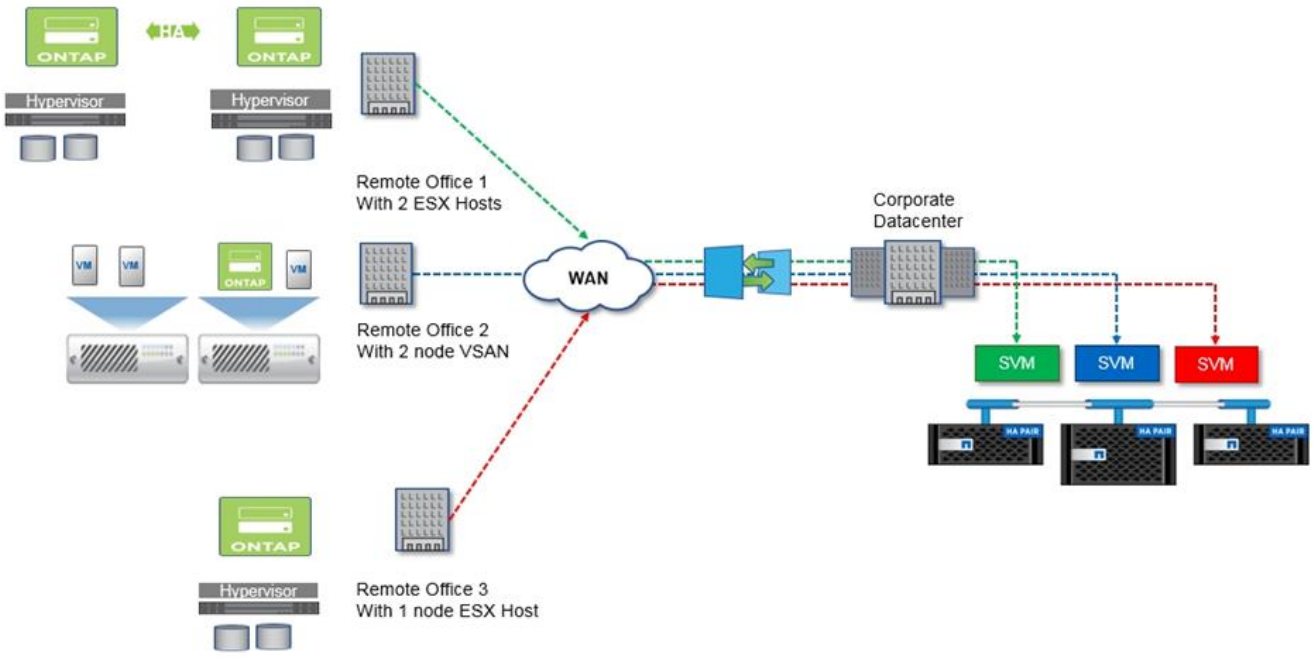
Because all VMware vSphere licenses are supported, you can choose the vSphere Remote Office Branch Office Standard or Advanced license instead of the Enterprise or Enterprise Plus license.

All vSphere and VSAN licenses are supported.

An ONTAP Select two-node cluster with a remote mediator is an attractive solution for small data centers. In this configuration, HA functionality is provided by ONTAP Select. The minimum networking requirement for a two-node ONTAP Select ROBO solution is four 1Gb links. A single 10Gb network connection is also supported. The vNAS ONTAP Select solution running on VSAN (including the two-node VSAN ROBO configuration) is another option. In this configuration, the HA functionality is provided by VSAN. Finally, a single-node ONTAP Select cluster replicating its data to a core location can provide a set of robust enterprise data management tools on top of a commodity server.

The following figure depicts a common remote office configuration using ONTAP Select on VM ESXi. Schedule-driven SnapMirror relationships periodically replicate the data from the remote office to a single consolidated engineered storage array located in the main data center.

Scheduled backup of remote office to corporate data center



ONTAP Select support for private cloud and data centers

ONTAP Select is ideally suited to support one or more private clouds within your organization. A common use case is to provide storage services for private clouds built on commodity servers.

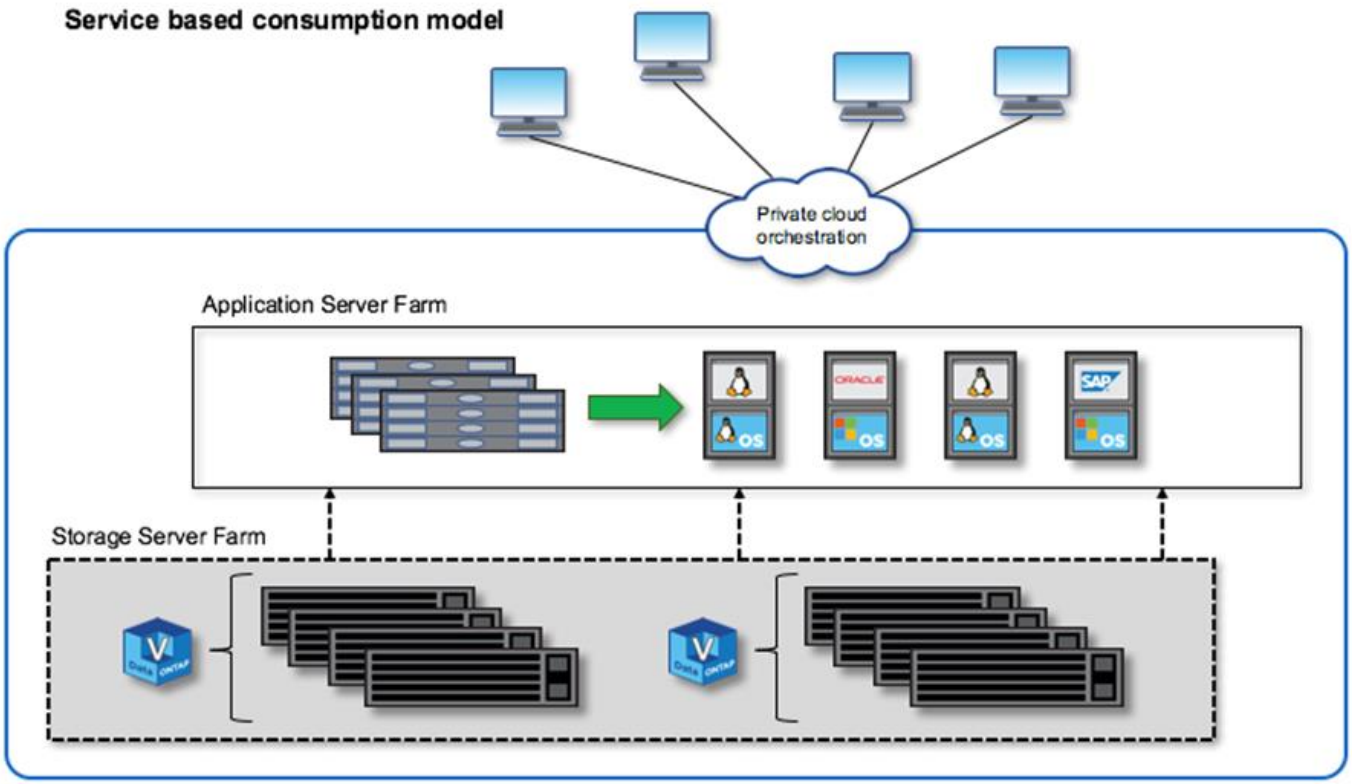
Like the public cloud, a private cloud provides flexibility as well as rapid setup and teardown. In addition, a

private cloud offers improved security and control.

The following figure shows how a storage farm provides computation and locally attached storage to the ONTAP Select VMs, which provide storage services upstream to an application stack. The entire workflow, from the provisioning of SVMs to the deployment and configuration of application VMs, is automated through a private cloud orchestration framework.

This is a service-oriented private cloud model. Using the HA version of ONTAP Select creates the same ONTAP experience you would expect on higher-cost FAS arrays. Storage server resources are consumed exclusively by the ONTAP Select VM, with application VMs hosted on separate physical infrastructure.

Private cloud built on DAS



ONTAP Select data protection and efficiency

ONTAP Select is built on ONTAP storage software to deliver enterprise storage services efficiently with a highly available, shared-nothing scale-out architecture. You can deploy a solution with one, two, four, six, eight, ten, or twelve nodes with up to 400TB of raw capacity for NFS, SMB/CIFS, and iSCSI connected storage per node. You can leverage native deduplication and compression to lower storage costs by increasing your effective capacity. The scale-out architecture enables high availability and non-disruptive data access for load balancing or to service hardware.

Snapshot data protection

ONTAP Select includes data protection features, including snapshot and SnapMirror software. You can quickly replicate your data to other ONTAP storage, whether on premises, at a remote site, or in the cloud. If you must

recover your data quickly, SnapRestore software can use local snapshots to recover entire file systems or data volumes in seconds, regardless of the capacity or number of files.

MetroCluster software defined storage

ONTAP Select MetroCluster software defined storage (SDS) offers enhanced protection and a cost effective implementation.

A two-node cluster can be stretched between two locations if certain minimum requirements are met. This architecture fits neatly in between hardware-based MetroCluster and single data-center clusters (hardware-defined or software-defined). The requirements for the ONTAP Select MetroCluster SDS highlight the general flexibility of software-defined storage solutions as well as the differences between it and the hardware-based MetroCluster SDS. No proprietary hardware is required.

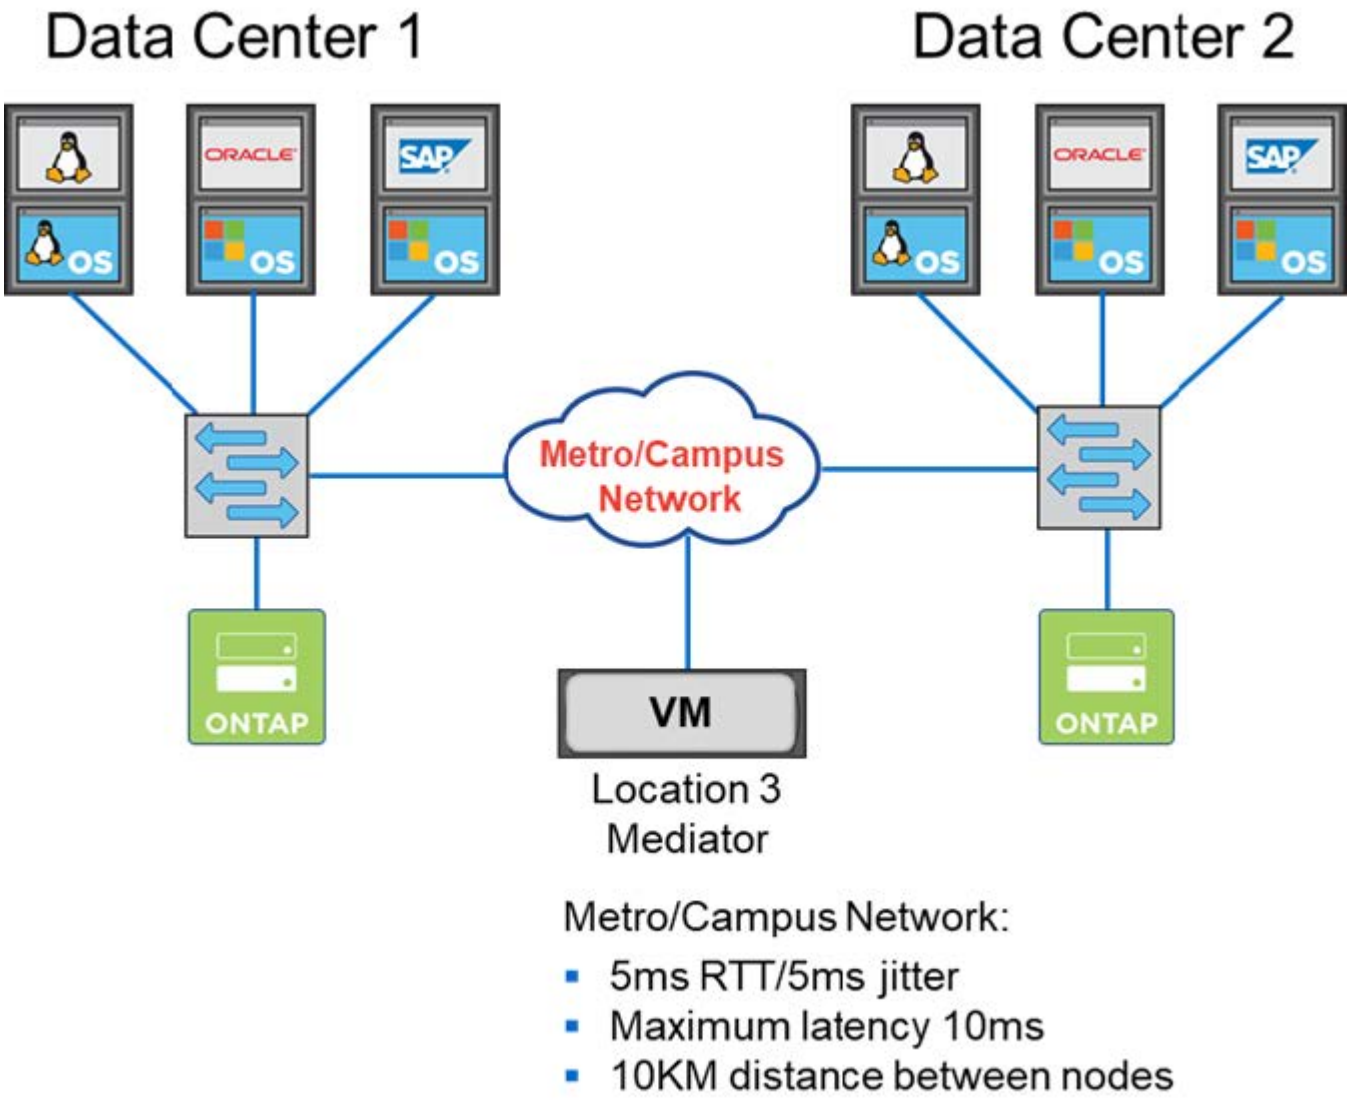
Unlike MetroCluster, ONTAP Select uses the existing network infrastructure and supports a network latency of up to 5ms RTT with a maximum jitter of up to 5ms, for a total of 10ms maximum latency. Separation requirements in the market space relate more to physical separation than the actual distance. In some instances, this can mean different buildings. In other instances, it can mean different rooms in the same building. Regardless of the actual physical placement, what defines a two-node cluster as a MetroCluster SDS is that each node uses a separate uplink switch.

As part of the two-node HA configuration, a mediator is required to properly identify the active node during a failover and avoid any split-brain scenario in which both nodes remain independently active during a network partition. This operation is identical to the regular two-node HA configuration previously available. For proper protection and failover during site failure, the mediator should be in a different site from the two HA nodes. The maximum latency between the mediator and each ONTAP Select node cannot exceed 125ms.

ONTAP Select MetroCluster SDS provides the following benefits:

- MetroCluster SDS provides another dimension (data center to data center) of protection for ONTAP Select. You can now take advantage of this extra level of protection in addition to leveraging all the benefits of software-defined storage and ONTAP.
- MetroCluster SDS provides business-critical data protection with a 0 RPO and automatic failover. Both the data storage and the application access points are automatically switched over to the surviving data center or node with zero intervention from IT.
- MetroCluster SDS is cost effective. It takes advantage of the existing networking infrastructure to enable stretched resiliency between the HA pair, and no additional hardware is required. It also provides active/active data access and data center redundancy in the same cluster.

MetroCluster SDS



For more best practices and other requirements, refer to the sections Two-node HA versus multi-node HA and Two-node stretched HA (MetroCluster SDS) best practices.

Related information

ONTAP Select support for ONTAP features

ONTAP Select terminology and key concepts

As you begin to explore ONTAP Select and plan a deployment, it is helpful to first become familiar with the terminology and key concepts.

ONTAP Select Deploy

ONTAP Select Deploy is the administration utility that you use to deploy ONTAP Select clusters. The Deploy utility runs in a dedicated Linux virtual machine. You can access the Deploy utility through the web user interface, CLI management shell, and REST API.

Kernel-based Virtual Machine

Kernel-based Virtual Machine (KVM) is a virtualization feature of the Linux kernel, which allows it to act as a

hypervisor platform. A wide range of guest operating systems are supported.

Hypervisor host versus ONTAP Select node

A *hypervisor host* is the core hardware platform that hosts an ONTAP Select virtual machine. When an ONTAP Select virtual machine is deployed and active on a hypervisor host, it is considered to be an *ONTAP Select node*.

ONTAP Select cluster

You can create an *ONTAP Select cluster* composed of one, two, four, six, eight, ten, or twelve nodes. Multi-node clusters always contain one or more HA pairs. For example, a four-node cluster consists of two HA pairs. A single node cluster does not provide HA capability.

Preparation of the hypervisor host environment

Before using the Deploy administration utility to deploy an ONTAP Select cluster, you need to prepare the hypervisor hosts where ONTAP Select will run, including the storage and networking environments. This host pre-configuration is done outside of the ONTAP Select product based on the current requirements and limitations.

Evaluation versus production deployments

Every ONTAP Select node runs with either an *evaluation license* or a *purchased license*. An evaluation license allows you to evaluate ONTAP Select prior to deploying it in a production environment. The evaluation license is automatically generated and applied. If you deploy a cluster in a production environment, you must purchase a license which involves choosing:

- Licensing model
- Storage capacity
- Platform license offering

Capacity Tiers licensing model

The Capacity Tiers licensing model is the original option when licensing storage for an ONTAP Select deployment. It is based on the ONTAP model used with NetApp AFF and FAS. A separate license is required for each node. The storage capacity is locked to the node and perpetual (no renewal required).

Capacity Pools licensing model

The Capacity Pools licensing model was introduced with ONTAP Select 9.5 using Deploy 2.10. A separate license is required for each storage Capacity Pool. The Capacity Pool license is locked to a License Manager instance (that is a Deploy instance) and must be renewed based on the terms of your purchase. You can license and use any number of Capacity Pools in your organization. However, because the Capacity Pools are shared by the ONTAP Select nodes, fewer licenses are typically required than Capacity Tiers licensing.

License Manager

The License Manager is a software component which supports Capacity Pools licensing. It is currently part of the Deploy administration utility. LM leases storage to the ONTAP Select nodes from the shared pools it manages. The *License Lock ID* is a numeric string uniquely identifying each LM instance, and therefore each Deploy instance. You must use both the Capacity Pool license serial number and LLID to generate a license file.

Platform license offerings

There are three license offerings available which determine the size capabilities of the ONTAP Select virtual machine when you purchase a license:

- Standard
- Premium
- Premium XL

For more information, refer to the two sections Plan and License.

Storage pools versus datastores

An ONTAP Select *storage pool* is a logical data container designed to abstract and hide the underlying physical storage. A storage pool is hypervisor-independent. When deployed on an ESXi hypervisor host, the ONTAP Select storage pool is synonymous with the VMware *datastore*.

Cluster MTU

Cluster MTU is a feature allowing you to configure the MTU size used on the internal network used with an ONTAP Select multi-node cluster. The Deploy administration utility adjusts the MTU size as you configure the HA pairs to accommodate your networking environment. You can also manually set the value.

ONTAP Select vNAS

The ONTAP Select vNAS solution allows an ONTAP Select node to access VMware datastores on external storage. With ONTAP Select vNAS, a local RAID controller is no longer needed; the RAID functionality is assumed to be provided by the remote storage. ONTAP Select vNAS can be configured in the following ways:

- VMware vSAN
- Generic external storage array

In both cases, the external storage must be configured prior to creating an ONTAP Select cluster or expanding the storage capacity of an existing node.

Node re-hosting on the ESXi VM

When you deploy a cluster that uses external storage available through the ONTAP Select vNAS solution (either VMware vSAN or a generic external storage array), the ESXi virtual machine hosting the ONTAP Select node can be moved through actions utilizing the following VMware features:

- vMotion
- High Availability (HA)
- Distributed Resource Scheduler (DRS)

  The ONTAP Select Deploy utility detects the movement of the virtual machine as part of executing an operation on the cluster, such as:

- cluster online
- cluster offline
- storage add

  When a virtual machine is moved, the Deploy utility updates its internal database and configures the new ESXi host. All actions performed on the ONTAP Select node are blocked until the movement of the virtual machine and Deploy updates are completed.

Open vSwitch for KVM

Open vSwitch (OVS) is a software implementation of a virtual switch supporting multiple networking protocols. OVS is open source and available according to the Apache License 2.0.

Mediator service

The ONTAP Select Deploy utility includes a mediator service which connects to the nodes in the active two-node clusters. This service monitors each HA pair and assists in managing failures.

> If you have one or more active two-node clusters, the ONTAP Select Deploy virtual machine administering the clusters must be running at all times. If the Deploy virtual machine is halted, the mediator service is unavailable and HA capability is lost for the two-node clusters.

MetroCluster SDS

MetroCluster SDS is a feature that provides an additional configuration option when deploying a two-node ONTAP Select cluster. Unlike a typical two-node ROBO deployment, the MetroCluster SDS nodes can be separated by a much greater distance. This physical separation enables additional use cases, such as disaster recovery. You must have a premium license or higher to use MetroCluster SDS. In addition, the network between the nodes must support a minimum latency requirement.

Credential store

The Deploy credential store is a secure database holding account credentials. It is used primarily to register hypervisor hosts as part of creating a new cluster. Refer to the Plan section for more information.

Storage efficiency

ONTAP Select provides storage efficiency options that are similar to the storage efficiency options present on FAS and AFF arrays. Conceptually, ONTAP Select with direct-attached storage (DAS) SSDs (using a premium license) is similar to an AFF array. Configurations using DAS with HDDs and all vNAS configurations should be considered similar to a FAS array. The main difference between the two configurations is that ONTAP Select with DAS SSDs supports inline aggregate level deduplication and aggregate level background deduplication. The remaining storage efficiency options are available for both configurations.

The vNAS default configurations enable a write optimization feature known as single instance data logging (SIDL). With ONTAP Select 9.6 and later releases, the background ONTAP storage efficiency features are qualified with SIDL enabled. Refer to the Deep dive section for more information.

Cluster refresh

After creating a cluster, you can make changes to the cluster or virtual machine configuration outside of the Deploy utility using ONTAP or hypervisor administration tools. You can also migrate a virtual machine which causes configuration changes. When these changes occur, the Deploy utility is not automatically updated and can become out of sync with the state of the cluster. You can use the cluster refresh feature to update the Deploy configuration database. Cluster refresh is available through the Deploy web user interface, CLI management shell, and REST API.

Software RAID

When using direct-attached storage (DAS), RAID functionality is traditionally provided through a local hardware RAID controller. You can instead configure a node to use *software RAID* where the ONTAP Select node provides the RAID functionality. If you use software RAID, a hardware RAID controller is no longer needed.

ONTAP Select image install

The Deploy administration utility contains only a single version of ONTAP Select. The version included is the most current available at the time of release. The ONTAP Select image install feature allows you to add earlier versions of ONTAP Select to your instance of the Deploy utility, which can then be used when deploying an ONTAP Select cluster. See Add ONTAP Select images for more information.

You should only add an ONTAP Select image with a version that is earlier than the original version included with your instance of Deploy. Adding later versions of ONTAP Select without also updating Deploy is not supported.

Administering an ONTAP Select cluster after it is deployed

After you deploy an ONTAP Select cluster, you can configure the cluster as you would a hardware-based ONTAP cluster. For example, you can configure an ONTAP Select cluster using System Manager or the standard ONTAP command line interface.

Related information

Add an ONTAP Select image to Deploy

Plan

ONTAP Select installation and deployment workflow

You can use the following workflow to deploy and administer an ONTAP Select cluster.

1. **Plan the deployment and prepare the environment**

Review the planning requirements for your ONTAP Select nodes, storage, networking, and hypervisor environment. Preparing your environment in advance helps you avoid deployment issues.

2. **Install the ONTAP Select Deploy administration utility**

Install the utility that you use to deploy and manage ONTAP Select clusters.

3. **Acquire the licenses required for a production deployment**

Purchase and apply the licenses needed for your production deployment. You can choose between the Capacity Tiers and Capacity Pools licensing models.

4. **Deploy an ONTAP Select cluster using the ONTAP Select Deploy administration utility**

Use the Deploy utility to create your ONTAP Select cluster:

- Deploy using the web UI
- Deploy using the CLI
- Deploy using the REST API

5. **Administer the ONTAP Select cluster using the standard NetApp management tools and interfaces**

Administer your ONTAP Select cluster using the same management tools and interfaces you use for other ONTAP clusters.

ONTAP Select

ONTAP Select requirements and planning considerations

There are several general requirements you should consider as part of planning an ONTAP Select deployment.

Required Linux knowledge and skills for KVM

Linux with the KVM hypervisor is a complex environment to work in. Before deploying ONTAP Select on KVM, you must have the necessary knowledge and skills.

Linux server distribution

You should have experience with the specific Linux distribution to be used for your ONTAP Select Deployment. Specifically, you should be able to perform the following tasks:

- Install the Linux distribution
- Configure the system using the CLI
- Add software packages as well as any dependencies

For more information on preparing your Linux server, including the required configuration and software packages, refer to the host configuration checklist. Refer to the hypervisor requirements for the currently supported Linux distributions.

KVM deployment and administration

You should be familiar with general virtualization concepts. In addition, there are several Linux CLI commands that you must use as part of installing and administering ONTAP Select in a KVM environment:

- `virt-install`
- `virsh`
- `lsblk`
- `lvs`
- `vgs`
- `pvs`

Networking and Open vSwitch configuration

You should be familiar with networking concepts and the configuration of network switches. In addition, you should have experience with Open vSwitch. You must use the following network commands a part of configuring the ONTAP Select network in a KVM environment:

- `ovs-vsctl`
- `ip`
- `ip link`
- `systemctl`

Cluster size and related considerations

There are several planning issues related to the cluster size that you should consider.

Number of nodes in the cluster

An ONTAP Select cluster is composed of one, two, four, six, eight, ten, or twelve nodes. You should determine the size of the cluster based on the application requirements. For example, if HA capability is needed for an enterprise deployment, then a multi-node cluster should be used.

Dedicated versus collocated

Based on the application type, you should determine if the deployment follows the dedicated or collocated model. Note that the collocated model can be more complex due to the workload diversity and tighter integration.

Hypervisor host considerations

There are several planning issues related to the hypervisor host that you should consider.

> You should not directly modify the configuration of an ONTAP Select virtual machine unless directed to do so by NetApp support. A virtual machine should only be configured and modified through the Deploy administration utility. Making changes to an ONTAP Select virtual machine outside of the Deploy utility without assistance from NetApp support can cause the virtual machine to fail and render it unusable.

Hypervisor independent

Both ONTAP Select and the ONTAP Select Deploy administration utility are hypervisor-independent. The following hypervisors are supported for both.

- VMware ESXi
- Kernel-based Virtual Machine (KVM)

> Beginning with ONTAP Select 9.14.1, support for KVM hypervisor has been reinstated. Previously, support for deploying a new cluster on a KVM hypervisor was removed in ONTAP Select 9.10.1 and support for managing existing KVM clusters and hosts, except to take offline or delete, was removed in ONTAP Select 9.11.1.

Refer to the hypervisor-specific planning information and release notes for additional details regarding the supported platforms.

Hypervisor for ONTAP Select nodes and administration utility

Both the Deploy administration utility and ONTAP Select nodes run as virtual machines. The hypervisor you choose for the Deploy utility is independent of the hypervisor you choose for the ONTAP Select nodes. You have complete flexibility when pairing the two:

- Deploy utility running on VMware ESXi can create and manage ONTAP Select clusters on either VMware ESXi or KVM
- Deploy utility running on KVM can create and manage ONTAP Select clusters on either VMware ESXi or KVM

One or more instances of ONTAP Select node per host

Each ONTAP Select node runs as a dedicated virtual machine. You can create multiple nodes on the same hypervisor host, with the following restrictions:

- Multiple nodes from a single ONTAP Select cluster cannot run on the same host. All the nodes on a specific host must be from different ONTAP Select clusters.
- You must use external storage.
- If you use software RAID, you can only deploy one ONTAP Select node on the host.

Hypervisor consistency for the nodes within a cluster

All of the hosts within an ONTAP Select cluster must run on the same version and release of the hypervisor software.

Number of physical ports on each host

You must configure each host to use one, two, or four physical ports. Although you have flexibility when configuring the network ports, you should follow these recommendations where possible:

- A host in a single-node cluster should have two physical ports.
- Each host in a multi-node cluster should have four physical ports

Integrating ONTAP Select with an ONTAP hardware-based cluster

You cannot add an ONTAP Select node directly to an ONTAP hardware-based cluster. However, you can optionally establish a cluster peering relationship between an ONTAP Select cluster and a hardware-based ONTAP cluster.

Storage considerations

There are several planning issues related to host storage that you should consider.

RAID type

When using direct-attached storage (DAS) on ESXi, you should decide whether to use a local hardware RAID controller or the software RAID feature included with ONTAP Select. If you use software RAID, see Storage and RAID considerations for more information.

Local storage

When using local storage managed by a RAID controller, you must decide the following:

- Whether to use one or more RAID groups
- Whether to use one or more LUNs

External storage

When using the ONTAP Select vNAS solution, you must decide where the remote datastores are located and how they are accessed. ONTAP Select vNAS supports the following configurations:

- VMware vSAN
- Generic external storage array

Estimate for the storage needed

You should determine how much storage is required for the ONTAP Select nodes. This information is required as part of acquiring the purchased licenses with storage capacity. Refer to Storage capacity restrictions for more information.



The ONTAP Select storage capacity corresponds to the total allowable size of the data disks attached to the ONTAP Select virtual machine.

Licensing model for production deployment

You must select the Capacity Tiers or Capacity Pools licensing model for each ONTAP Select cluster deployed in a production environment. Review the section *License* for more information.

Authentication using the credential store

The ONTAP Select Deploy credential store is a data base holding account information. Deploy uses the account credentials to perform host authentication as part of cluster creation and management. You should be aware of how the credential store is used as part of planning an ONTAP Select deployment.



The account information is stored securely in the data base using the Advanced Encryption Standard (AES) encryption algorithm and SHA-256 hashing algorithm.

Types of credentials

The following types of credentials are supported:

- host

  The **host** credential is used to authenticate a hypervisor host as part of deploying an ONTAP Select node directly to ESXi or KVM.

- vcenter

  The **vcenter** credential is used to authenticate a vCenter server as part of deploying an ONTAP Select node to ESXi when the host is managed by VMware vCenter.

Access

The credential store is accessed internally as part of performing normal administrative tasks using Deploy, such as adding a hypervisor host. You can also manage the credential store directly through the Deploy web user interface and CLI.

Related information

- Storage and RAID considerations

ONTAP Select VMware hypervisor and hardware considerations

There are several hardware requirements and planning issues you should consider related to the VMware environment.

Hypervisor requirements

There are several requirements related to the hypervisor where ONTAP Select runs.

> You should review the ONTAP Select Release Notes for any additional known restrictions or limitations.

VMware licensing

To deploy an ONTAP Select cluster, your organization must have a valid VMware vSphere license for the hypervisor hosts where ONTAP Select runs. You should use the licenses that are appropriate for your deployment.

Software compatibility

You can deploy ONTAP Select on KVM and ESXi hypervisors.

KVM

ONTAP Select supports the following KVM hypervisor versions:

- KVM on Red Hat Enterprise Linux (RHEL) 10.1, 10.0, 9.7, 9.6, 9.5, 9.4, 9.2, 9.1, 9.0, 8.8, 8.7, and 8.6
- KVM on Rocky Linux 10.1, 10.0, 9.7, 9.6, 9.5, 9.4, 9.3, 9.2, 9.1, 9.0, 8.9, 8.8, 8.7, and 8.6

ESXi

ONTAP Select supports the following ESXi hypervisor versions:

- VMware ESXi 9.0
- VMware ESXi 8.0 U3
- VMware ESXi 8.0 U2
- VMware ESXi 8.0 U1 (build 21495797)
- VMware ESXi 8.0 GA (build 20513097)

NetApp supports ONTAP Select on the identified versions of ESXi as long as VMware also continues to support the same versions.

ESXi 7.0 GA reached end of availability status. If you have ONTAP Select clusters with this version, you must upgrade to the supported versions as per the Interoperability Matrix Tool (IMT).

VMware vCenter and standalone ESXi hosts

If an ESXi hypervisor host is managed by a vCenter server, you must register the host to the Deploy administration utility using the vCenter credentials. You cannot register the host as a standalone host using the ESXi credentials.

Core hardware requirements

The physical hypervisor host where you deploy ONTAP Select must meet several hardware requirements. You can choose any platform for the hypervisor host, as long as it meets the minimum hardware requirements. The following vendors provide supported hardware platforms: Cisco, Dell, HP, Fujitsu, Lenovo, and Supermicro.

ONTAP Select supports CPU models based on Intel Xeon Sandy Bridge and Advanced Micro Devices (AMD) processors.

Refer to the Interoperability Matrix Tool for more information.

Basic hardware requirements

There are several common hardware requirements that apply to all platforms regardless of the node instance type or license offering.

Processor

The supported microprocessors include Intel Xeon server processors and AMD processors.

Refer to Intel Xeon Processors and AMD EPYC Server CPUs for more information.

Ethernet configuration

There are several supported Ethernet configurations based on the cluster size.

Cluster size	Minimum requirements	Recommended requirements
Single node cluster	2 x 1GbE	2 x 10GbE
Two-node cluster or MetroCluster SDS	4 x 1GbE or 1 x 10GbE	2 x 10GbE

Cluster size	Minimum requirements	Recommended requirements
Four, six, eight, ten, or twelve-node cluster	2 x 10GbE	4 x 10GbE or 2 x 25/40GbE

Additional hardware requirements based on the instance type

There are several additional hardware requirements based on the node instance type.

Node size	CPU cores	Memory	Required platform license offering
Small	Six physical cores or greater, with four reserved for ONTAP Select	24GB or greater, with 16GB reserved for ONTAP Select	Standard, premium, or premium XL
Medium	Ten physical cores or greater, with eight reserved for ONTAP Select	72GB or greater, with 64GB reserved for ONTAP Select	Premium or premium XL
Large	Eighteen physical cores or greater, with sixteen reserved for ONTAP Select	136GB or greater, with 128GB reserved for ONTAP Select	Premium XL

There are additional disk requirements based on the platform license. See Storage and RAID for more information.

Related information

Learn about the platform license offerings for ONTAP Select

ONTAP Select storage and RAID considerations

There are several planning issues related to ONTAP Select host storage that you should consider.

External storage support information is outlined for ESXi and KVM hosts. For more information, refer to VMware ESXi requirements and KVM requirements.

Hardware RAID controller requirements

The RAID controller on the hypervisor host where you deploy ONTAP Select must meet several requirements.

A host where ONTAP Select runs requires local physical drives when using a hardware RAID controller or the software RAID capability provided with ONTAP Select. If you use the ONTAP Select vNAS solution to access external storage, a local RAID controller and software RAID capability are not used.

The minimum requirements for the RAID controller include:

- 12 Gbps throughput

- 512 MB internal battery-backed or flash (SuperCAP) cache
- Configured in write back mode:
  - Enable failback mode to “write through” (if supported)
  - Enable “always read ahead” policy (if supported)
- All local disks behind the RAID controller should be configured as a single RAID group; use multiple RAID controllers if needed:
  - Disable the local drive cache for RAID group, which is fundamental to preserving data integrity.
- Perform LUN configuration based on the following guidelines:
  - If the RAID group size exceeds the maximum LUN size of 64TB, configure multiple equal-sized LUNs consuming all the available storage within the RAID group.
  - If the RAID group size is smaller than the maximum LUN size of 64TB, configure one LUN consuming all available storage within the RAID group.

Software RAID requirements

When deploying an ONTAP Select cluster on the hypervisor, you can use the software RAID capability provided by ONTAP Select instead of a local hardware RAID controller. Review the following requirements and restrictions before deploying a cluster using software RAID.

General requirements

The environment for a software RAID deployment must meet the following core requirements:

- VMware ESXi 8.0 GA (build 20513097) or later
- ONTAP Select premium license or higher
- Local SSD drives only
- Separation of system disks from the root and data aggregates
- No hardware RAID controller on the host



If a hardware RAID controller is present, refer to the Deep dive storage section for additional configuration requirements.

ESXi specific requirements

- VMware ESXi 8.0 GA (build 20513097) or later
- VMware VMotion, HA, and DRS are not supported
- You cannot use software RAID with a node that has been upgraded from ONTAP Select 9.4 or earlier. If this is the case, you need to create a new node for software RAID deployment.

KVM specific requirements

There are also specific software package configuration requirements. Refer to the preparation of the Linux server step for more information.

Media expectations for KVM

The SSD flash storage devices used must meet the following additional requirements:

- The SSD devices must accurately and persistently report themselves to the Linux host through the following methods:
  - # cat /sys/block/<device>/queue/rotational

    The value reported for these commands must be '0'.
- It is expected that the devices are connected to an HBA or, in some cases, to a RAID controller configured to operate in JBOD mode. When using a RAID controller, the device function must be passed through the host without overlaying any RAID functionality. When using a RAID controller in JBOD mode, you should review the RAID documentation or contact the vendor as needed to make sure the device reports the rotational speed as '0'.
- There are two separate storage components:
  - Virtual machine storage

    This is an LVM pool (storage pool) containing the system data used to host the ONTAP Select virtual machine. The LVM pool must be backed by a high endurance flash device, and can be either SAS, SATA, or NVMe. An NVMe device is recommended for improved performance.
  - Data disks

    This is a set of SAS or SATA SSD drives used for data management. The SSD devices should be enterprise grade and durable. The NVMe interface is not supported.
- All devices must be formatted with 512BPS.

ONTAP Select node configuration

You must configure each ONTAP Select node and hypervisor host as follows to separate the system disks from the root and data aggregates:

- Create a system storage pool
  You must create a storage pool for the ONTAP Select system data. You must attach the storage pool as part of configuring the ONTAP Select node.
- Attach necessary physical disks
  The hypervisor host must have the required SSD disks attached and available for use by the ONTAP Select virtual machine. These drives hold the root and data aggregates. You must attach the storage disks as part of configuring the ONTAP Select node.

Storage capacity restrictions

As part of planning an ONTAP Select deployment, you should be aware of the restrictions related to storage allocation and use.

The most important storage restrictions are presented below. You should also review the Interoperability Matrix Tool for more detailed information.



ONTAP Select enforces several restrictions related to storage allocation and use. Before you deploy an ONTAP Select cluster or purchase a license, you should be familiar with these restrictions. Refer to the License section for more information.

Calculate raw storage capacity

The ONTAP Select storage capacity corresponds to the total allowable size of the virtual data and root disks attached to the ONTAP Select virtual machine. You should consider this when allocating capacity.

Minimum storage capacity for a single-node cluster

The minimum size of the storage pool allocated for the node in a single-node cluster is:

- Evaluation: 500 GB
- Production: 1.0 TB

The minimum allocation for a production deployment consists of 1 TB for user data, plus approximately 266 GB used by various ONTAP Select internal processes, which is considered required overhead.

Minimum storage capacity for a multi-node cluster

The minimum size of the storage pool allocated for each node in a multi-node cluster is:

- Evaluation: 1.9 TB
- Production: 2.0 TB

The minimum allocation for a production deployment consists of 2 TB for user data, plus approximately 266 GB used by various ONTAP Select internal processes, which is considered required overhead.



Each node in an HA pair must have the same storage capacity.

When estimating the amount of storage for an HA pair, you must consider that all aggregates (root and data) are mirrored. As a result, each plex of the aggregate consumes an equal amount of storage.

For example, when a 2TB aggregate is created, it allocates 2TB to two plex instances (2TB for plex0 and 2TB for plex1) or 4TB of the total licensed amount of storage.

Storage capacity and multiple storage pools

You can configure each ONTAP Select node to use up to 400 TB of storage when using local direct-attached storage, VMware vSAN, or external storage arrays. However, a single storage pool has a maximum size of 64 TB when using direct-attached storage or external storage arrays. Therefore, if you plan to use more than 64 TB of storage in these situations, you must allocate multiple storage pools as follows:

- Assign the initial storage pool during the cluster creation process
- Increase the node storage by allocating one or more additional storage pools



A 2% buffer is left unused in each storage pool and does not require a capacity license. ONTAP Select does not use this storage unless you specify a capacity cap. If you specify a capacity cap, then ONTAP Select uses that amount of storage unless the amount specified falls in the 2% buffer zone. The buffer is needed to prevent occasional errors that occur when attempting to allocate all of the space in a storage pool.

Storage capacity and VMware vSAN

When using VMware vSAN, a datastore can be larger than 64 TB. However, you can only initially allocate up to

64 TB when creating the ONTAP Select cluster. After the cluster is created, you can allocate additional storage from the existing vSAN datastore. The vSAN datastore capacity that can be consumed by ONTAP Select is based on the VM storage policy set.

Best practices

You should consider the following recommendations regarding the hypervisor core hardware:

- All drives in a single ONTAP Select aggregate should be the same type. For example, you should not mix HDD and SSD drives in the same aggregate.

Additional disk drive requirements based on the platform license

The drives you choose are limited based on the platform license offering.

> The disk drive requirements apply when using a local RAID controller and drives, as well as software RAID. These requirements do not apply to external storage accessed through the ONTAP Select vNAS solution.

Standard

- 8 to 60 internal HDD (NL-SAS, SATA, 10K SAS)

Premium

- 8 to 60 internal HDD (NL-SAS, SATA, 10K SAS)
- 4 to 60 internal SSD

Premium XL

- 8 to 60 internal HDD (NL-SAS, SATA, 10K SAS)
- 4 to 60 internal SSD
- 4 to 14 internal NVMe

> Software RAID with local DAS drives is supported with the premium license (SSD only) and premium XL license (SSD or NVMe).

NVMe drives with software RAID

You can configure software RAID to use NVMe SSD drives. Your environment must meet the following requirements:

- ONTAP Select with a supported Deploy administration utility
- Premium XL platform license offering or a 90-day evaluation license
- VMware ESXi version 8.0 or later
- NVMe devices conforming to specification 1.0 or later

You need to manually configure the NVMe drives before using them. See Configure a host to use NVMe drives for more information.

External storage requirements

VMware ESXi external storage requirements for ONTAP Select

ONTAP Select vNAS is a solution allowing the ONTAP Select data stores to be external to the ESXi hypervisor host where the ONTAP Select virtual machine runs. These remote data stores can be accessed through VMware vSAN or a generic external storage array.

Basic requirements and restrictions

The ONTAP Select vNAS solution can be used with an ONTAP Select cluster of any size.

All related storage components, including hardware, software, and feature requirements, must adhere to the requirements described in the Interoperability Matrix Tool. In addition, ONTAP Select supports all external storage arrays described in the VMware Storage/SAN Compatibility documentation, including iSCSI, NAS (NFSv3), Fibre Channel, and Fibre Channel over Ethernet. External array support is limited by the ESXi version supported by ONTAP Select.

The following VMware features are supported when deploying a cluster with ONTAP Select vNAS:

- VMotion
- High Availability (HA)
- Distributed Resource Scheduler (DRS)

> These VMware features are supported with single-node and multi-node ONTAP Select clusters. When deploying a multi-node cluster, you should make sure that two or more nodes from the same cluster do not run on the same hypervisor host.

The following VMware features are not supported:

- Fault Tolerance (FT)
- Virtual datastore (VVOL)

Configuration requirements

If you plan to use a VMFS datastore on an external storage array (iSCSI, Fibre Channel, Fibre Channel over Ethernet), you must create a VMFS storage pool before configuring ONTAP Select to use the storage. If you use an NFS datastore, there is no need to create a separate VMFS datastore. All vSAN datastores must be defined within the same ESXi cluster.

> You must provide a capacity limit for every datastore on VMware vSAN or an external storage array when configuring a host or performing a storage add operation. The capacity you specify must be within the allowed storage limits of the external storage. An error will occur if you do not provide a capacity limit or the external storage runs out of space during the disk creation operation.

Best practices

Review the VMware documentation and follow the applicable best practices identified for ESXi hosts. In addition:

- Define dedicated network ports, bandwidth, and vSwitch configurations for the ONTAP Select networks and external storage (VMware vSAN and generic storage array traffic when using iSCSI or NFS)

- Configure the capacity option to restrict storage utilization (ONTAP Select cannot consume the entire capacity of an external vNAS datastore)
- Verify that all generic external storage arrays use the available redundancy and HA features where possible

KVM external storage requirements for ONTAP Select

You can configure ONTAP Select on the KVM hypervisor with an external storage array.

Basic requirements and restrictions

If you use an external array for the ONTAP Select storage pools, the following configuration restrictions apply:

- You must define as the logical pool type using CLVM.
- You must provide a storage capacity limit.
- The configuration only supports FC, Fibre Channel over Ethernet (FCoE), and iSCSI protocols.
- The configuration does not recognize thinly provisioned storage.

> The storage capacity you specify must be within the allowed storage limits of the external storage. An error occurs if you do not provide a capacity limit or the external storage runs out of space during the disk creation operation.

Best practices

You should follow these best practices:

- Define dedicated network ports, bandwidth, and vSwitch configurations for the ONTAP Select networks and external storage
- Configure the capacity option to restrict storage utilization (ONTAP Select cannot consume the entire capacity of an external storage pool)
- Verify that all external storage arrays use the available redundancy and high availability (HA) features where possible

ONTAP Select networking requirements

You must configure the hypervisor network correctly before deploying ONTAP Select.

Virtual switch options

You must configure a virtual switch on each of the ONTAP Select hosts to support the external network and internal network (multi-node clusters only). As part of deploying a multi-node cluster, you should test the network connectivity on the internal cluster network.

> To learn more about how to configure a vSwitch on a hypervisor host and the high speed interface feature, refer to the Deep dive networking section.

Upgrade to VMXNET3 (ESXi only)

Beginning with ONTAP Select 9.5 using Deploy 2.10, VMXNET3 is the default network driver included with new cluster deployments on VMware ESXi. If you upgrade an older ONTAP Select node to version 9.5 or later, the

driver is not automatically upgraded.

Cluster MTU

A separate internal network is used to connect the ONTAP Select nodes in a multi-node cluster. Typically the MTU size for this network is 9000. However, there are situations where this MTU size is too large for the network connecting the ONTAP Select nodes. To accommodate the smaller frames, the MTU size used by ONTAP Select on the internal network can be in the range of 7500-9000 bytes.

The MTU size is displayed in the Cluster Details section of the cluster creation page. The value is determined by the Deploy administration utility as follows:

1. Initial default of 9000.
2. As you add the hosts and networks for the HA pairs, the MTU value is reduced as needed, based on the configuration of the vSwitches in the network.
3. The final cluster MTU value for the cluster is set after you have added all the HA pairs and are ready to create the cluster.



You can manually set the cluster MTU value if needed, based on the design of your network.

Two-NIC host with standard vSwitch (ESXi only)

In order to improve ONTAP Select performance in a two-NIC configuration, you should isolate the internal and external network traffic using two port groups. This recommendation applies to the following specific configuration:

- ONTAP Select multi-node cluster
- Two NICs (NIC1 and NIC2)
- Standard vSwitch

In this environment, you should configure the traffic using two port groups as follows:

Port group 1

- Internal network (cluster, RSM, HA-IC traffic)
- NIC1 is active
- NIC2 in standby

Port group 2

- External network (data and management traffic)
- NIC1 is standby
- NIC2 in active

Refer to the Deep dive networking section for more information about two-NIC deployments.

Four-NIC host with standard vSwitch (ESXi only)

In order to improve ONTAP Select performance in a four-NIC configuration, you should isolate the internal and external network traffic using four port groups. This recommendation applies to the following specific configuration:

- ONTAP Select multi-node cluster
- Four NICs (NIC1, NIC2, NIC3, and NIC4)
- Standard vSwitch

In this environment, you should configure the traffic using four port groups as follows:

Port group 1

- Internal network (cluster, RSM traffic)
- NIC1 is active
- NIC2, NIC3, NIC4 in standby

Port group 2

- Internal network (cluster, HA-IC traffic)
- NIC3 is active
- NIC1, NIC2, NIC4 in standby

Port group 3

- External network (data and management traffic)
- NIC2 is active
- NIC1, NIC3, NIC4 in standby

Port group 4

- External network (data traffic)
- NIC4 is active
- NIC1, NIC2, NIC3 in standby

Refer to the Deep dive networking section for more information about four-NIC deployments.

Network traffic requirements

You must make sure that your firewalls are configured properly to allow the network traffic to flow among the various participants in an ONTAP Select deployment environment.

Participants

There are several participants or entities that exchange network traffic as part of an ONTAP Select deployment. These are introduced, and then used in the summary description of the network traffic requirements.

- Deploy
  ONTAP Select Deploy administration utility
- vSphere (ESXi only)
  Either a vSphere server or ESXi host, depending on how the host is managed in your cluster deployment
- Hypervisor server
  ESXi hypervisor host or Linux KVM host
- OTS node
  An ONTAP Select node

- OTS cluster
  An ONTAP Select cluster
- Admin WS
  Local administrative workstation

Summary of network traffic requirements

The following table describes the network traffic requirements for an ONTAP Select deployment.

Protocol / Port	ESXi / KVM	Direction	Description
TLS (443)	ESXi	Deploy to vCenter server (managed) or ESXi (managed or unmanaged)	VMware VIX API
902	ESXi	Deploy to vCenter server (managed) or ESXi (unmanaged)	VMware VIX API
ICMP	ESXi or KVM	Deploy to hypervisor server	Ping
ICMP	ESXi or KVM	Deploy to each OTS node	Ping
SSH (22)	ESXi or KVM	Admin WS to each OTS node	Administration
SSH (22)	KVM	Deploy to hypervisor server nodes	Access hypervisor server
TLS (443)	ESXi or KVM	Deploy to OTS nodes and clusters	Access ONTAP
TLS (443)	ESXi or KVM	Each OTS node to Deploy	Access Deploy (Capacity Pools licensing)
iSCSI (3260)	ESXi or KVM	Each OTS node to Deploy	Mediator/Mailbox disk

ONTAP Select two-node clusters with HA

Deploying a two-node cluster with HA involves the same planning and configuration used with other cluster node configurations. However, there are several differences you should be aware of when creating a two-node cluster.

Target environment

The two-node cluster consists of one HA pair and has been specifically designed for remote office and branch office deployments.

While designed primarily for the remote and branch office environment, you can also deploy a two-node cluster in the data center if needed.

Licensing

You can deploy a two-node cluster using any VMware vSphere license. However, the VMware ROBO Standard and Advanced licenses are ideal for remote and branch office deployments.

Mediator service

When a cluster consists of two nodes, it is not possible to achieve the quorum required when a node fails or loses communication. To resolve these types of split-brain situations, every instance of the ONTAP Select Deploy utility includes a mediator service. This service connects to each node in the active two-node clusters to monitor the HA pairs and assist in managing failures. The mediator service maintains the HA state information at a dedicated iSCSI target associated with each two-node cluster.

If you have one or more active two-node clusters, the ONTAP Select Deploy virtual machine administering the clusters must be running at all times. If the ONTAP Select Deploy virtual machine is halted or fails, the mediator service is unavailable and HA capability is lost for the two-node clusters.

Location of the cluster and mediator service

Because the two-node clusters are typically deployed in a remote or branch office, they can be remote from the corporate data center and the ONTAP Select Deploy utility providing administrative support. With this configuration, the management traffic between the ONTAP Select Deploy utility and cluster flows over the WAN. Refer to the release notes for more information about limitations and restrictions.

Back up the Deploy configuration data

It is a best practice to back up the ONTAP Select Deploy configuration data on a regular basis, including after creating a cluster. This becomes particularly important with two-node clusters, because of the mediator configuration data included with the backup.

Static IP address assigned to Deploy

You must assign a static IP address to the ONTAP Select Deploy administration utility. This requirement applies to all ONTAP Select Deploy instances that manage one or more ONTAP Select two-node clusters.

ONTAP Select remote and branch office deployments

You can deploy ONTAP Select in a remote office/branch office (ROBO) environment. As part of planning a ROBO deployment, you must select the configuration supporting your objectives.

There are two primary configurations available when deploying ONTAP Select in a ROBO environment.

You can use any VMware vSphere license when deploying ONTAP Select.

ONTAP Select two-node cluster with ONTAP HA

The ONTAP Select two-node cluster consists of one HA pair and is ideal for ROBO deployments.

ONTAP Select single-node cluster with VMware support

You can deploy an ONTAP Select single-node cluster in a ROBO environment. While a single node lacks native HA capability, you can deploy the cluster in one of the following ways to provide storage protection:

- Shared external storage using VMware HA
- VMware vSAN

If you use vSAN, you must have a VMware vSAN ROBO license.

ONTAP Select MetroCluster SDS deployment requirements

MetroCluster SDS is a configuration option when creating a two-node ONTAP Select cluster. It is similar to a Remote Office/Branch Office (ROBO) deployment. This enhanced two-node deployment provides additional use case scenarios. You should be aware of the requirements and restrictions as part of preparing to deploy MetroCluster SDS.

Before deploying MetroCluster SDS, verify that the following requirements are met.

Licensing

Each node must have a premium or higher ONTAP Select license.

Hypervisor platforms

MetroCluster SDS can be deployed on the same VMware ESXi and KVM hypervisors as supported for a two-node cluster in a ROBO environment.

> Beginning with ONTAP Select 9.14.1, support for KVM hypervisor has been reinstated. Previously, support for deploying a new cluster on a KVM hypervisor was removed in ONTAP Select 9.10.1 and support for managing existing KVM clusters and hosts, except to take offline or delete, was removed in ONTAP Select 9.11.1.

Network configuration

Layer 2 connectivity is required between the participating sites. Both 10GbE and 1GbE are supported, including the following configurations:

- 1 x 10GbE
- 4 x 1GbE

> The data serving ports and interconnect ports must be connected to the same first switch.

Latency between the nodes

The network between the two nodes must support a mean latency of 5 ms with an additional 5 ms periodic jitter. Before deploying the cluster, you must test the network using the procedure described in the Deep dive networking section.

Mediator service

As with all two-node ONTAP Select clusters, there is a separate mediator service contained in the Deploy virtual machine that monitors the nodes and assists in managing failures. With the enhanced distance available with MetroCluster SDS, this creates three distinct sites in the network topology. Latency on the link between the mediator and a node should be 125 ms round-trip or less.

Storage

Direct-attached storage (DAS) is supported using either HDD and SSD disks. vNAS is also supported, including external storage arrays and vSAN in a VMware environment.

> When deploying MetroCluster SDS, you cannot use vSAN in a distributed or "stretched" topology.

Static IP address assigned to Deploy

You must assign a static IP address to the Deploy administration utility. This requirement applies to all Deploy instances that manage one or more ONTAP Select two-node clusters.

ONTAP Select VMware vCenter server on ESXi

You must define a vCenter server account and associate it with a role containing the necessary administrative privileges.

You also need the fully qualified domain name or IP address of the vCenter server managing the ESXi hypervisor hosts where ONTAP Select is deployed.

Administrative privileges

The minimum administrative privileges needed to create and manage an ONTAP Select cluster are presented below.

Datastore

- Allocate space
- Browse datastore
- Low level file operations
- Update virtual machine files
- Update virtual machine metadata

Host

Configuration

- Network configuration
- System management

Local operations

- Create virtual machine
- Delete virtual machine
- Reconfigure virtual machine

Network

- Assign network

Virtual machine

Configuration

All privileges in the category.

Interaction

All privileges in the category.

Inventory

All privileges in the category.

Provisioning

All privileges in the category.

vApp

All privileges in the category.

Related information

Learn about VMware vSphere privileges for vSAN ESA in vCenter

ONTAP Select Deploy

ONTAP Select Deploy general requirements and planning

There are several general requirements that you should consider as part of planning to install the ONTAP Select Deploy administration utility.

Pairing the Deploy utility with the ONTAP Select clusters

You have several options when pairing an instance of the Deploy utility with the ONTAP Select clusters.

> In all deployment scenarios, a single ONTAP Select cluster and the nodes in the cluster can be managed by only one instance of the Deploy administration utility. A cluster cannot be managed by two or more different instances of the Deploy utility.

One instance of the utility for each ONTAP Select cluster

You can deploy and manage each ONTAP Select cluster using a dedicated instance of the Deploy utility. With this one-to-one configuration, there is a clear separation between each of the utility-to-cluster pairings. This configuration provides a high level of isolation with smaller failure domains.

One instance of the utility for multiple ONTAP Select clusters

You can deploy and manage multiple ONTAP Select clusters in your organization using a single instance of the Deploy utility. With this one-to-many configuration, all processing and configuration data is managed by the same instance of the Deploy utility.

> One instance of the Deploy utility can administer up to 400 ONTAP Select nodes or 100 clusters.

Requirements related to the KVM environment

Before installing the Deploy administration utility in a KVM hypervisor environment, you should review the basic requirements and prepare for the deployment.

Requirements and restrictions for a deployment

There are several requirements and restrictions that you should consider when installing the ONTAP Select Deploy utility in a KVM environment.

Linux KVM host server hardware requirements

There are several minimum resource requirements that your Linux KVM hypervisor host must meet. Verify that the hosts where ONTAP Select is deployed meet the following basic requirements:

- Linux server:
  - The hardware and software must be 64-bit
  - The server must adhere to the same supported versions as defined for an ONTAP Select node
- Virtual CPUs (2)
- Virtual memory (4GB)
- Storage (40GB)

- "Dynamic Host Configuration Protocol (DHCP) is enabled (you can also assign a static IP address)

Network connectivity

Verify that the Deploy virtual machine network interface is configured and can to connect to the ONTAP Select hosts that it manages.

Support for IP version 4

ONTAP Select Deploy only supports IP version 4 (IPv4). IP version 6 (IPv6) is not supported. This restriction affects ONTAP Select in the following ways:

- You must assign an IPv4 address to the management LIF of the Deploy VM.
- Deploy cannot create ONTAP Select nodes configured to use IPv6 on the ONTAP LIFs.

Required configuration information

As part of your deployment planning, you should determine the required configuration information before installing the ONTAP Select Deploy administration utility.

Name of the Deploy VM

The name to use for the VM.

Name of the Linux KVM host

The Linux KVM host where the Deploy utility is installed.

Name of the storage pool

The storage pool holding the VM files (approximately 40GB is required).

Network for the VM

The network where the Deploy VM is connected.

Optional network configuration information

The Deploy VM is configured using the DHCP by default. However, if needed, you can manually configure the network interface for the VM.

Host name

The name of the host.

Host IP address

The static IPv4 address.

Subnet mask

The subnetwork mask, which is based on the network that the VM is a part of.

Gateway

The default gateway or router.

Primary DNS server

The primary domain name server.

Secondary DNS server

The secondary domain name server.

Search domains

The search domains to use.

Authentication using the credential store

The ONTAP Select Deploy credential store is a data base holding account information. Deploy uses the account credentials to perform host authentication as part of cluster creation and management. You should be aware of how the credential store is used as part of planning an ONTAP Select deployment.

The account information is stored securely in the database using the AES encryption algorithm and SHA-256 hashing algorithm.

Types of credentials

The following types of credentials are supported:

- **Host**: Used to authenticate a hypervisor host as part of deploying an ONTAP Select node directly to VMware ESXi
- **vCenter**: Used to authenticate a vCenter server as part of deploying an ONTAP Select node to ESXi when the host is managed by VMware vCenter

Access

The credential store is accessed internally as part of performing normal administrative tasks using Deploy, such as adding a hypervisor host. You can also manage the credential store directly through the Deploy web user interface and CLI.

ONTAP Select Deploy hypervisor host considerations

There are several planning issues related to the hypervisor host that you should consider.

You should not directly modify the configuration of an ONTAP Select virtual machine unless directed to do so by NetApp support. A virtual machine should only be configured and modified through the Deploy administration utility. Making changes to an ONTAP Select virtual machine outside of the Deploy utility without assistance from NetApp support can cause the virtual machine to fail and render it unusable.

Hypervisor independent

Both ONTAP Select and the ONTAP Select Deploy administration utility are hypervisor independent.

The following hypervisors are supported for both ONTAP Select and the ONTAP Select Deploy administration:

- VMware ESXi
- Kernel-based Virtual Machine (KVM)

Refer to the hypervisor-specific planning information and release notes for additional details regarding the supported platforms.

Hypervisor for ONTAP Select nodes and administration utility

Both the Deploy administration utility and the ONTAP Select nodes run as virtual machines. The hypervisor you choose for the Deploy utility is independent of the hypervisor you choose for the ONTAP Select nodes. You have complete flexibility when pairing the two:

- A Deploy utility running on VMware ESXi can create and manage ONTAP Select clusters on either VMware ESXi or KVM.
- A Deploy utility running on KVM can create and manage ONTAP Select clusters on either VMware ESXi or KVM.

One or more instances of ONTAP Select node per host

Each ONTAP Select node runs as a dedicated virtual machine. You can create multiple nodes on the same hypervisor host, with the following restrictions:

- Multiple nodes from a single ONTAP Select cluster cannot run on the same host. All the nodes on a specific host must be from different ONTAP Select clusters.
- You must use external storage.
- If you use software RAID, you can only deploy one ONTAP Select node on the host.

Hypervisor consistency for the nodes within a cluster

All of the hosts within an ONTAP Select cluster must run on the same version and release of the hypervisor software.

Number of physical ports on each host

You must configure each host to use one, two, or four physical ports. Although you have flexibility when configuring the network ports, you should follow these recommendations where possible:

- A host in a single-node cluster should have two physical ports.
- Each host in a multi-node cluster should have four physical ports

Integrate ONTAP Select with an ONTAP hardware-based cluster

You cannot add an ONTAP Select node directly to an ONTAP hardware-based cluster. However, you can optionally establish a cluster peering relationship between an ONTAP Select cluster and a hardware-based ONTAP cluster.

VMware hypervisor environment

There are several requirements and restrictions specific to the VMware environment that you should consider before installing the ONTAP Select Deploy utility in a VMware environment.

ESXi host server hardware requirements

Your ESXi hypervisor host must meet several minimum resource requirements. Check that the hosts where ONTAP Select is deployed meet the following basic requirements:

- ESXi server:
  - Hardware and software must be 64-bit
  - Must adhere to the same supported versions as defined for an ONTAP Select node
- Virtual CPUs (2)
- Virtual memory (4 GB)
- Storage (40 GB)
- DHCP enabled (can also assign a static IP address)

Network connectivity

Check that the ONTAP Select Deploy virtual machine network interface is configured and has a single

management IP address. You can use DHCP to dynamically assign an IP address or manually configure a static IP address.

Depending on your deployment decisions, the Deploy VM must be able to connect to the vCenter server, ESXi hypervisor hosts, and ONTAP Select nodes it manages. You must configure your firewalls to allow the required traffic.

Deploy uses the VMware VIX API to communicate with the vCenter server and ESXi hosts. Initially, it establishes a connection using SOAP over SSL on TCP port 443. After this, a connection is opened using SSL on port 902. In addition, Deploy issues PING commands to verify that there is an ESXi host at the IP address you specify.

Additionally, Deploy must be able to communicate with the ONTAP Select node and cluster management IP addresses using the following protocols:

- PING command (ICMP)
- SSH (port 22)
- SSL (port 443)

Support for IP version 4

ONTAP Select Deploy only supports IP version 4 (IPv4). IP version 6 (IPv6) is not supported. This restriction affects ONTAP Select in the following ways:

- You must assign an IPv4 address to the management LIF of the Deploy virtual machine.
- Deploy cannot create ONTAP Select nodes configured to use IPv6 on the ONTAP LIFs.

VMware vCenter server considerations

VMware vSphere privileges for ESA

To create and manage vSAN Express Storage Architecture (ESA) in vCenter, you need the following vSphere Content Library privileges:

> To access the Content Library privileges, an administrator needs to grant you access to Global Permissions in vCenter Server.

- Add library item
- Create local library
- Delete library item
- Delete local library
- Download files
- Read storage
- Update files
- Update library
- Update library item
- Update local library
- View configuration settings

Summary of best practices for ONTAP Select deployment

There are best practices that you should consider as part of planning an ONTAP Select deployment.

Storage

You should consider the following best practices for storage.

All-Flash or Generic Flash arrays

ONTAP Select virtual NAS (vNAS) deployments using all-flash VSAN or generic flash arrays should follow the best practices for ONTAP Select with non-SSD DAS storage.

External storage

You should adhere to the following recommendations:

- Define dedicated network ports, bandwidth, and vSwitch configurations for the ONTAP Select networks and external storage
- Configure the capacity option to restrict storage utilization (ONTAP Select cannot consume the entire capacity of an external storage pool)
- Verify that all external storage arrays use the available redundancy and HA features where possible

Hypervisor core hardware

All of the drives in a single ONTAP Select aggregate should be the same type. For example, you should not mix HDD and SSD drives in the same aggregate.

RAID controller

The server RAID controller should be configured to operate in writeback mode. If write workload performance issues are seen, check the controller settings and make sure that writethrough or writearound is not enabled.

If the physical server contains a single RAID controller managing all locally attached disks, NetApp recommends creating a separate LUN for the server OS and one or more LUNs for ONTAP Select. In the event of boot disk corruption, this best practice allows the administrator to recreate the OS LUN without affecting ONTAP Select.

The RAID controller cache is used to store all incoming block changes, not just those targeted toward the NVRAM partition. Therefore, when choosing a RAID controller, select one with the largest cache available. A larger cache allows less frequent disk flushing and an increase in performance for the ONTAP Select VM, the hypervisor, and any compute VMs collocated on the server.

RAID groups

The optimal RAID-group size is eight to 12 drives. The maximum number of drives per RAID group is 24.

The maximum number of NVME drives supported per ONTAP Select node is 14.

A spare disk is optional, but recommended. NetApp also recommends using one spare per RAID group; however, global spares for all RAID groups can be used. For example, you can use two spares for every three RAID groups, with each RAID group consisting of eight to 12 drives.

ONTAP Select receives no performance benefits by increasing the number of LUNs within a RAID group. Multiple LUNs should only be used to follow best practices for SATA/NL-SAS configurations or to bypass hypervisor file system limitations.

VMware ESXi hosts

NetApp recommends using ESXi 8.0 or later and an NVMe disk for the datastore hosting the system disks. This configuration provides the best performance for the NVRAM partition.

> When installing on ESXi 8.0 or later, ONTAP Select uses the vNVMEe driver regardless of whether the system disk resides on an SSD or on an NVMe disk. This sets the VM hardware level to 13, which is compatible with ESXi 8.0 and later.

Define dedicated network ports, bandwidth, and vSwitch configurations for the ONTAP Select networks and external storage (VMware vSAN and generic storage array traffic when using iSCSI or NFS).

Configure the capacity option to restrict storage utilization (ONTAP Select cannot consume the entire capacity of an external vNAS datastore).

Assure that all generic external storage arrays use the available redundancy and HA features where possible.

VMware Storage vMotion

Available capacity on a new host is not the only factor when deciding whether to use VMware Storage vMotion with an ONTAP Select node. The underlying storage type, host configuration, and network capabilities should be able to sustain the same workload as the original host.

Networking

You should consider the following best practices for networking.

Duplicate MAC addresses

To eliminate the possibility of having multiple Deploy instances assign duplicate MAC addresses, one Deploy instance per layer-2 network should be used to create or manage an ONTAP Select cluster or node.

EMS messages

The ONTAP Select two-node cluster should be carefully monitored for EMS messages indicating that storage failover is disabled. These messages indicate a loss of connectivity to the mediator service and should be rectified immediately.

Latency between nodes

The network between the two nodes must support a mean latency of 5 ms with an additional 5 ms periodic jitter. Before deploying the cluster, test the network using the procedure described in the ONTAP Select Product Architecture and Best Practices technical report.

Load balancing

To optimize load balancing across both the internal and the external ONTAP Select networks, use the Route Based on Originating Virtual Port load-balancing policy.

Multiple layer-2 networks

If data traffic spans multiple layer-2 networks and the use of VLAN ports is required or when you are using multiple IPspaces, VGT should be used.

Physical switch configuration

VMware recommends that STP be set to Portfast on the switch ports connected to the ESXi hosts. Not setting STP to Portfast on the switch ports can affect the ONTAP Select ability to tolerate uplink failures. When using LACP, the LACP timer should be set to fast (1 second). The load-balancing policy should be set to Route Based on IP Hash on the port group and Source and Destination IP Address and TCP/UDP port and VLAN on the LAG.

Virtual switch options for KVM

You must configure a virtual switch on each of the ONTAP Select hosts to support the external network and internal network (multi-node clusters only). As part of deploying a multi-node cluster, you should test the network connectivity on the internal cluster network.

To learn more about how to configure an Open vSwitch on a hypervisor host, refer to the ONTAP Select on KVM Product Architecture and Best Practices technical report.

HA

You should consider the following best practices for high availability.

Deploy backups

It is a best practice to regularly back up the Deploy configuration data, including after creating a cluster. This becomes particularly important with two-node clusters, because the mediator configuration data is included with the backup.

After creating or deploying a cluster, you should back up the ONTAP Select Deploy configuration data.

Mirrored aggregates

Although the existence of the mirrored aggregate is needed to provide an up-to-date (RPO 0) copy of the primary aggregate, take care that the primary aggregate does not run low on free space. A low-space condition in the primary aggregate might cause ONTAP to delete the common Snapshot copy used as the baseline for storage giveback. This works as designed to accommodate client writes. However, the lack of a common Snapshot copy on failback requires the ONTAP Select node to do a full baseline from the mirrored aggregate. This operation can take a significant amount of time in a shared-nothing environment.

NetApp recommends that you maintain at least 20% free space for mirrored aggregates for optimal storage performance and availability. Although the recommendation is 10% for non-mirrored aggregates, the filesystem can use the additional 10% of space to absorb incremental changes. Incremental changes increase space utilization for mirrored aggregates due to ONTAP's copy-on-write Snapshot-based architecture. Failure to adhere to these best practices might have a negative impact on performance. High availability takeover is only supported when data aggregates are configured as mirrored aggregates.

NIC aggregation, teaming, and failover

ONTAP Select supports a single 10Gb link for two-node clusters; however, it is a NetApp best practice to have hardware redundancy through NIC aggregation or NIC teaming on both the internal and the external networks of the ONTAP Select cluster.

If a NIC has multiple application-specific integrated circuits (ASICs), select one network port from each ASIC when building network constructs through NIC teaming for the internal and external networks.

NetApp recommends that the LACP mode be active on both the ESXi and the physical switches. Furthermore, the LACP timer should be set to fast (1 second) on the physical switch, ports, port channel interfaces, and on the VMNICs.

When using a distributed vSwitch with LACP, NetApp recommends that you configure the load-balancing policy to Route Based on IP Hash on the port group, Source and Destination IP Address, TCP/UDP Port, and VLAN on the LAG.

Two-node stretched HA (MetroCluster SDS) best practices

Before you create a MetroCluster SDS, use the ONTAP Deploy connectivity checker to make sure that the network latency between the two data centers falls within the acceptable range.

There is an extra caveat when using virtual guest tagging (VGT) and two-node clusters. In two-node cluster configurations, the node management IP address is used to establish early connectivity to the mediator before ONTAP is fully available. Therefore, only external switch tagging (EST) and virtual switch tagging (VST) tagging is supported on the port group mapped to the node management LIF (port e0a). Furthermore, if both the management and the data traffic are using the same port group, only EST and VST are supported for the entire two-node cluster.

License

Options

Learn about evaluation licenses for ONTAP Select deployments

You can deploy ONTAP Select with an evaluation license or a purchased license. The license you choose must apply to each of the nodes in the ONTAP Select cluster and therefore to the entire cluster. You can use an evaluation license if you want to evaluate ONTAP Select before making the decision to purchase. The evaluation license is included with the ONTAP Select Deploy administration utility and is automatically applied to each ONTAP Select node as part of an evaluation deployment.

To download the ONTAP Select Deploy administration utility, you need the following:

- A registered NetApp Support Site account. If you don't have an account, see User registration.
- To accept the End User License agreement for an ONTAP Select deployment with an evaluation license.

There are several considerations when deploying and supporting an evaluation cluster:

- You can only use the cluster for evaluation purposes. You must not use a cluster with an evaluation license in a production environment.
- You should use the ONTAP Select Deploy administration utility as follows when configuring each host:
  - Don't supply a serial number
  - Configure to use an evaluation license

Licensing characteristics

The ONTAP Select evaluation license has the following characteristics:

- A production license with storage capacity isn't required
- The node serial number is twenty digits and automatically generated by ONTAP Select Deploy

  (you don't acquire it directly from NetApp)
- The evaluation period provided by the license can be up to 90 days
- The maximum storage allocated by each node is the same as a production license

Upgrade to a production license

You can upgrade an ONTAP Select evaluation cluster to use a production license. You should be aware of the following restrictions:

- You must use the Deploy administration utility to perform the license upgrade
- You can use a Capacity Tier license, however Capacity Pools licensing isn't supported

- Each node must have enough storage allocated to support the minimum required for a production license, based on the cluster size

See Convert an evaluation license to a production license for more information.

What's next?

- Access the ONTAP Select evaluation software
- Deploy a 90-day evaluation instance of an ONTAP Select cluster

Related information

- Learn about licenses for production deployments

Production licenses

Learn about ONTAP Select purchased licenses for production deployments

After you determine that ONTAP Select is suitable for your organization, you can purchase the licenses needed to support a production deployment. You must choose either the Capacity Pools or Capacity Tiers licensing model as well as the storage capacity for each deployment.

Common licensing characteristics

The *Capacity Pools* and *Capacity Tiers* licensing models are very different in several respects. However, the two licensing models share several common characteristics, including:

- You must purchase one or more licenses as needed when deploying ONTAP Select in a production environment.
- The storage capacity for a license is allocated in 1 TB increments.
- All platform licensing offerings are supported (standard, premium, and premium XL).
- You should contact your NetApp account team or partner for assistance as needed when acquiring the necessary licenses.
- You must upload the license files to the Deploy administration utility, which then applies the licenses based on the licensing model.
- After installing and applying a license, you can add additional capacity by contacting your NetApp account team or partner to procure an updated license.
- An ONTAP Select node that is initially deployed with a purchased license cannot be converted to an evaluation license.

Capacity Tiers licensing model

There are several characteristics unique to the Capacity Tiers licensing model, including:

- You must purchase a license for each ONTAP Select node.
- The minimum amount you can purchase is 1 TB.
- Each Capacity Tier license has a storage capacity and is locked to a specific node.
- The storage capacity identifies the raw capacity and corresponds to the total allowable size of the data disks available to the ONTAP Select virtual machine.

- NetApp generates a nine-digit license serial number for each ONTAP Select node.
- The storage allocated to a node is perpetual (no renewal required), but you must renew the support contract associated with the license.
- Each ONTAP Select node in an high-availability (HA) pair requires a license that is at least as large as the raw storage attached to the node.
- The node serial number is nine digits and equal to the license serial number.
- You can apply the license file during cluster deployment or within 30 days after creating a cluster.

Capacity Pools licensing model

There are several characteristics unique to the Capacity Pools licensing model, including:

- You must purchase a license for each shared Capacity Pool.
- The minimum amount you can purchase is 2 TB.
- Each Capacity Pool license has a storage capacity and is locked to a specific License Manager instance within a Deploy administration utility.
- The storage capacity identifies the raw capacity and corresponds to the total allowable size of the data aggregates that you can create in the ONTAP Select virtual machine.
- A nine-digit license serial number is generated by NetApp for each Capacity Pool.
- The storage allocated to a Capacity Pool is valid only for a specific time based on the purchase (renewal required).
- The License Manager generates a twenty-digit node serial number based on the Capacity Pool license serial number.
- All ONTAP Select nodes in a cluster share the same license capacity. This means the license must be large enough to provide for the storage allocated to all nodes in the cluster.
- Each node automatically leases storage capacity for its local data aggregates from a shared Capacity Pool.

Platform license offerings for ONTAP Select

You can purchase an ONTAP Select Capacity Tier or Capacity Pool license in the standard, premium, or premium XL level. These license offerings determine the capabilities of the hosts where you deploy ONTAP Select.

What a platform license offering provides

A specific license offering determines the capabilities of the hypervisor host in two areas:

- Instance type (CPU, memory)
- Additional features

The license offerings are arranged in an ascending order of capabilities from standard to premium XL. In general, the license option you choose grants you the capabilities of that level and all lower levels. For example, the premium level provides the capabilities of both premium and standard.

The following table compares the capabilities for the standard, premium, and premium XL license offerings.

License supports …		Standard	Premium	Premium XL
Instance type		Small only	Small or medium	Small, medium, or large
Hard disk drives (HDD) in a …	Hardware RAID configuration	Yes	Yes	Yes
	vNAS configuration	Yes	Yes	Yes
Solid state drives (SSDs) in a …	Hardware RAID configuration	No	Yes	Yes
	Software RAID configuration	No	Yes	Yes
	vNAS configuration	Yes	Yes	Yes
NVMe drives in a …	Hardware RAID configuration	No	No	Yes
	Software RAID configuration	No	No	Yes
	vNAS configuration	Yes	Yes	Yes
MetroCluster SDS		No	Yes	Yes

Kernel-Based Virtual Machines (KVM) do not support the large instance type.

Compare hardware support for platform license offerings

The standard, premium, and premium XL license offerings support a wide range of hardware and software. For current information on hardware and software versions, refer to the Interoperability Matrix Tool.

Core item

Core item type	Description
Host protocols	NFS, SMB/CIFS, iSCSI, and NVMe over TCP
Deployment options	Single node Two-node cluster (HA pair) Four, six, eight, ten, or twelve-node cluster
Supported capacity (per node)	Up to 400TB raw data (ESXi and KVM)

Hardware

Hardware type	Description		
Instance size	Small	Medium	Large

CPU family	Intel Xeon E5-26xx v3 (Haswell) or later AMD EPYC Embedded 9004 or later and 8004	Intel Xeon E5-26xx v3 (Haswell) or later AMD EPYC Embedded 9004 or later and 8004	Intel Xeon E5-26xx v3 (Haswell) or later AMD EPYC Embedded 9004 or later and 8004
ONTAP Select CPU / memory	4 virtual CPUs (vCPUs) / 16GB of RAM	8 vCPUs / 64GB of RAM	16 vCPUs / 128GB of RAM
Host CPU / memory minimum requirements [1]	6 cores / 24GB of RAM	10 cores / 72GB of RAM	18 cores / 136GB of RAM
Network (per node)	Minimum of two 1GbE ports for a single node cluster Minimum of four 1GbE ports for a two-node cluster (HA pair) Minimum of two 10GbE ports for a four, six, eight, ten, or twelve-node cluster		

[1] Assumes two cores and 8GB of RAM for the hypervisor.

Storage type

The following table provides the minimum license type required for the specified storage.

Storage type	**Description**		
License type	**Standard**	**Premium**	**Premium XL**
Instance size	**Small**	**Small and medium**	**Small, medium, and large**
Local DAS with hardware RAID controller	8 - 60 drives	8 - 60 drives	8 - 60 drives
HDD (SAS, NL-SAS, SATA)	Not applicable	4 – 60 drives	4 – 60 drives
SSD (SAS)	Not applicable	4 – 60 drives (SSD only)	4 – 60 drives (SSD only)
Local DAS with software RAID	Not applicable	Applicable	4 - 14 drives (NVMe only)
External arrays [1]	Datastores hosted on external arrays are connected through FC, FCoE, iSCSI, and NFS (NFS isn't supported for KVM). These datastores provide high availability and resilience.		

[1] External array protocol support reflects networked storage connectivity.

Software

Software type	**Description**

Hypervisor support (VMware)	VMware vSphere 9.0 VMware vSphere 8.0GA and updates 1 through 3 VMware vSphere 7.0GA and updates 1 through 3C
Hypervisor support (KVM)	Red Hat Enterprise Linux 64-bit (KVM) 10.1, 10.0, 9.7, 9.6, 9.5, 9.4, 9.3, 9.2, 9.1, 9.0, 8.8, 8.7, and 8.6 Rocky Linux (KVM) 10.1, 10.0, 9.7, 9.6, 9.5, 9.4, 9.3, 9.2, 9.1, 9.0, 8.9, 8.8, 8.7, and 8.6
Management software	NetApp Active IQ Unified Manager management suite ONTAP Select Deploy Utility SnapCenter (optional)

What's next?

Learn about purchasing an ONTAP Select license

Capacity Pools licensing model

Operational details for ONTAP Select Capacity Pools licensing model

The Capacity Pools licensing model is different from the Capacity Tiers model. Instead of a dedicating storage capacity to each individual node, the storage capacity is allocated to a pool and shared among many nodes. Additional components and processes have been created to support the Capacity Pools model.

License Manager

The License Manager runs as a separate process within each instance of the Deploy administration utility. Some of the functions provided by LM include:

- Generate a unique twenty-digit serial number for each node based on the Capacity Pool license serial number
- Create leases for capacity from the shared Capacity Pools based on requests from the ONTAP Select nodes
- Report pool usage information through the Deploy user interface

Lease characteristics

The storage allocated for every data aggregate at a node using a Capacity Pool license must have an associated lease. The node requests a storage lease and if the capacity is available, the License Manager responds with a lease. Each lease has the following explicit or implicit attributes:

- License Manager

  Every ONTAP Select node is associated with one License Manager instance

- Capacity Pool

Every ONTAP Select node is associated with one Capacity Pool

- Storage allocation

  A specific capacity value is assigned in the lease

- Expiration date and time

  Leases have a duration of between one hour and seven days depending on the user configuration.

License Lock ID

Each instance of the License Manager, and therefore each corresponding Deploy utility instance, is identified with a unique 128-bit number. This number is combined with the nine-digit Capacity Pool license serial number to lock the pool to a specific License Manager instance (which is effectively a Deploy instance). You must provide both values at the NetApp support site as part of generating the NetApp License File (NLF).

You can determine the License Lock ID for your Deploy instance using the web user interface in the following ways:

- Getting Started page

  Deploy displays this page when you first sign in. You can also open the drop down box at the top right of the page and select **Getting Started** to display this page. The LLID displays in the Add Licenses section.

- Select the **Administration** tab at the top of the page, then select **Systems & AutoSupport**.

Basic lease operations

An ONTAP Select node must locate or request a valid capacity lease every time a data aggregate is created, expanded, or changed. You can use a lease obtained from a previous request that is still valid, or you can request a new lease if needed. The following steps are taken by the ONTAP Select node to locate a Capacity Pool lease:

1. If an existing lease is located at the node, it is used as long as all of the following are true:
   - Lease has not expired
   - Storage request for the aggregate does not exceed the lease capacity
2. If an existing lease cannot be located, the node requests a new lease from the License Manager.

Return storage capacity to a Capacity Pool

Storage capacity is allocated from a Capacity Pool as needed and each new request can reduce the available storage in the pool. Storage capacity is returned to the pool in several situations, including:

- Lease for a data aggregate expires and is not renewed by the node
- Data aggregate is deleted



If an ONTAP Select virtual machine is deleted, any active leases remain in effect until they expire. When this occurs, the capacity is returned to the pool.

Node serial numbers for ONTAP Select Capacity Pools licensing model

With the Capacity Tiers licensing model, the nine-digit node serial number is the same as the license serial number assigned to the node. However, the serial numbers assigned to nodes using the Capacity Pools licensing model have a different format.

The serial number of a node using Capacity Pools licensing has the following format:

```
999 pppppppppp nnnnnnnn
```

Spaces have been added for clarity, but are not part of the actual serial number.

The following table describes each section of the node serial number from left to right.

Section	Description
'999'	Constant three-digit value reserved by NetApp.
pppppppppp	Variable nine-digit license serial number assigned to the Capacity Pool by NetApp
nnnnnnnn	Variable eight-digit value generated by the License Manager for each node using the Capacity Pool

When opening a case with NetApp support involving a node that uses a Capacity Pool license, you must provide the nine-digit Capacity Pool license serial number. You cannot provide the full twenty-digit node serial number. You can derive the license serial number from the node serial number by skipping the first three digits of the node serial number ('999') and extracting the next nine digits (pppppppppp).

Deployment restrictions for ONTAP Select Capacity Pools licensing

The restrictions that apply when using the Capacity Pool licensing model are presented below.

Consistent licensing model per cluster

All of the nodes within a single ONTAP Select cluster must use the same licensing model, either Capacity Tiers or Capacity Pools. You cannot mix the licensing types for the nodes within a single cluster.

All nodes in a cluster use the same License Manager instance

All the nodes with a Capacity Pool license in an ONTAP Select cluster must use the same License Manager instance. Because there is one instance of License Manager within each Deploy instance, this restriction is a restatement of the existing requirement that the same Deploy instance must manage all nodes in a cluster.

One Capacity Pool per node

Each node can lease storage from exactly one Capacity Pool. A node cannot use two or more pools.

Same pool for nodes in an HA pair

Both nodes in a single HA pair must lease storage from the same Capacity Pool. However, different HA pairs within the same cluster can lease storage from different pools managed by the same License Manager.

Storage license duration

You must choose a license duration when acquiring the storage license from NetApp. For example, a license

could be valid for one year.

Data aggregate lease duration

When an ONTAP Select node requests a storage lease for a data aggregate, the License Manager provides a lease for a specific duration based on the configuration of the Capacity Pool. You can configure the lease duration for each pool between one hour and seven days. The default lease duration is 24 hours.

Static IP address assigned to Deploy

You must assign a static IP address to the Deploy administration utility when Capacity Pools licensing is used.

Summary of benefits for ONTAP Select Capacity Pools licensing

There are several benefits when using the Capacity Pools licensing model instead of the Capacity Tiers licensing model.

More efficient use of storage capacity

When using Capacity Tiers licensing, you allocate a fixed storage capacity to each node. Any unused space cannot be shared with the other nodes and is effectively wasted. With Capacity Pools licensing, each node only consumes the capacity it needs, based on the size of the data aggregates.

Capacity is anchored in a central pool where it can be shared among many nodes in your organization.

Significantly reduced administrative overhead resulting in lower cost

If you use Capacity Tier licenses, you must obtain and install a license for each node. When using Capacity Pools, there is one license for each shared pool. This can dramatically reduce the administrative overhead and result in lower cost.

Improved usage metrics

The Deploy web user interface provides enhanced usage information for the Capacity Pools. You can quickly determine how much storage is used and available in a Capacity Pool, which nodes are using storage from a pool, and what pools a cluster is allocating capacity from.

ONTAP Select Capacity Tiers and Capacity Pools licensing comparison

The following table compares the two production licensing models supported by ONTAP Select.

	Capacity Tiers	Capacity Pools
License serial number	NetApp generates nine digits and assigns them to a node	NetApp generates nine digits and assigns them to a Capacity Pool
License lock	Locked to ONTAP Select node	Locked to License Manager instance
License duration	Perpetual (no renewal required)	Fixed duration based on purchase (renewal required)
Lease duration for data aggregate	Not applicable	One hour to seven days
Node serial number	Nine digits and equal to license serial number	Twenty digits and generated by License Manager

	Capacity Tiers	Capacity Pools
Support	Add-on and time-limited	Included and co-termed
License types	Standard, premium, premium XL	Standard, premium, premium XL
Evaluation license available	Yes	Yes
Evaluation to production upgrade	Yes	No
ONTAP Select virtual machine resize (small to medium, medium to large)	Yes	Yes
Enforcement: license expired	N/A	Yes (no grace period)
Minimum license	1 TB	2 TB
Maximum managed space	400 TB per node	400 TB per node

Purchase a license

License purchasing workflow for ONTAP Select

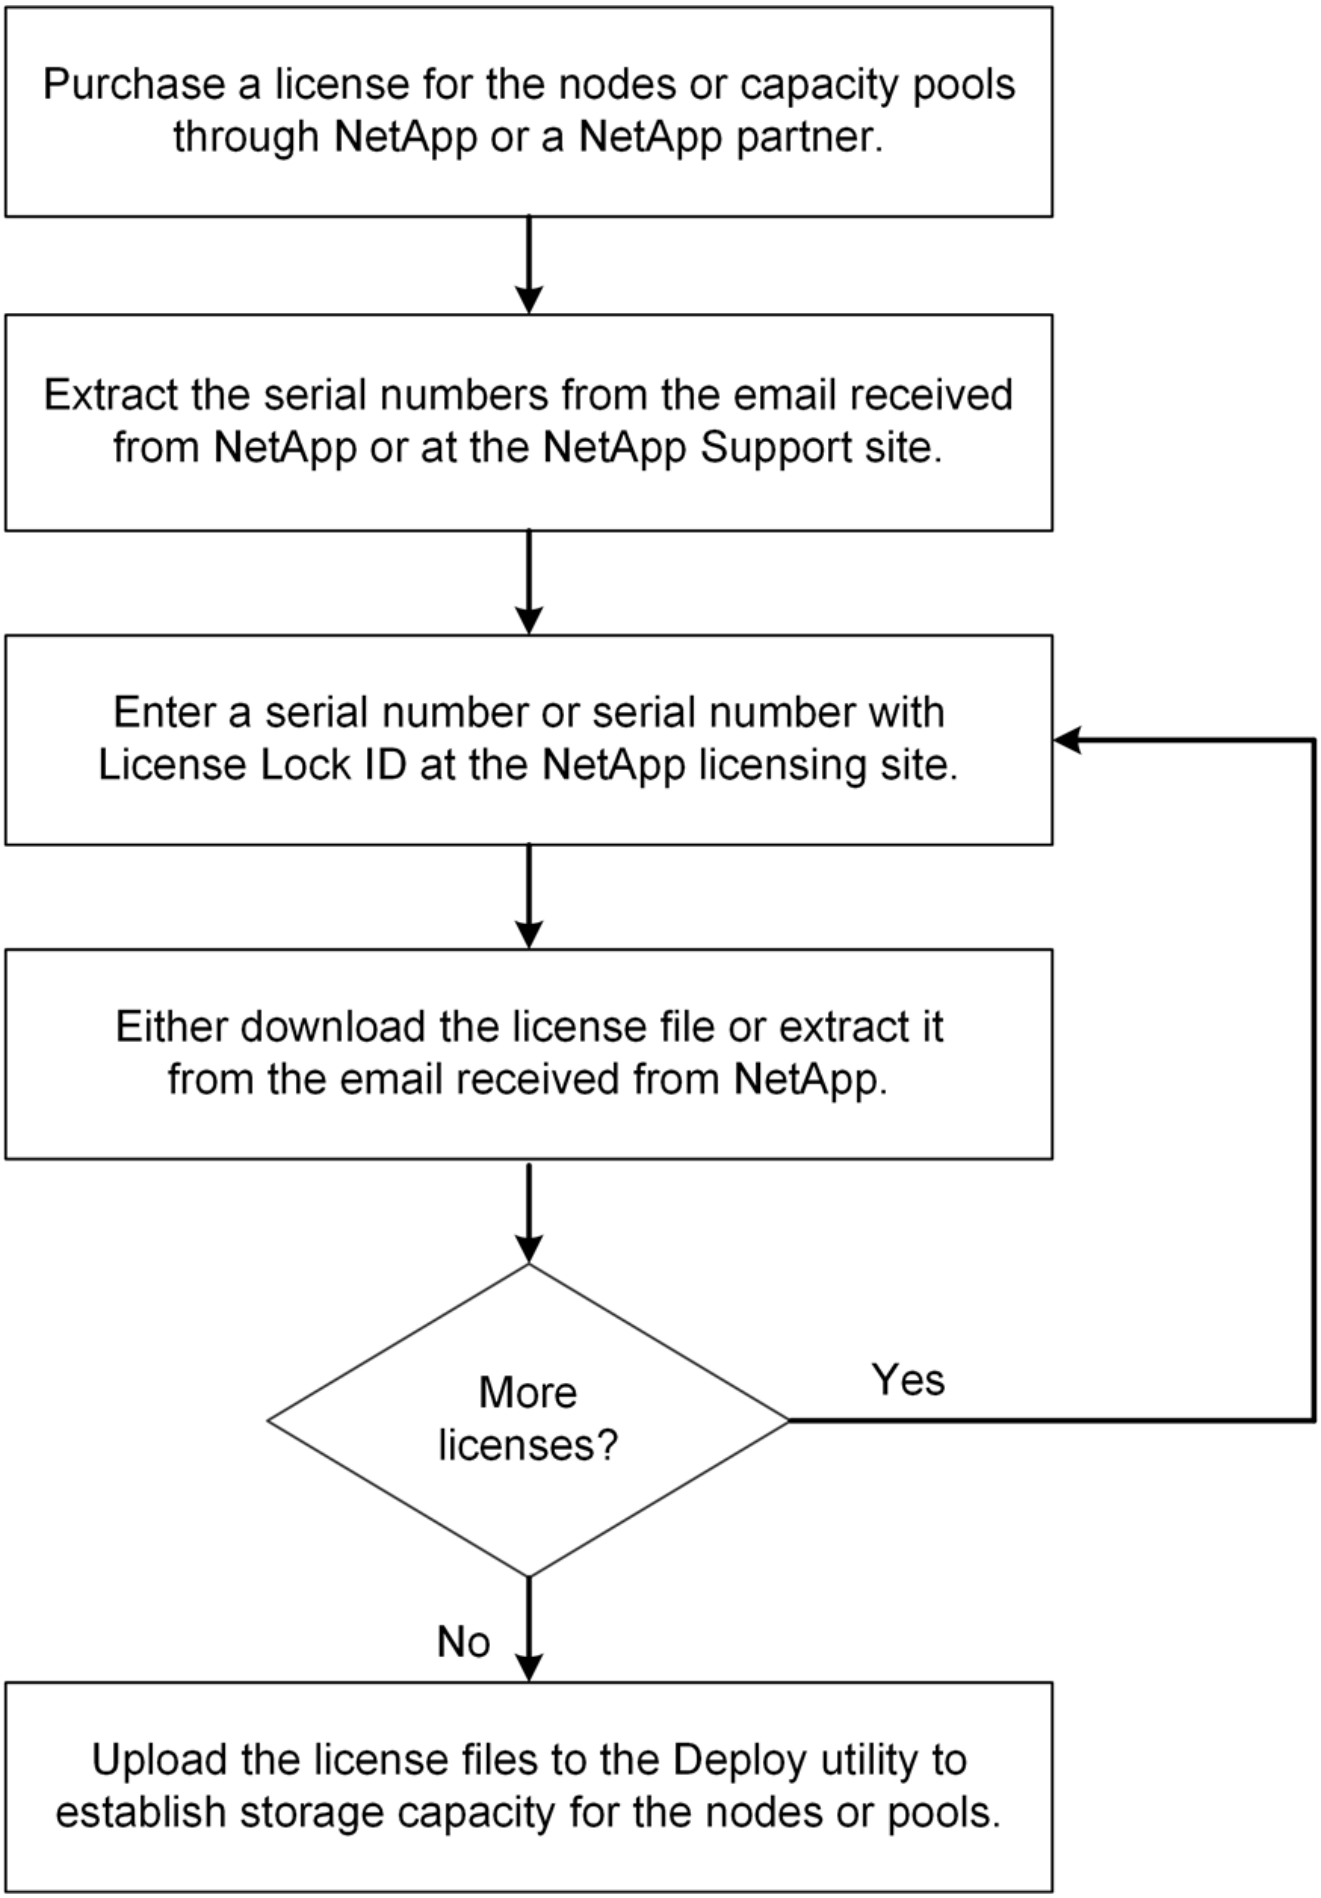
The following workflow illustrates the process of purchasing and applying a license for your ONTAP Select deployment. When purchasing a license, you must select the licensing model and storage capacity.

The exact process varies based on whether you are using a Capacity Tier or Capacity Pool license:

Nine-digit license serial number
: Serial number applies to either a node (Capacity Tiers) or a storage pool (Capacity Pools)

License Lock ID
: You must have the License Lock ID for your Deploy instance when using a Capacity Pool license

Licensing web site
: You obtain a Capacity Tier and Capacity Pool license at different web sites.


Purchase a license for the nodes or capacity pools through NetApp or a NetApp partner.
Extract the serial numbers from the email received from NetApp or at the NetApp Support site.
Enter a serial number or serial number with License Lock ID at the NetApp licensing site.
Either download the license file or extract it from the email received from NetApp.
More licenses?
Yes
No
Upload the license files to the Deploy utility to establish storage capacity for the nodes or pools.

What's next?

Learn about the considerations for sizing an ONTAP Select license:

- Capacity Tiers license
- Capacity Pools license

Related information

Learn about managing ONTAP Select licenses

Capacity Tiers

Determine Capacity Tiers license size for ONTAP Select

The Capacity Tiers license is based on the underlying disk capacity managed by each ONTAP Select node and not the usable capacity of the node, HA-pair, or cluster. Each Capacity Tiers license is assigned to an individual node serial number for a specified capacity and requires an accompanying support contract.

The licensed capacity for a Capacity Tiers license is in full TB amounts, with a minimum of 1 TB. You calculate the licensed capacity by multiplying the amount of usable space by the overhead factor. The overhead factor estimates the ONTAP Select software overhead capacity to be taken from the usable space and it applies to determining the ONTAP Select capacity for NVMe, SSD, HDD, or virtual drives. The overhead factor differs between a single-node and multi-node (HA-pair based) clusters:

- Overhead factor for single-node clusters is 1.13
- Overhead factor for multi-node clusters is 2.67

You need to calculate the licensed capacity based on the usable space to ensure you have enough for your deployment. For example, you want 10 TB of usable capacity for an HA pair (5 TB per node):

- With an HA pair, the data is duplicated between both nodes for redundancy, so for 10 TB of usable data in an HA pair, you need to have 10 TB per node, in case one of the nodes goes offline.
- The factor of 2.67 doubles the capacity of each node and adds the overhead capacity (0.67) needed for ONTAP Select to function.
- The license capacity per node is 14 TB (5 x 2.67 = 13.35). For the HA pair, you require two 14 TB licenses to have 10 TB of usable capacity.

Description	Single node	HA pair
Usable space	10 TB	5 TB per node 10 TB for HA pair
Licensed space	12 TB (10 x 1.13 = 11.3, round up to next full TB)	Two 14 TB licenses (5 x 2.67 = 13.35, round up to next full TB)

Description	Single node	HA pair
Maximum usable space per node **Note**: The maximum managed space per node or HA pair is 400 TB	353 TB (400 / 1.13)	Per node in an HA pair: • 149 TB (400 / 2.67) Per HA pair in a multi-node cluster: • Two-node: 298 TB (2 x 149) • Four-node: 596 TB (4 x 149) • Six-node: 894 TB (6 x 149) • Eight-node: 1,192 TB (8 x 149) • Ten-node: 1,490 TB (10 x 149) • Twelve-node: 1,788 TB (12 x 149)

What's next

Acquire a Capacity Tier license for ONTAP Select

Acquire an ONTAP Select Capacity Tier license

When using Capacity Tiers licensing, you need to acquire a license file for each ONTAP Select node. The license file defines the storage capacity for the node and is locked to the node through the unique nine-digit serial number assigned by NetApp.

Before you begin

- You must have the nine-digit license serial number assigned to the ONTAP Select node by NetApp.
- You should wait at least 24 hours after the shipped date of your purchase order before acquiring a license file.
- You have reviewed the considerations for sizing a Capacity Tiers license.

About this task

You must perform this task for each ONTAP Select node requiring a Capacity Tiers license.

Steps

1. Access the ONTAP Select license site.
2. Sign in using your NetApp account credentials.
3. On the **License Generator** page, select the desired license offering from the dropdown box.
4. Fill in the remaining fields on the **License Generator** page, including the **Product Serial #**.

   The "product serial #" is the serial number for the ONTAP Select node.

5. Select **Submit**.
6. After the request has been validated, select the delivery method for the license.

   You can select either **Download License** or **Email License**.

7. Confirm that you received the license file based on your selected delivery method.

After you finish

You must upload the license file to the Deploy administration utility before it can be applied to an ONTAP Select node.

Capacity Pools

Determine Capacity Pools license size for ONTAP Select

The licensed capacity for a Capacity Pool license is in full TB amounts, with a minimum of 2 TB. You calculate the licensed capacity by multiplying the amount of usable space by the overhead factor. The overhead factor estimates the ONTAP Select software overhead capacity to be taken from the usable space and it applies to determining the ONTAP Select capacity for NVMe, SSD, HDD, or virtual drives. The overhead factor differs between a single-node and multi-node (HA-pair based) clusters:

- Overhead factor for single-node clusters is 1.13
- Overhead factor for multi-node clusters is 2.67

You need to calculate the licensed capacity based on the usable space to ensure you have enough for your deployment. For example, you want 20 TB of usable capacity for an HA pair (10 TB per node):

- With an HA pair, the data is duplicated between both nodes for redundancy, so for 20 TB of usable data in an HA pair, you need to have 20 TB per node, in case one of the nodes goes offline.
- The factor of 2.67 doubles the capacity of each node and adds the overhead capacity (0.67) needed for ONTAP Select to function.
- The license capacity per node is 27 TB (10 x 2.67 = 26.7). For the HA pair, you require two 27 TB licenses to have 20 TB of usable capacity.

Description	Single-node	HA pair or mix of HA pair and single-node
Usable space	20 TB	10 TB per single-node cluster 20 TB for HA pair
Licensed space	23 TB (20 x 1.13 = 22.6, round up to next full TB)	54 TB (20 x 2.67 = 53.4, round up to next full TB)
Maximum usable space per node **Note**: The maximum managed space per node or HA pair is 400 TB	353 TB (400 / 1.13)	Per single-node in a cluster: 149 TB (400 / 2.67) Per HA pair in a cluster: 298 TB (2 x 149)

What's next

Acquire a Capacity Pools license for ONTAP Select

Acquire an ONTAP Select Capacity Pool license

You must acquire a license file for each Capacity Pool used by the ONTAP Select nodes. The license file defines the storage capacity and expiration for the pool. It is locked to the License Manager through a combination of the unique license serial number assigned by NetApp and the License Lock ID associated with the Deploy instance.

Before you begin

- You must have the nine-digit license serial number assigned to the Capacity Pool by NetApp.
- You should wait at least 24 hours after the shipped date of your purchase order before acquiring a license file.
- You have reviewed the considerations for sizing a Capacity Pools license.

About this task

You must perform this task for each Capacity Pool used by the ONTAP Select nodes.

Steps

1. Log in to the NetApp Support site.
2. Select **Systems** > **Software Licenses**.
3. Type the license serial number for the Capacity Pool and select **Go!**.
4. On the license details page, navigate to the **Product Details** column.
5. Select **Get NetApp License File** on the appropriate row.
6. Type the License Lock ID for your ONTAP Select Deploy instance and select **Submit**.
7. Select the appropriate delivery method and select **Submit**.
8. Select **OK** on the delivery confirmation window.

After you finish

You must upload the license file to the Deploy administration utility before an ONTAP Select node can use the Capacity Pool.

ONTAP Select support for ONTAP features

ONTAP Select supports most ONTAP features. Many of the ONTAP features are licensed automatically for each node when you deploy a cluster. However some features require a separate license.

ONTAP features that have hardware-specific dependencies are generally not supported with ONTAP Select.

ONTAP features licensed by default

The following ONTAP features are supported with ONTAP Select and licensed by default:

- Autonomous Ransomware Protection (ARP) (manual updates)
- CIFS

- Deduplication and compression
- External key management using the Key Management Interoperability Protocol (KMIP)

  Learn about enabling and configuring external key management.

- FlexCache

  ONTAP Select supports FlexCache cache volumes with write-back. FlexCache origin volumes with write-back are not supported.

  Learn more about FlexCache write-back with ONTAP Select

- FlexClone
- iSCSI
- NDMP
- NetApp Volume Encryption (non-restricted countries only)
- NFS
- NVMe over TCP
- ONTAP multitenancy capability
- ONTAP S3
- S3 SnapMirror
- SnapLock Select

  SnapLock Select is the implementation of ONTAP SnapLock designed for ONTAP Select and includes SnapLock Enterprise. It does not include SnapLock Compliance.

  Learn about ONTAP SnapLock

- SnapMirror

  ONTAP Select does not support consistency groups in Asynchronous SnapMirror.

- SnapRestore
- SnapVault
- Storage VM disaster recovery (SVM DR)

  ONTAP Select supports SVM DR as both a source and destination with a maximum of 16 relationships. SVM DR support is limited to the use of the source ONTAP version to versions +2. For example, the ONTAP Select 9.16.1 source can connect to destination ONTAP versions 9.16.1, 9.17.1, or 9.18.1.

ONTAP features that are separately licensed

You need to acquire a separate license for any ONTAP feature that is not licensed by default, including:

- FabricPool

  You do not need a FabricPool license when utilizing StorageGRID.

- MetroCluster SDS (ONTAP Select premium license offering)

- SnapMirror cloud

  You need to open a non-technical support case to obtain a SnapMirror cloud license.

  Learn about SnapMirror cloud backups to object storage

Related information

- Compare ONTAP Select and ONTAP 9
- Learn about the NetApp ONTAP master license keys

Install

Pre-installation checklist

Host preparation checklist

Prepare a KVM host for ONTAP Select

Prepare each KVM hypervisor host where an ONTAP Select node is deployed. As you prepare the hosts, you carefully assess the deployment environment to make sure that the hosts are correctly configured and ready to support the deployment of an ONTAP Select cluster.

> The ONTAP Select Deploy administration utility doesn't perform the required network and storage configuration of the hypervisor hosts. You must manually prepare each host before deploying an ONTAP Select cluster.

Step 1: Prepare the KVM hypervisor host

You need to prepare each of the Linux KVM servers where an ONTAP Select node is deployed. You must also prepare the server where the ONTAP Select Deploy administration utility is deployed.

Steps

1. Install Red Hat Enterprise Linux (RHEL).

   Install the RHEL operating system using the ISO image. Refer to the software compatibility information for hypervisors for a list of the supported RHEL versions. During installation, configure the system as follows:

   a. Select Default as the security policy.
   b. Choose the Virtualized Host software selection.
   c. Verify that the destination is the local boot disk and not a RAID LUN used by ONTAP Select.
   d. Verify that the host management interface is up after you boot the system.

   > You can edit the correct network configuration file under /etc/sysconfig/network-scripts and then bring up the interface by using the `ifup` command.

2. Install additional packages required for ONTAP Select.

   ONTAP Select requires several additional software packages. The exact list of packages varies based on the version of Linux that you're using. As a first step, verify that the yum repository is available on your server. If it isn't available, you can retrieve it by using the `wget your_repository_location` command.

   > Some of the required packages might already be installed if you chose Virtualized Host for the software selection during installation of the Linux server. You might need to install the openvswitch package from source code as described in the Open vSwitch documentation.

   For additional information about the necessary packages and other configuration requirements, refer to the Interoperability Matrix Tool.

3. Configure PCI passthrough for NVMe disks.

   If you're using NVMe disks in your configuration, you need to configure PCI passthrough (DirectPath IO) to provide direct access for the KVM host to the locally attached NVMe disks in the ONTAP Select cluster. You need direct access to perform the following tasks:

   - Configure the KVM host to use NVMe drives
   - Use software RAID after you deploy the cluster

   Refer to the Red Hat Documentation for instructions on how to configure PCI passthrough (DirectPath IO) for a KVM hypervisor.

4. Configure the storage pools.

   An ONTAP Select storage pool is a logical data container that abstracts the underlying physical storage. You must manage the storage pools on the KVM hosts where ONTAP Select is deployed.

Step 2: Create a storage pool

Create at least one storage pool at each ONTAP Select node. If you use software RAID instead of a local hardware RAID, storage disks are attached to the node for the root and data aggregates. In this case, you must still create a storage pool for the system data.

Before you begin

Verify that you can sign in to the Linux CLI on the host where ONTAP Select is deployed.

About this task

The ONTAP Select Deploy administration utility expects the target location for the storage pool to be specified as `/dev/<pool_name>`, where `<pool_name>` is a unique pool name on the host.

> The entire capacity of the LUN is allocated when a storage pool is created.

Steps

1. Display the local devices on the Linux host and choose the LUN that will contain the storage pool:

   ```
   lsblk
   ```

   The appropriate LUN is likely to be the device with the largest storage capacity.

2. Define the storage pool on the device:

   ```
   virsh pool-define-as <pool_name> logical --source-dev <device_name>
   --target=/dev/<pool_name>
   ```

   For example:

   ```
   virsh pool-define-as select_pool logical --source-dev /dev/sdb
   --target=/dev/select_pool
   ```

3. Build the storage pool:

```
virsh pool-build <pool_name>
```

4. Start the storage pool:

```
virsh pool-start <pool_name>
```

5. Configure the storage pool to automatically start at system boot:

```
virsh pool-autostart <pool_name>
```

6. Verify that the storage pool has been created:

```
virsh pool-list
```

Step 3: Optionally, delete a storage pool

You can delete a storage pool when it's no longer needed.

Before you begin

Verify that you can sign in to the Linux CLI where ONTAP Select is deployed.

About this task

The ONTAP Select Deploy administration utility expects the target location for the storage pool to be specified as `/dev/<pool_name>`, where `<pool_name>` is a unique pool name on the host.

Steps

1. Verify that the storage pool is defined:

```
virsh pool-list
```

2. Destroy the storage pool:

```
virsh pool-destroy <pool_name>
```

3. Undefine the configuration for the inactive storage pool:

```
virsh pool-undefine <pool_nanme>
```

4. Verify that the storage pool has been removed from the host:

```
virsh pool-list
```

5. Verify that all logical volumes for the storage pool volume group have been deleted.
   a. Display the logical volumes:

   ```
   lvs
   ```

   b. If any logical volumes exist for the pool, delete them:

   ```
   lvremove <logical_volume_name>
   ```

6. Verify that the volume group has been deleted:
   a. Display the volume groups:

   ```
   vgs
   ```

   b. If a volume group exists for the pool, delete it:

   ```
   vgremove <volume_group_name>
   ```

7. Verify that the physical volume has been deleted:
   a. Display the physical volumes:

   ```
   pvs
   ```

   b. If a physical volume exists for the pool, delete it:

   ```
   pvremove <physical_volume_name>
   ```

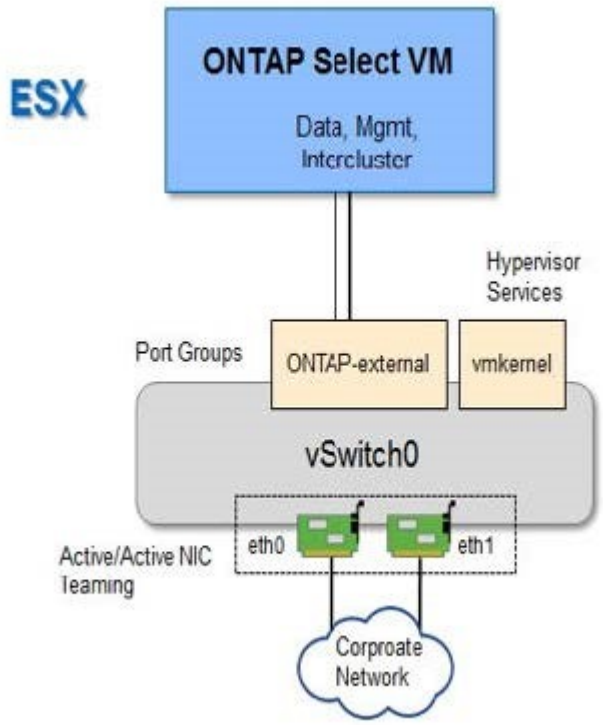
Step 4: Review the ONTAP Select cluster configuration

You can deploy ONTAP Select as either a multi-node cluster or a single-node cluster. In many cases, a multi-node cluster is preferable because of the additional storage capacity and high-availability (HA) capability.

The following figures illustrate the ONTAP Select networks used with a single-node cluster and four-node cluster for an ESXi host.

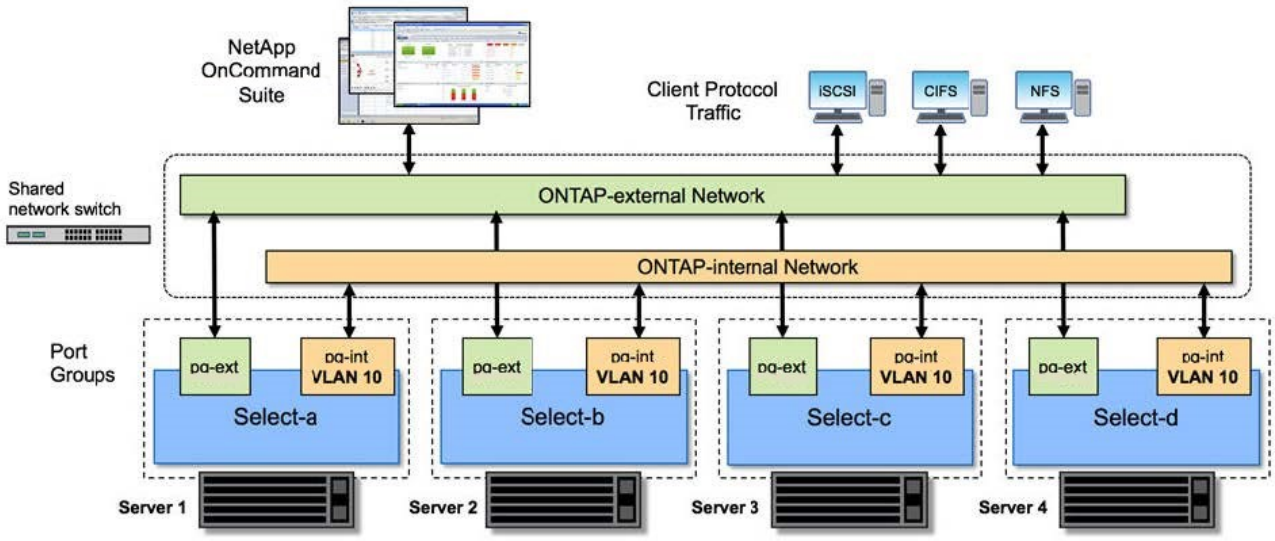
Single-node cluster

The following figure illustrates a single-node cluster. The external network carries client, management, and cross-cluster replication traffic (SnapMirror/SnapVault).



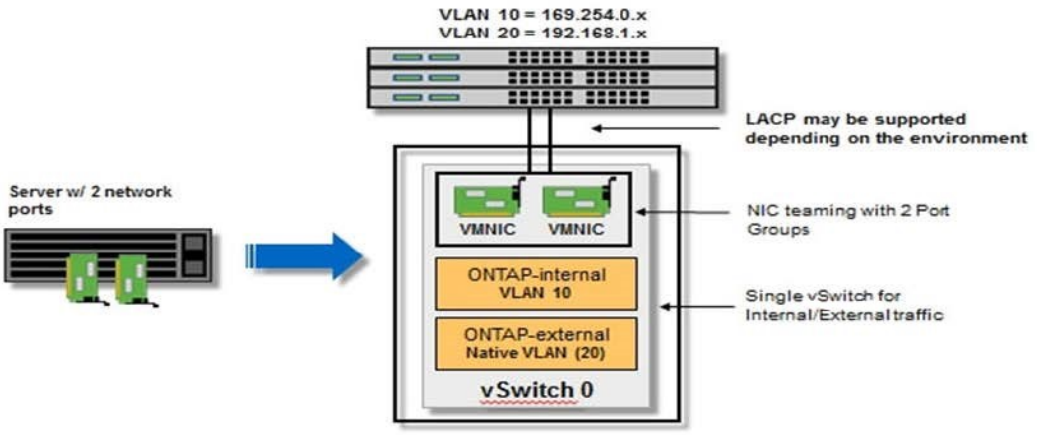
Four-node cluster

The following figure illustrates a four-node cluster showing two networks. The internal network enables communication among the nodes in support of the ONTAP cluster network services. The external network carries client, management, and cross-cluster replication traffic (SnapMirror/SnapVault).



Single node within a four-node cluster

The following figure illustrates the typical network configuration for a single ONTAP Select virtual machine within a four-node cluster. There are two separate networks: ONTAP-internal and ONTAP-external.



Step 5: Configure Open vSwitch

Use Open vSwitch to configure a software-defined switch on each KVM host node.

Before you begin

Verify that the network manager is disabled and the native Linux network service is enabled.

About this task

ONTAP Select requires two separate networks, both of which utilize port bonding to provide HA capability for the networks.

Steps

1. Verify that Open vSwitch is active on the host:
   a. Determine if Open vSwitch is running:

      ```
      systemctl status openvswitch
      ```

   b. If Open vSwitch is not running, start it:

      ```
      systemctl start openvswitch
      ```

2. Display the Open vSwitch configuration:

   ```
   ovs-vsctl show
   ```

   The configuration appears empty if Open vSwitch has not already been configured on the host.

3. Add a new vSwitch instance:

```
ovs-vsctl add-br <bridge_name>
```

For example:

```
ovs-vsctl add-br ontap-br
```

4. Bring the network interfaces down:

```
ifdown <interface_1>
ifdown <interface_2>
```

5. Combine the links using the Link Aggregation Control Protocol (LACP):

```
ovs-vsctl add-bond <internal_network> bond-br <interface_1>
<interface_2> bond_mode=balance-slb lacp=active other_config:lacp-
time=fast
```

You only need to configure a bond if there is more than one interface.

6. Bring the network interfaces up:

```
ifup <interface_1>
ifup <interface_2>
```

Prepare an ESXi host for ONTAP Select

Prepare each ESXi hypervisor host where an ONTAP Select node is deployed. As you prepare the hosts, you carefully assess the deployment environment to make sure that the hosts are correctly configured and ready to support the deployment of an ONTAP Select cluster.

The ONTAP Select Deploy administration utility doesn't perform the required network and storage configuration of the hypervisor hosts. You must manually prepare each host before deploying an ONTAP Select cluster.

Step 1: Prepare the ESXi hypervisor host

Verify the configuration for the ESXi host and the firewall ports.

Steps

1. Verify that each ESXi is configured with the following:
   - A pre-installed and supported hypervisor
   - A VMware vSphere license
2. Verify that the same vCenter server can manage all the hosts where an ONTAP Select node is deployed within the cluster.
3. Verify that the firewall ports are configured to allow access to vSphere. These ports must be open to support serial port connectivity to the ONTAP Select virtual machines.

   Recommended

   NetApp recommends that you open the following firewall ports to allow access to vSphere:

   - Ports 7200 – 7400 (both inbound and outbound traffic)

   Default

   By default, VMware allows access on the following ports:

   - Port 22 and ports 1024 – 65535 (inbound traffic)
   - Ports 0 – 65535 (outbound traffic)

   For more information, refer to the Broadcom VMware vSphere documentation.

4. Familiarize yourself with the vCenter rights that are required. See VMware vCenter server for more information.

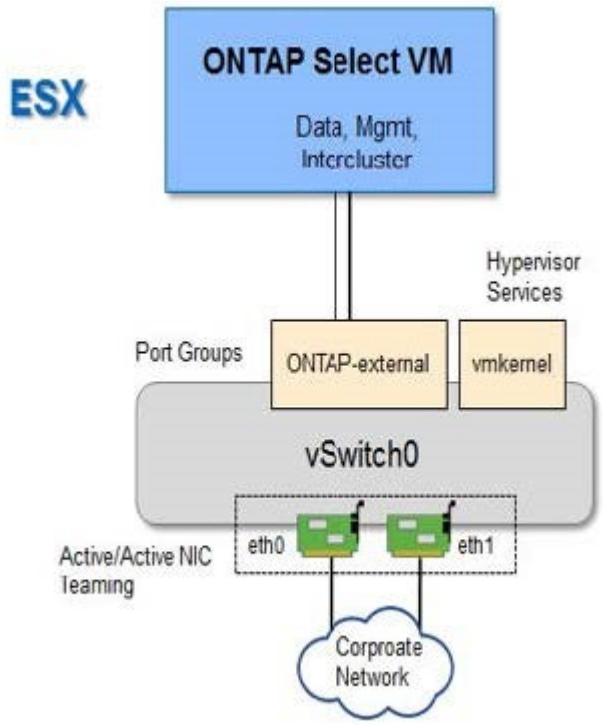
Step 2: Review the ONTAP Select cluster configuration

You can deploy ONTAP Select as either a multi-node cluster or a single-node cluster. In many cases, a multi-node cluster is preferable because of the additional storage capacity and high-availability (HA) capability.

The following figures illustrate the ONTAP Select networks used with a single-node cluster and four-node cluster.

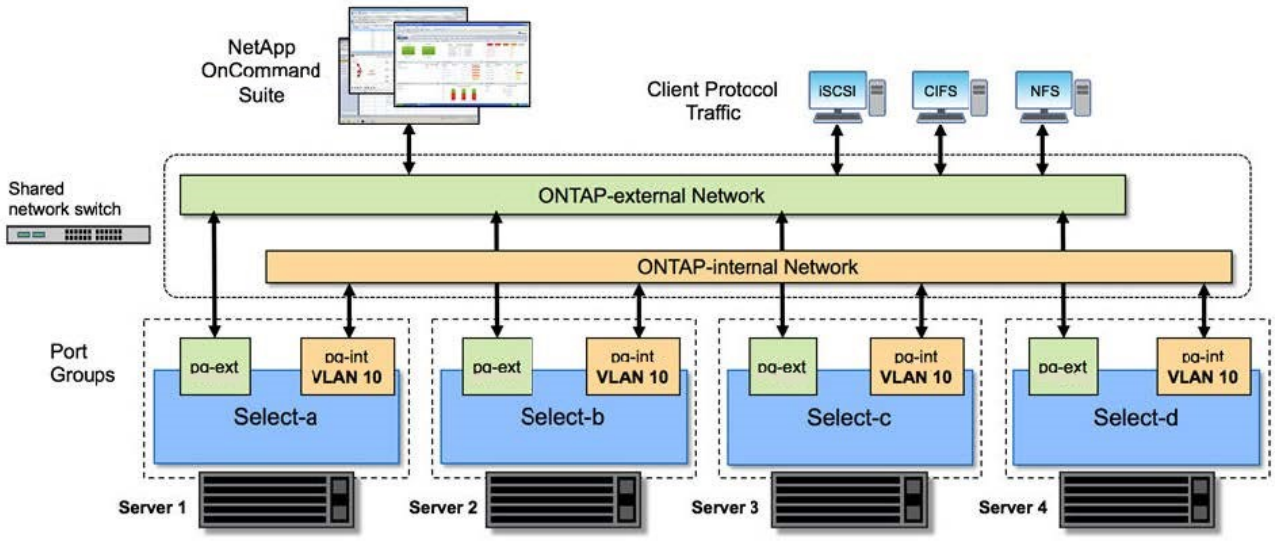
Single-node cluster

The following figure illustrates a single-node cluster. The external network carries client, management, and cross-cluster replication traffic (SnapMirror/SnapVault).



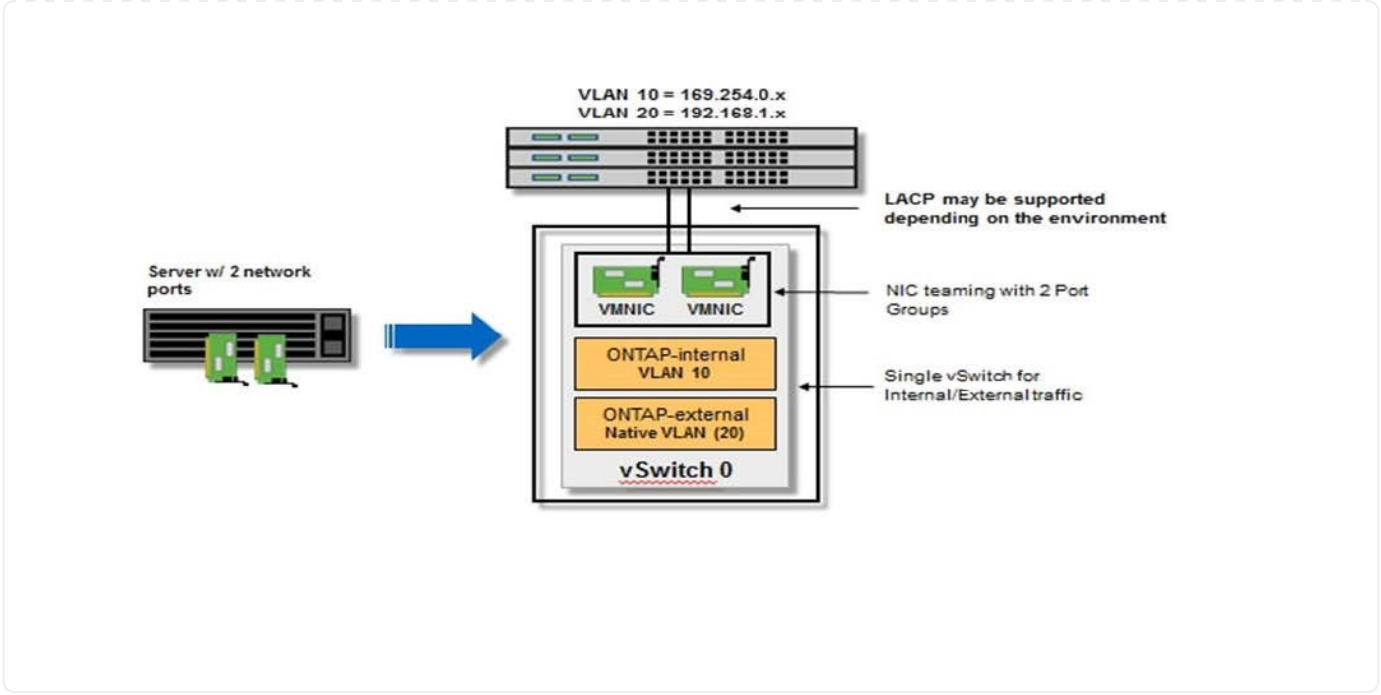
Four-node cluster

The following figure illustrates a four-node cluster showing two networks. The internal network enables communication among the nodes in support of the ONTAP cluster network services. The external network carries client, management, and cross-cluster replication traffic (SnapMirror/SnapVault).



Single node within a four-node cluster

The following figure illustrates the typical network configuration for a single ONTAP Select virtual machine within a four-node cluster. There are two separate networks: ONTAP-internal and ONTAP-external.



Required information for ONTAP Select Deploy utility installation

Before installing the Deploy administration utility in a hypervisor environment, review the required configuration information and optional network configuration information to prepare for successful deployment.

Required configuration information

As part of your deployment planning, you should determine the required configuration information before installing the ONTAP Select Deploy administration utility.

Required information	Description
Name of the Deploy virtual machine	Identifier to use for the virtual machine.
Name of the hypervisor host	Identifier for the VMware ESXi or KVM hypervisor host where the Deploy utility is installed.
Name of the data store	Identifier for the hypervisor data store holding the virtual machine files (approximately 40GB is required).
Network for the virtual machine	Identifier for the Network where the Deploy virtual machine is connected.

Optional network configuration information

The Deploy virtual machine is configured using DHCP by default. However, if needed, you can manually configure the network interface for the virtual machine.

Network information	Description
Host name	Identifier of the host machine.
Host IP address	Static IPv4 address of the host machine.

Network information	Description
Subnet mask	Subnetwork mask, based on the network the virtual machine is a part of.
Gateway	Default gateway or router.
Primary DNS server	Primary Domain Name Server.
Secondary DNS server	Secondary Domain Name Server.
Search domains	List of the search domains to use.

Required information for ONTAP Select installation

As part of preparing to deploy an ONTAP Select cluster in a VMware environment, collect the information required when using the ONTAP Select Deploy administration utility to deploy and configure the cluster.

Some of the information you collect applies to the cluster itself, while other information applies to the individual nodes in the cluster.

Cluster-level information

You must collect information related to the ONTAP Select cluster.

Cluster information	Description
Name of the cluster	Unique identifier of the cluster.
Licensing mode	Evaluation or purchased licensing.
IP configuration for the cluster	IP configuration for the clusters and nodes, including: * Management IP address of the cluster * Subnet mask * Default gateway

Host-level information

You must collect information related to each of the nodes in the ONTAP Select cluster.

Cluster information	Description
Name of the host	Unique identifier of the host.
Domain name of the host	Fully qualified domain name of the host.
IP configuration for the nodes	Management IP address for each node in the cluster.
Mirror node	Name of the associated node in the HA pair (multi-node clusters only).
Storage pool	Name of the storage pool that is used.
Storage disks	List of disks if using software RAID.
Serial number	If you are deploying with a purchased license, the unique nine-digit serial number provided by NetApp.

Configure an ONTAP Select host to use NVMe drives

If you plan to use NVMe drives with software RAID, you need to configure the ESXi or KVM host to recognize the drives.

Use VMDirectPath I/O Pass-through on the NVMe devices to maximize data efficiency. This setting exposes the drives to the ONTAP Select virtual machine, allowing ONTAP to have direct PCI access to the device.

Step 1: Configure the host

Configure the ESXi or KVM host to recognize the drives.

Before you begin

Make sure your deployment environment meets the following minimum requirements:

- For an ESXi host, ONTAP Select 9.7 or later with a supported Deploy administration utility
- For a KVM host, ONTAP Select 9.17.1 or later with a supported Deploy administration utility
- Premium XL platform license offering or a 90-day evaluation license
- The ESXi or KVM host is running a supported hypervisor version:

  ESXi

  ESXi is supported on the following hypervisor versions:

  - VMware ESXi 9.0
  - VMware ESXi 8.0 U3
  - VMware ESXi 8.0 U2
  - VMware ESXi 8.0 U1 (build 21495797)
  - VMware ESXi 8.0 GA (build 20513097)

  KVM

  KVM is supported on the following hypervisor versions:

  - Red Hat Enterprise Linux (RHEL) 10.1, 10.0, 9.7, 9.6, 9.5, 9.4, 9.2, 9.1, 9.0, 8.8, 8.7, and 8.6
  - Rocky Linux 10.1, 10.0, 9.7, 9.6, 9.5, 9.4, 9.3, 9.2, 9.1, 9.0, 8.9, 8.8, 8.7, and 8.6

- NVMe devices conforming to specification 1.0 or later

Follow the host preparation checklist, and review the required information for the deploy utility installation and the ONTAP Select installation for more information.

About this task

You must perform this procedure before creating a new ONTAP Select cluster. You can also perform the procedure to configure additional NVMe drives for an existing software RAID NVMe cluster. In this case, after configuring the drives, you must add them through Deploy as you would additional SSD drives. The main difference is that Deploy detects the NVMe drives and reboots the nodes. When adding NVMe drives to an existing cluster, note the following about the reboot process:

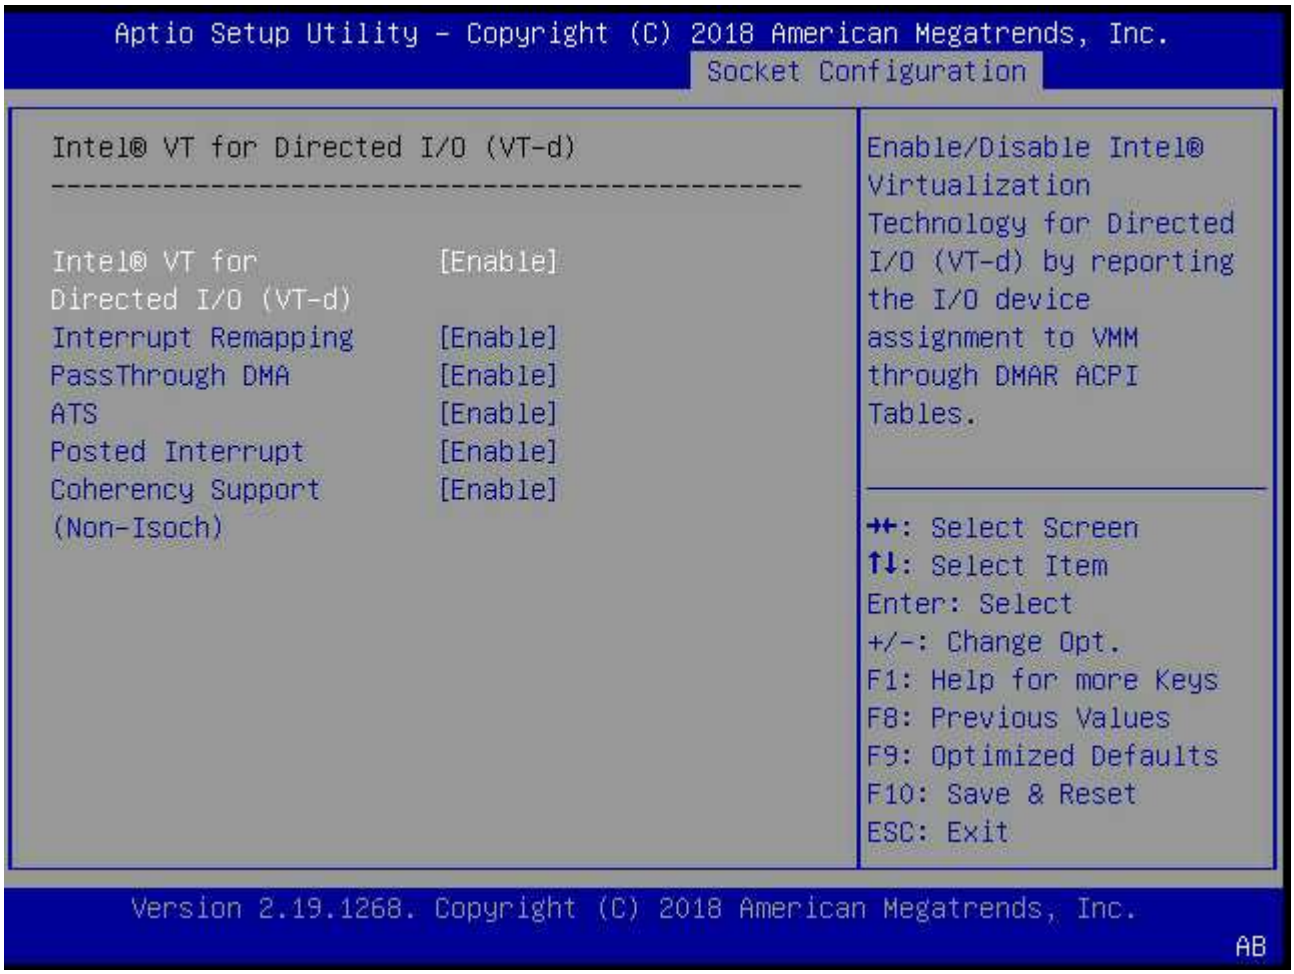
- Deploy handles the reboot orchestration.

- HA takeover and giveback is performed in an orderly fashion, but it can be time consuming to resynchronize the aggregates.
- A single-node cluster will incur downtime.

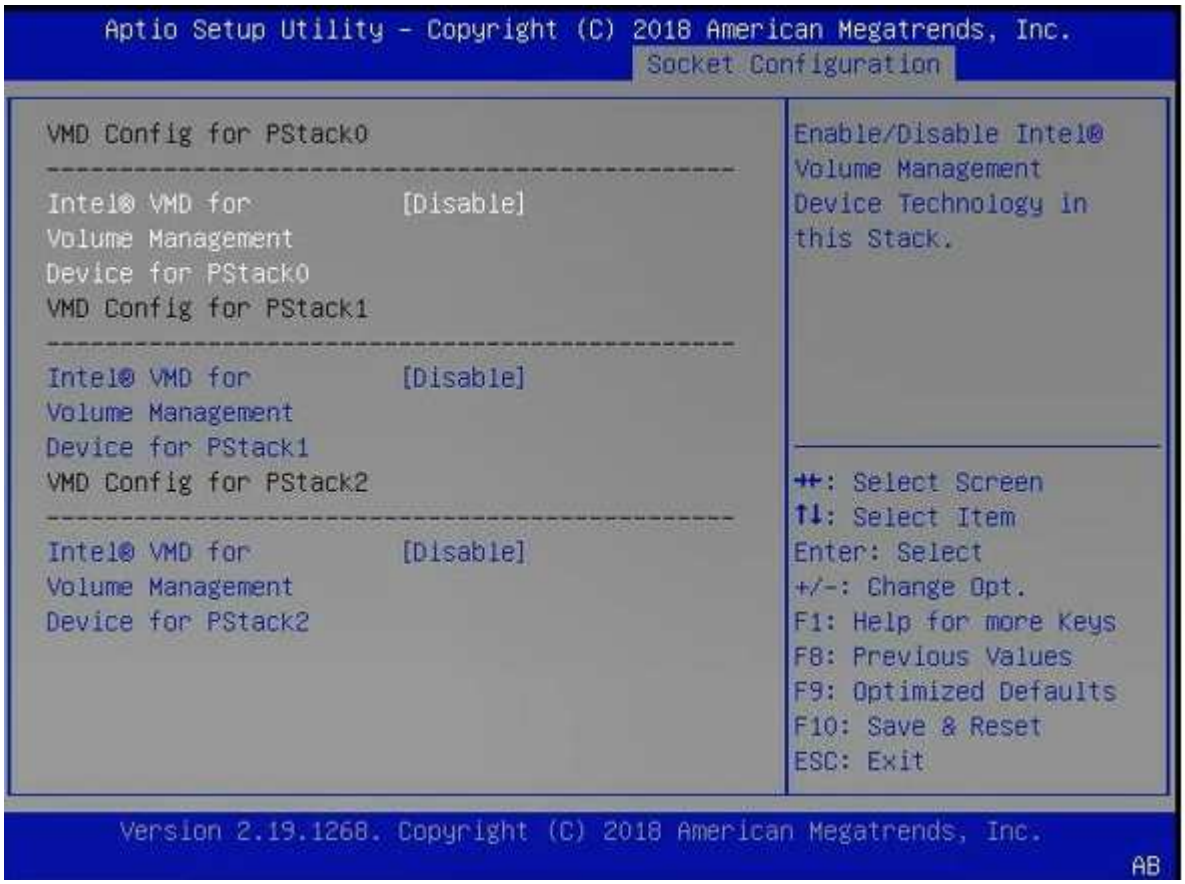
See Increase storage capacity for additional information.

Steps

1. Access the **BIOS configuration** menu on the host to enable support for I/O virtualization.
2. Enable the **Intel VT for Directed I/O (VT-d)** setting.

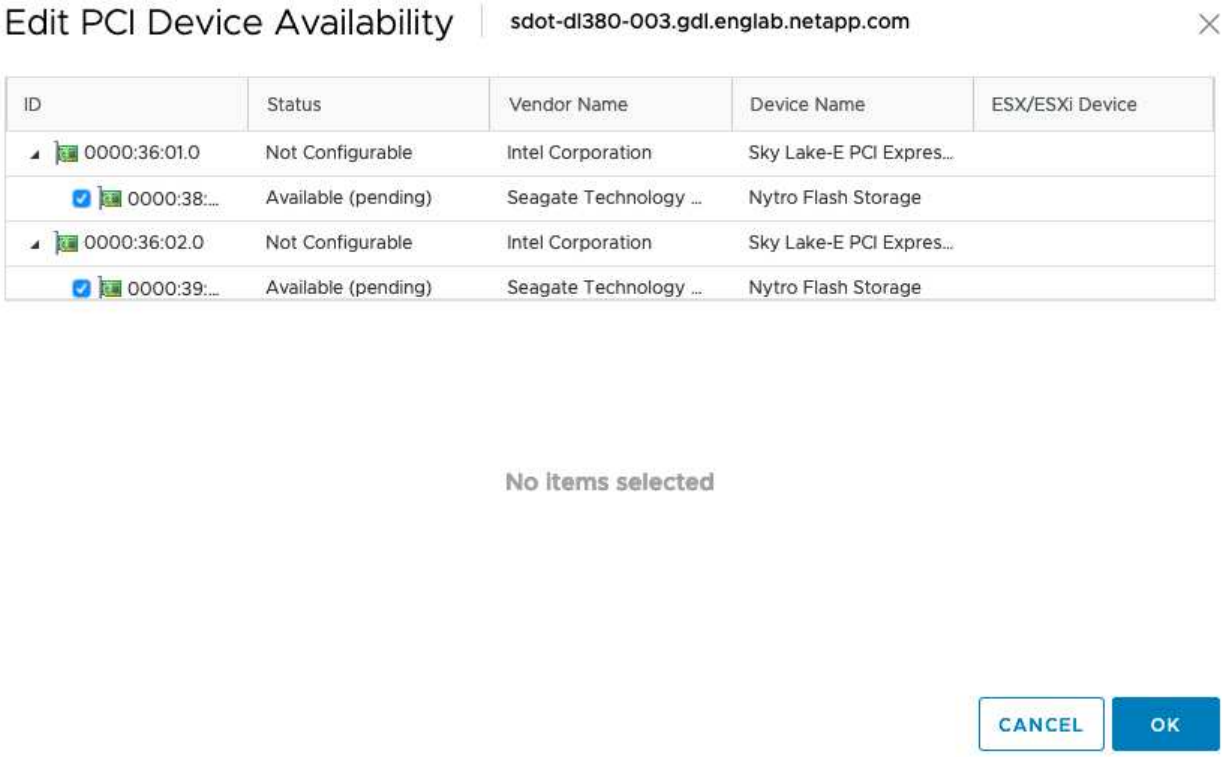


3. Some servers offer support for **Intel Volume Management Device (Intel VMD)**. When enabled, this makes the available NVMe devices invisible to the ESXi or KVM hypervisor; disable this option before proceeding.



4. Configure the NVMe drives for pass-through to virtual machines.
   a. In vSphere, open the host **Configure** view and select **Edit** under **Hardware: PCI devices**.
   b. Select the NVMe drives you want to use for ONTAP Select.

      The following example output shows the available drives for an ESXi host:

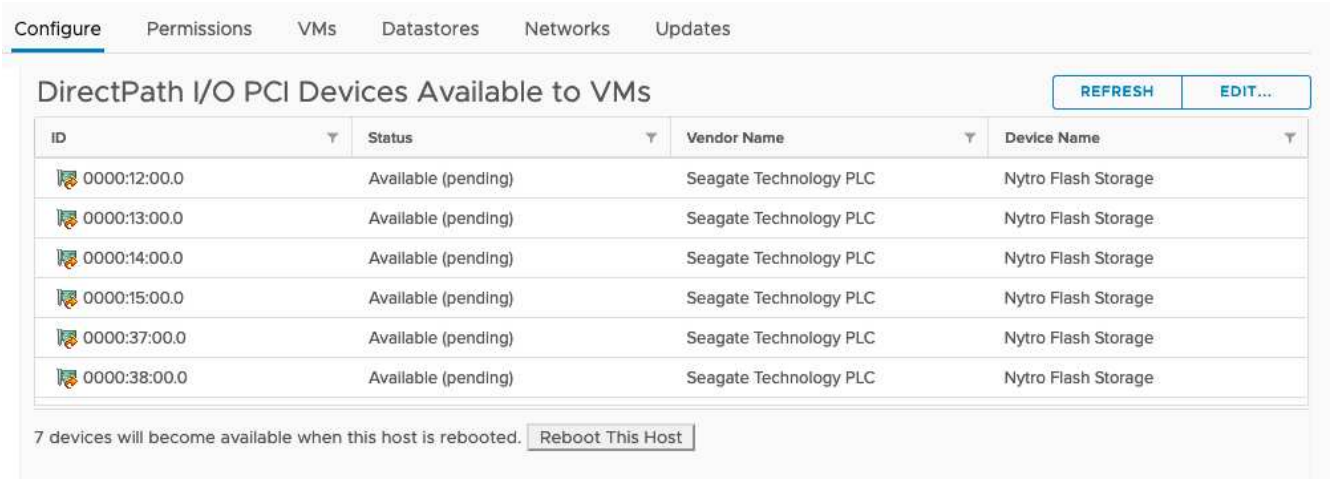


You need a VMFS datastore that is also backed by an NVMe device to host the ONTAP Select VM system disks and virtual NVRAM. Leave at least one NVMe drive available for this purpose when configuring the others for PCI pass-through.

c. Select **OK**. The selected devices indicate **Available (pending)**.

5. Select **Reboot This Host**.

The following example output is for an ESXi host:



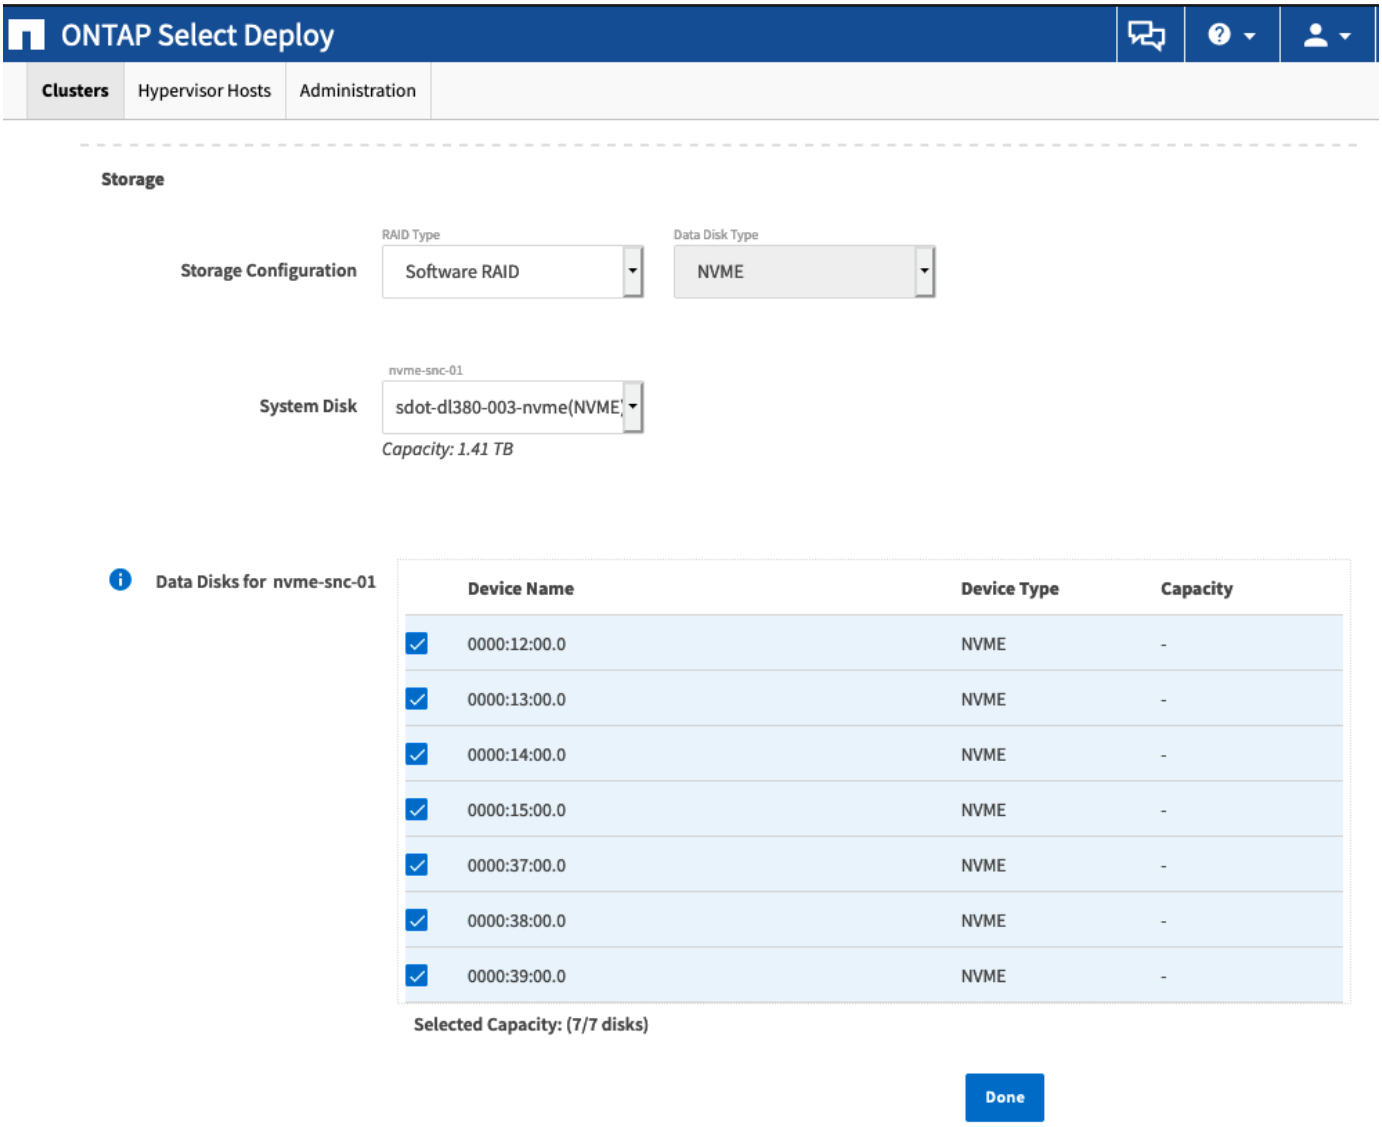
Step 2: Install the ONTAP Select Deploy utility

After the hosts are prepared, you can install the ONTAP Select Deploy utility. Deploy guides you through creating ONTAP Select storage clusters on your newly prepared hosts. During this process, Deploy detects the presence of the NVMe drives configured for pass-through and automatically selects them for use as ONTAP

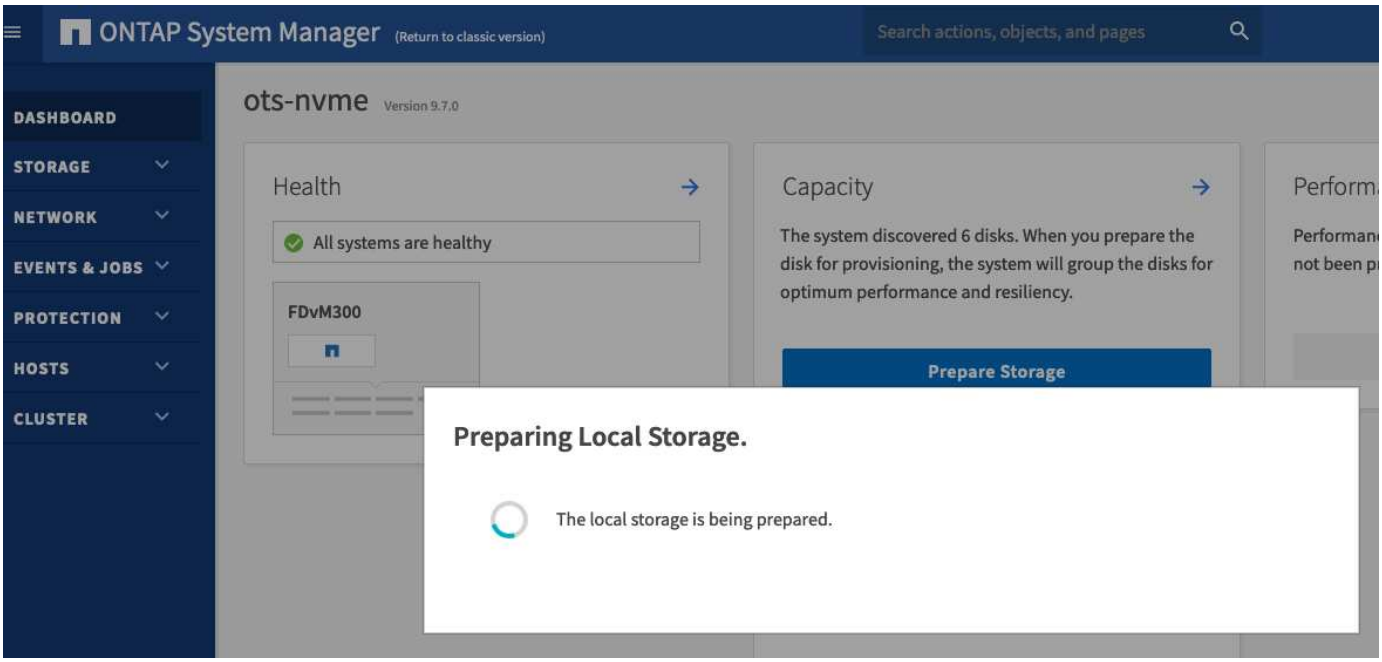
data disks. You can adjust the default selection if needed.

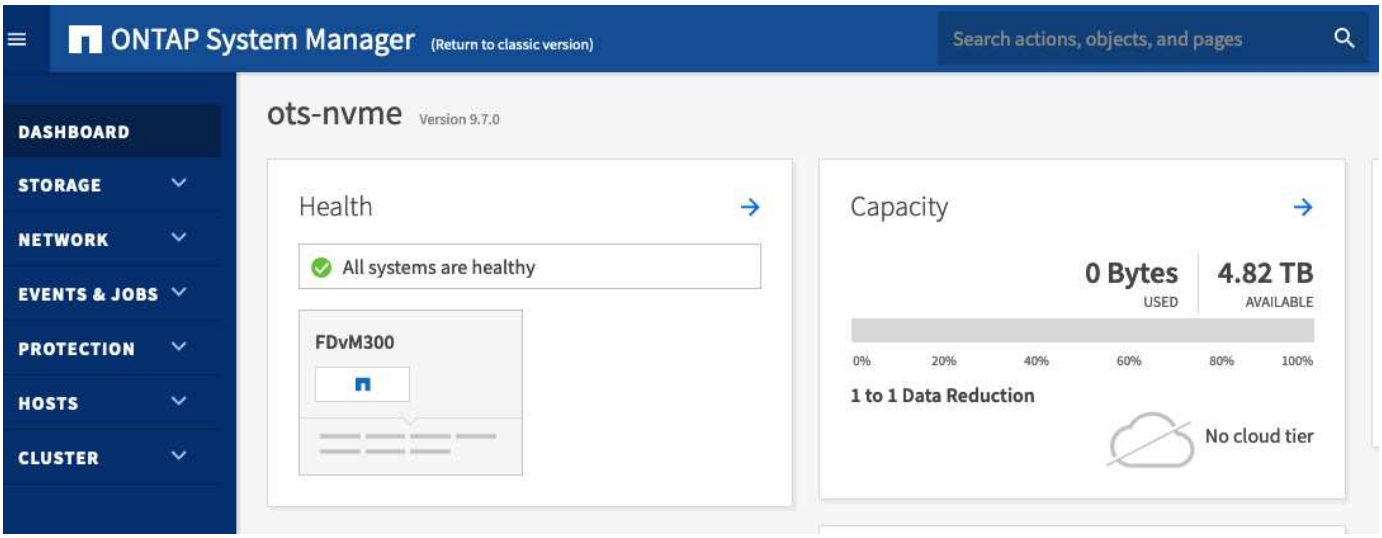
A maximum of 14 NVMe devices are supported for each ONTAP Select node.

The following example output is for an ESXi host:



After the cluster is successfully deployed, ONTAP System Manager allows you to provision the storage according to best practices. ONTAP automatically enables flash-optimized storage efficiency features that make the best use of your NVMe storage.





Install ONTAP Select Deploy

You need to install the ONTAP Select Deploy administration utility and use the utility to create an ONTAP Select cluster.

Before you begin

You have reviewed the pre-installation checklist.

Step 1: Download the virtual machine image

Download the ONTAP Select package from the NetApp support site.

Before you begin

You have a registered NetApp Support Site account.

About this task

The ONTAP Select Deploy administration utility is packaged as a virtual machine (VM) based on the Open Virtualization Format (OVF) standard. The single compressed file has the suffix `OVA` for ESXi hosts and `TGZ` for KVM hosts. The VM provides the Deploy server and installation images for ONTAP Select nodes.

Steps

1. Access the NetApp Support Site downloads page.
2. Scroll down and select **ONTAP Select Deploy**.
3. Select the desired ONTAP Select version.
4. Review the End User License Agreement (EULA) and select **Accept & Continue**.
5. Select and download the appropriate **ONTAP Select Deploy Install** package. Respond to all prompts as needed.

Step 2: Verify the ONTAP Select Deploy OVA signature

Verify the ONTAP Select Open Virtualization Appliance (OVA) signature before installing the installation package.

Before you begin

Verify that your system meets the following requirements:

- OpenSSL versions 1.0.2 to 3.0 for basic verification
- Public internet access for Online Certificate Status Protocol (OCSP) verification

Steps

1. Obtain the following files from the product download page on the NetApp support site:

File	Description
`ONTAP-Select-Deploy-Production.pub`	The public key used to verify the signature.
`csc-prod-chain-ONTAP-Select-Deploy.pem`	The public certification authority (CA) chain of trust.
`csc-prod-ONTAP-Select-Deploy.pem`	The certificate used to generate the key.
`ONTAPdeploy.ova`	The product installation executable for ONTAP Select.
`ONTAPdeploy.ova.sig`	The SHA-256 algorithm is hashed and then signed by the Remote Support Agent (RSA) using the `csc-prod` key and signature for the installer.

2. Verify that the `ONTAPdeploy.ova.sig` file is using the associated certificates and validation commands.
3. Verify the signature:

```
openssl dgst  -sha256 -verify ONTAP-Select-Deploy-Production.pub
-signature ONTAPdeploy.ova.sig  ONTAPdeploy.ova
```

Step 3: Deploy the virtual machine

Install and start the ONTAP Select Deploy VM using the OVF VM image. As part of the installation process, you configure the network interface to use DHCP or a static IP configuration.

Before you begin

For an ESXi hypervisor, you must prepare to deploy the ONTAP Select Deploy VM:

- Enable the OVF functionality in your browser by installing the VMware Client Integration Plugin or performing similar configuration as needed
- Enable the DHCP in the VMware environment if you will dynamically assign an IP address to the Deploy VM

For ESXi and KVM hypervisors, you must have the configuration information to be used when creating the VM, including the name of the VM, external network, and host name. When defining a static network configuration, you need the following additional information:

- IP address of the Deploy VM
- Netmask
- IP address of the gateway (router)
- IP address of the primary DNS server
- IP address of the second DNS server
- DNS search domains

About this task

If you use vSphere, the Deploy OVF template wizard includes a form to provide all of the Deploy configuration information, including the network configuration. However, if you choose not to use this form, you can use the console of the Deploy VM to configure the network instead.

Steps

The steps you follow depend on whether you use an ESXi or KVM hypervisor.

ESXi

1. Access the vSphere client and sign in.
2. Navigate to the appropriate location in the hierarchy and select **Deploy OVF Template**.
3. Select the OVA file and complete the Deploy OVF Template wizard. Select the options that are appropriate for your environment.

   You must define the password for the administrator account. You need to provide this password when signing in to the Deploy utility.

4. After the VM is deployed, select the new VM. If it isn't already powered on based on your input to the deployment wizard, power it on manually.
5. If needed, you can configure the Deploy network using the VM console:
   a. Select the **Console** tab to access the ESXi host setup shell and monitor the power on process.
   b. Wait for the following prompt:

      Host name :

c. Type the host name and select **Enter**.

d. Wait for the following prompt:

   Provide a password for the admin user:

e. Type the password and select **Enter**.

f. Wait for the following prompt:

   Use DHCP to set networking information? [n]:

g. Type **n** to define a static IP configuration or **y** to use the DHCP, and select **Enter**.

h. If you choose a static configuration, provide all the required network configuration information.

KVM

1. Sign in to the CLI at the Linux server:

```
ssh root@<ip_address>
```

2. Create a new directory and extract the raw VM image:

```
mkdir /home/select_deploy25
cd /home/select_deploy25
mv /root/<file_name> .
tar -xzvf <file_name>
```

3. Create and start the KVM VM running the Deploy administration utility:

```
virt-install --name=select-deploy --vcpus=2 --ram=4096 --os
-variant=debian10 --controller=scsi,model=virtio-scsi --disk
path=/home/deploy/ONTAPdeploy.raw,device=disk,bus=scsi,format=raw
--network "type=bridge,source=ontap-
br,model=virtio,virtualport_type=openvswitch" --console=pty --import
--noautoconsole
```

4. If needed, you can configure the Deploy network using the VM console:

   a. Connect to the VM console:

```
virsh console <vm_name>
```

   b. Wait for the following prompt:

```
Host name :
```

c. Type the host name and select **Enter**.

d. Wait for the following prompt:

```
Use DHCP to set networking information? [n]:
```

e. Type **n** to define a static IP configuration or **y** to use the DHCP, and select **Enter**.

f. If you choose a static configuration, provide all network configuration information as required.

Step 4: Sign in to the Deploy web interface

Sign in to the web user interface to confirm the Deploy utility is available and perform the initial configuration.

Steps

1. Point your browser to the Deploy utility using the IP address or domain name:

   ```
   https://<ip_address>/
   ```

2. Provide the administrator (admin) account name and password and sign in.
3. If the **Welcome to ONTAP Select** pop-up window is displayed, review the prerequisites and select **OK** to continue.
4. If this is the first time signing in and you did not install Deploy using the wizard available with vCenter, provide the following configuration information when prompted:
   - New password for the administrator account (required)
   - AutoSupport (optional)
   - vCenter server with account credentials (optional)

What's next:

Deploy an ONTAP Select cluster

Related information

- Learn about signing in to Deploy using SSH
- Learn about deploying a 90-day evaluation instance of an ONTAP Select cluster

Deploy an ONTAP Select cluster using the web UI

You can use the web user interface provided with the ONTAP Select Deploy administration utility to deploy a single-node or multi-node ONTAP Select cluster.

When you create an ONTAP Select cluster using the Deploy utility web interface, you are guided through a specific sequence of steps. The exact process varies depending on whether you deploy a single-node or multi-node cluster.



You can also deploy ONTAP Select clusters using the Deploy utility CLI.

Step 1: Prepare for the deployment

Prepare for the deployment to make sure it's successful.

Steps

1. Initial planning.

   Review the Plan and License sections. Based on this review, you can make decisions about the cluster, including:

   - Hypervisor
   - Number of nodes
   - License type
   - Platform size (instance type)
   - ONTAP Select version

2. Prepare the host.

   You must prepare the hypervisor hosts where the ONTAP Select nodes will run and have the needed storage license files based on your licensing model. To view the preparation requirements:

   a. Sign in to the Deploy web UI.

   b. Select  at the top of the page.

   c. Select **Prerequisites**.

   d. Scroll down to review the requirements and select **OK**.

3. Acquire the license files.

   If you plan to deploy the cluster in a production environment, you must acquire the storage license files based on your licensing model.

4. Deploy installation and account credentials.

   Install the Deploy administration utility and perform initial configuration. You need to have the password for the Deploy administrator account that was configured as part of the installation process.

5. Optionally, install earlier ONTAP Select node images.

   By default, the Deploy administration utility contains the most current version of ONTAP Select at the time of release. If you want to deploy clusters using an earlier version of ONTAP Select, you need to add the ONTAP Select image to your Deploy instance.

6. Learn about the "Getting started" launch page.

   The initial page **Getting Started with ONTAP Select Deploy** guides you through the multi-step process of creating a cluster. There are five major steps, including:

   - Add licenses
   - Add hosts to inventory
   - Create a cluster
   - Network precheck

- Deploy the cluster

> You can perform the same steps independently by selecting the tabs at the top of the page (Clusters, Hypervisor Hosts, Administration).

7. Review the network checker.

   If you are deploying a multi-node cluster, you should be familiar with the network checker. You can run the network connectivity checker using the web UI or the CLI.

Step 2: Create a single-node or multi-node cluster

You can use the ONTAP Select Deploy web user interface to deploy a single-node or multi-node ONTAP Select cluster.

Before you begin

Verify that you have installed the Deploy administration and completed the initial configuration (password, AutoSupport, and vCenter).

About this task

An ONTAP Select cluster with one or more nodes is created for a production deployment.

Steps

The steps you follow depend on whether you want to create a single-node cluster or a multi-node cluster. A multi-node cluster can have two, four, six, eight, ten, or twelve nodes.

Single-node cluster

1. Sign in to the Deploy utility through the web interface using the administrator account (admin).
2. If the **Welcome to ONTAP Select** pop-up window is displayed, confirm you have met the configuration prerequisites and select **OK**.
3. If the **Getting Started** cluster launch page is not displayed, select  at the top of the page and select **Getting Started**.
4. On the **Getting Started** page, select **Upload**, and then select a license from your local workstation and select **Open** to upload the license.
5. Select **Refresh** and confirm the license has been added.
6. Select **Next** to add a hypervisor host and then select **Add**.

   You can add the hypervisor host directly or by connecting to a vCenter server. Provide the appropriate host details and credentials as needed.

7. Select **Refresh** and confirm the **Type** value for the host is **ESX** or **KVM**.

   Any account credentials you provide are added to the Deploy credential database.

8. Select **Next** to begin the cluster creation process.
9. In the **Cluster Details** section, provide all the required information describing the cluster and select **Done**.
10. Under **Node Setup**, provide the node management IP address and select the license for the node;

you can upload a new license if needed. You also can change the node name if needed.

11. Provide the **Hypervisor** and **Network** configuration.

    There are three node configurations which define the virtual machine size and available feature set. These instance types are supported by the standard, premium, and premium XL offerings of the purchased license, respectively. The license you select for the node must match or exceed the instance type.

    Select the hypervisor host as well as the management and data networks.

12. Provide the **Storage** configuration and select **Done**.

    You can select the drives based on your platform license level and host configuration.

13. Review and confirm the configuration of the cluster.

    You can change the configuration by selecting  in the applicable section.

14. Select **Next** and provide the ONTAP administrator password.
15. Select **Create Cluster** to begin the cluster creation process and then select **OK** in the pop-up window.

    It can take up to 30 minutes for the cluster to be created.

16. Monitor the multi-step cluster creation process to confirm the cluster is created successfully.

    The page is automatically refreshed at regular intervals.

Multi-node cluster

1. Sign in to the Deploy utility through the web interface using the administrator account (admin).
2. If the **Welcome to ONTAP Select** popup window is displayed, confirm that you have met the configuration prerequisites and select **OK**.
3. If the **Getting Started** cluster launch page is not displayed, select  at the top of the page and select **Getting Started**.
4. On the **Getting Started** page, select **Upload** and select a license from your local workstation and select **Open** to upload the license. Repeat to add additional licenses.
5. Select **Refresh** and confirm the licenses have been added.
6. Select **Next** to add all hypervisor hosts and then select **Add**.

    You can add the hypervisor hosts directly or by connecting to a vCenter server. Provide the appropriate host details and credentials as needed.

7. Select **Refresh** and confirm the **Type** value for the host is **ESX** or **KVM**.

    Any account credentials you provide are added to the Deploy credential database.

8. Select **Next** to begin the cluster creation process.
9. In the **Cluster Details** section, select the desired **Cluster Size**, provide all the required information describing the clusters, and select **Done**.
10. Under **Node Setup**, provide the node management IP addresses and select the licenses for each node; you can upload a new license if needed. You also can change the node names if needed.

11. Provide the **Hypervisor** and **Network** configuration.

    There are three node configurations which define the virtual machine size and available feature set. These instance types are supported by the standard, premium, and premium XL offerings of the purchased license, respectively. The license you select for the nodes must match or exceed the instance type.

    Select the hypervisor hosts as well as the management, data, and internal networks.

12. Provide the **Storage** configuration and select **Done**.

    You can select the drives based on your platform license level and host configuration.

13. Review and confirm the configuration of the cluster.

    You can change the configuration by selecting  in the applicable section.

14. Select **Next** and run the Network Precheck by selecting **Run**. This validates that the internal network selected for ONTAP cluster traffic is functioning correctly.
15. Select **Next** and provide the ONTAP administrator password.
16. Select **Create Cluster** to begin the cluster creation process, and then select **OK** in the popup window.

    It can take up to 45 minutes for the cluster to be created.

17. Monitor the multi-step cluster creation process to confirm that the cluster is created successfully.

    The page is automatically refreshed at regular intervals.

Step 3: Complete the deployment

After deploying the cluster, confirm the ONTAP Select AutoSupport feature is configured and then back up the ONTAP Select Deploy configuration data.



If the cluster creation operation is initiated but fails to complete, the ONTAP administrative password you define might not be applied. If this occurs, you can determine the temporary administrative password for the ONTAP Select cluster by using the following CLI command:

```
(ONTAPdeploy) !/opt/netapp/tools/get_cluster_temp_credentials
--cluster-name my_cluster
```

Initial state of the ONTAP Select cluster after deployment

You should be aware of the initial state of a cluster after it has been deployed and configure the cluster as needed for your environment.

An ONTAP Select cluster has several characteristics after it is created.



Restricting roles and permissions for the ONTAP administrator account can limit ONTAP Select Deploy's ability to manage the cluster. For more information, refer to the Knowledge Base article

OTS Deploy cluster refresh fails with error.

LIFs

There are two types of customer-specified LIFs assigned:

- Cluster management (one per cluster)
- Node management (one per node)

A multi-node cluster has an internal network with autogenerated LIFs.

SVMs

Three SVMs are active:

- Admin SVM
- Node SVM
- System (cluster) SVM

Data SVMs aren't created as part of the ONTAP Select cluster deployment. They must be created by the cluster administrator after deployment. For more information, see Create an SVM.

Aggregates

The root aggregate is created.

Features

All features are licensed and available. Both SnapLock and FabricPool require separate licenses.

Related information

- Types of SVMs contained in a cluster
- ONTAP features licensed by default

Administer

Learn about administering ONTAP Select

After creating an ONTAP Select cluster, you can support the deployment by performing various administrative tasks. There are a few general considerations to be aware of.

In general, the procedures you can perform using the Deploy web interface fall into one of three categories.

Deploy an ONTAP Select cluster

You can deploy a single-node or multi-node cluster. See Deploy an ONTAP Select cluster for more information.

Perform a procedure with an existing ONTAP Select cluster

The administrative procedures are organized in various categories, such as *Security* and *Clusters*.

Perform a procedure on the Deploy utility

There are several procedures specific to Deploy (such as changing the administrator's password).

Administer ONTAP Select

There are many different administrative procedures available as part of supporting ONTAP Select. In addition, there are procedures specific to the Deploy administrative utility. The most important of these procedures are presented below. In general, all use the Deploy web user interface.

You can also use the command line interface to administer ONTAP Select.

Perform additional ONTAP configuration

After an ONTAP Select cluster is deployed, you can configure and manage the cluster just as you would a hardware-based ONTAP system. For example, you can use ONTAP System Manager or the ONTAP CLI to configure the ONTAP Select cluster.

NetApp client software

You can connect to ONTAP Select using the following supported NetApp client software:

- ONTAP System Manager
- Active IQ Unified Manager
- OnCommand Insight
- OnCommand Workflow Automation
- SnapCenter
- Virtual Storage Console for VMware vSphere

To identify the supported versions of the client software, review the Interoperability Matrix Tool. If the client software supports ONTAP 9, then the same version is also supported with ONTAP Select.

The use of SnapCenter and the corresponding plug-ins requires server-based licenses. Storage system licensing of the SnapCenter plug-ins is not currently supported with ONTAP Select.

Any other NetApp client software that is not included in the list is not supported by ONTAP Select.

Possible configuration options

There are several options available when configuring the cluster, including the following:

- Creating the networking configuration
- Laying out your aggregates
- Creating the data storage VMs (SVMs)

Purchased licenses with storage capacity

If you decided not to install the license files with storage capacity as part of deploying the ONTAP Select cluster, you must acquire and install the license files before the grace period expires for clusters running with a purchased license.

Mirrored aggregates

The Deploy administration utility creates data spare disks on each ONTAP Select node from the usable datastore space (such as, Pool0 and Pool1). To implement high availability for your data on a multi-node cluster, you must create a mirrored aggregate using these spares.



High availability takeover is only supported when data aggregates are configured as mirrored aggregates.

Upgrade ONTAP Select nodes

After deploying an ONTAP Select cluster, you can upgrade the ONTAP image at each node in the cluster as needed.



You cannot use the Deploy administration utility to perform upgrades of existing ONTAP Select nodes. The Deploy utility can only be used to create new ONTAP Select clusters.

General procedure

At a high level, you should use the following steps to upgrade an existing ONTAP Select node.

Steps

1. Access the NetApp Support Site downloads page.
2. Scroll down and select **ONTAP Select Image**.
3. Select the desired release of the install image.
4. Review the End User License Agreement (EULA) and select **Accept & Continue**.
5. Select and download the appropriate **ONTAP Select Image Upgrade** package. Respond to all prompts as needed.

   Review the release notes for additional information and any required procedures before upgrading an ONTAP Select node.

6. Upgrade the ONTAP Select node using the standard ONTAP upgrade procedures with the ONTAP Select upgrade file. For information on supported upgrade paths, refer to the Supported ONTAP upgrade paths.

Revert an ONTAP Select node

You cannot revert an ONTAP Select node to a version prior to the one on which it was originally installed. For example:

ONTAP Select 9.16.1 is initially installed

You can upgrade the node to version 9.17.1 and then revert back to version 9.16.1 if needed.

ONTAP Select 9.17.1 is initially installed

Reversion is not possible because there was no earlier version previously installed.

Use the VMXNET3 network driver

VMXNET3 is the default network driver included with new cluster deployments on VMware ESXi. If you upgrade an existing ONTAP Select node running ONTAP Select 9.4 or earlier, the network driver is not automatically upgraded. You must manually upgrade to VMXNET3. You should contact NetApp support for assistance with the upgrade.

Related information

ONTAP upgrade overview

Manage ONTAP Select diagnostics and support

There are several related diagnostic and support tasks you can perform as part of administering ONTAP Select.

Configure the Deploy system

You should set the basic system configuration parameters that affect how the Deploy utility operates.

About this task

The Deploy configuration data is used by AutoSupport.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Administration** tab at the top of the page.
3. Select **Settings & AutoSupport**, then select .
4. Provide the configuration data as appropriate for your environment and select **Modify**.

   If you use a proxy server, you can configure the proxy URL as follows:
   `http://USERNAME:PASSWORD@<FQDN|IP>:PORT`

   Example
   `http://user1:mypassword@proxy.company-demo.com:80`

Display the ONTAP Select Deploy event messages

The ONTAP Select Deploy utility includes an event logging facility that provides information about the activity of the system. You should view the contents of the event log to debug any issues or when directed to do so by support.

About this task

You can filter the list of event messages based on several characteristics, including:

- Status
- Type
- Category
- Instance
- Time
- Description

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Administration** tab at the top of the page.
3. Select **Events & Jobs**, then select **Events**.
4. Optionally, select **Filter** and create a filter to limit the event messages displayed.

Enable AutoSupport

You can enable and disable the AutoSupport feature as needed.

About this task

AutoSupport is the primary troubleshooting tool used by NetApp in supporting ONTAP Select. Therefore, you should not disable AutoSupport unless absolutely necessary. If you do disable AutoSupport, data is still collected but not transmitted to NetApp.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Administration** tab at the top of the page.
3. Select **Settings & AutoSupport** and then select ⋮.
4. Enable or disable the AutoSupport feature as needed.

Generate and download an AutoSupport package

ONTAP Select includes the ability to generate an AutoSupport package. You should generate a package to debug any issues or when directed to do so by support.

About this task

You can generate the following AutoSupport packages under the direction and guidance of NetApp support:

- Deploy logs
  Log files created by the ONTAP Select Deploy utility
- Troubleshooting
  Troubleshooting and debugging information about the hypervisor hosts and ONTAP Select nodes
- Performance
  Performance information about the hypervisor hosts and ONTAP Select nodes

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Administration** tab at the top of the page.
3. Select **Settings & AutoSupport** and then select ⋮.
4. Select **Generate**.
5. Select the type and provide a description for the package; you can optionally provide a case number.
6. Select **Generate**.

   Each AutoSupport package is assigned a unique sequence identification number.

7. Optionally, select the correct package under **AutoSupport History**, then select the download icon to save the AutoSupport file to your local workstation.

Secure an ONTAP Select deployment

There are several related tasks you can perform as part of securing an ONTAP Select deployment.

Change the Deploy administrator password

You can change the password for the Deploy virtual machine administrator account as needed using the web user interface.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the figure icon at the top right of the page, then select **Change Password**.
3. Provide the current and new password as prompted and select **Submit**.

Add a management server account

You can add a management server account to the Deploy credential store database.

Before you begin

You should be familiar with the types of credentials and how they are used by ONTAP Select Deploy.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Administration** tab at the top of the page.
3. Select **Management Servers** > **Add vCenter**.
4. Enter the following information and select **Add**.

In this field …	Do the following …
Name/IP Address	Provide the domain name or IP address of the vCenter server.
Username	Enter the account user name to access vCenter.
Password	Enter the password for the associated user name.

5. Optionally, select ⋮ after the new management server is added, then select one of the following:
    - Update credentials
    - Verify credentials
    - Remove management server

Configure MFA

Beginning with ONTAP Select 9.13.1, multifactor authentication (MFA) is supported for the ONTAP Select Deploy administrator account:

- ONTAP Select Deploy CLI MFA login using YubiKey Personal Identity Verification (PIV) or Fast IDentity Online (FIDO2) authentication
- ONTAP Select Deploy CLI MFA login using ssh-keygen

ONTAP Select Deploy CLI MFA login using YubiKey PIV or FIDO2 authentication

YubiKey PIV

Configure the YubiKey PIN and generate or import the Remote Support Agent (RSA) or Elliptic Curve Digital Signature Algorithm (ECDSA) private key and certificate with the steps in TR-4647: Multifactor authentication in ONTAP.

- For Windows: The **YubiKey PIV Client configuration for Windows** section of the technical report.
- For MacOS: The **YubiKey PIV client configuration For MAC OS and Linux** section of the technical report.

FIDO2

If you choose to opt for YubiKey FIDO2 authentication, configure the YubiKey FIDO2 PIN using the YubiKey Manager and generate the FIDO2 key with a PuTTY-CAC (Common Access Card) for Windows or ssh-keygen for MacOS. The steps to do this are in the technical report TR-4647: Multifactor authentication in ONTAP.

- For Windows: The **YubiKey FIDO2 client configuration for Windows** section of the technical report.
- For MacOS: The **YubiKey FIDO2 client configuration For Mac OS and Linux** section of the technical report.

Obtain the YubiKey PIV or FIDO2 public key

Obtaining the public key depends on whether you're a Windows or MacOS client, and if you are using PIV or FIDO2.

For Windows:

- Export the PIV public key using the **Copy to Clipboard** feature under SSH → Certificate as described in the section **Configuring the Windows PuTTY-CAC SSH Client for YubiKey PIV Authentication** on page 16 of TR-4647.
- Export the FIDO2 public key using the **Copy to Clipboard** feature under SSH → Certificate as described in the section **Configuring the Windows PuTTY-CAC SSH Client for YubiKey FIDO2 Authentication** on page 30 of TR-4647.

For MacOS:

- The PIV public key should be exported using the `ssh-keygen -e` command as described in the section **Configure the Mac OS or Linux SSH Client for YubiKey PIV authentication** on page 24 of TR-4647.
- The FIDO2 public key is in the `id_ecdsa_sk.pub` file or `id_edd519_sk.pub` file, depending on whether you use ECDSA or EDD519, as described in the section **Configure the MAC OS or Linux SSH client for YubiKey FIDO2 authentication** on page 39 of TR-4647.

Configure the public key in ONTAP Select Deploy

SSH is used by the administrator account for the public key authentication method. The command used is the same whether the authentication method is the standard SSH public key authentication or YubiKey PIV or FIDO2 authentication.

For hardware-based SSH MFA, the authentication factors in addition to the public key configured on ONTAP Select Deploy are as follows:

- The PIV or FIDO2 PIN
- Possession of the YubiKey hardware device. For FIDO2, this is confirmed by physically touching the YubiKey during the authentication process.

Before you begin

Set the PIV or FIDO2 public key which is configured for the YubiKey. The ONTAP Select Deploy CLI command `security publickey add -key` is the same for PIV or FIDO2 and the public key string is different.

The public key is obtained from:

- The **Copy to Clipboard** function for PuTTY-CAC for PIV and FIDO2 (Windows)
- Exporting the public key in an SSH compatible format using the `ssh-keygen -e` command for PIV
- The public key file located in the `~/.ssh/id_***_sk.pub` file for FIDO2 (MacOS)

Steps

1. Find the generated key in the `.ssh/id_***.pub` file.
2. Add the generated key to ONTAP Select Deploy using the `security publickey add -key <key>` command.

```
(ONTAPdeploy) security publickey add -key "ssh-rsa <key>
user@netapp.com"
```

3. Enable MFA Authentication with the `security multifactor authentication enable` command.

```
(ONTAPdeploy) security multifactor authentication enable
MFA enabled Successfully
```

Log in to ONTAP Select Deploy using YubiKey PIV Authentication over SSH

You can log in to ONTAP Select Deploy using YubiKey PIV Authentication over SSH.

Steps

1. After the YubiKey token, the SSH client, and ONTAP Select Deploy are configured, you can use MFA YubiKey PIV authentication over SSH.
2. Log in to ONTAP Select Deploy. If you are using the Windows PuTTY-CAC SSH client, a dialog will pop-up prompting you to enter your YubiKey PIN.
3. Log in from your device with the YubiKey connected.

Example output:

```
login as: admin
Authenticating with public key "<public_key>"
Further authentication required
<admin>'s password:

NetApp ONTAP Select Deploy Utility.
Copyright (C) NetApp Inc.
All rights reserved.

Version: NetApp Release 9.13.1 Build:6811765 08-17-2023 03:08:09

(ONTAPdeploy)
```

ONTAP Select Deploy CLI MFA login using ssh-keygen

The `ssh-keygen` command is a tool for creating new authentication key pairs for SSH. The key pairs are used for automating logins, single sign-on, and for authenticating hosts.

The `ssh-keygen` command supports several public key algorithms for authentication keys.

- The algorithm is selected with the `-t` option
- The key size is selected with the `-b` option

Example output:

```
ssh-keygen -t ecdsa -b 521
ssh-keygen -t ed25519
ssh-keygen -t ecdsa
```

Steps

1. Find the generated key in the `.ssh/id_***.pub` file.
2. Add the generated key to ONTAP Select Deploy using the `security publickey add -key <key>` command.

```
(ONTAPdeploy) security publickey add -key "ssh-rsa <key>
user@netapp.com"
```

3. Enable MFA Authentication with the `security multifactor authentication enable` command.

```
(ONTAPdeploy) security multifactor authentication enable
MFA enabled Successfully
```

4. Log in to the ONTAP Select Deploy system after enabling MFA. You should receive an output similar to the following example.

```
[<user ID> ~]$ ssh <admin>
Authenticated with partial success.
<admin>'s password:

NetApp ONTAP Select Deploy Utility.
Copyright (C) NetApp Inc.
All rights reserved.

Version: NetApp Release 9.13.1 Build:6811765 08-17-2023 03:08:09

(ONTAPdeploy)
```

Migrate from MFA to single-factor authentication

MFA can be disabled for the Deploy administrator account using the following methods:

- If you can log in to the Deploy CLI as an administrator using Secure Shell (SSH), disable MFA by running the `security multifactor authentication disable` command from the Deploy CLI.

```
(ONTAPdeploy) security multifactor authentication disable
MFA disabled Successfully
```

- If you cannot log in to the Deploy CLI as an administrator using SSH:
  1. Connect to the Deploy virtual machine (VM) video console through vCenter or vSphere.
  2. Log in to the Deploy CLI using the administrator account.
  3. Run the `security multifactor authentication disable` command.

```
Debian GNU/Linux 11 <user ID> tty1

<hostname> login: admin
Password:

NetApp ONTAP Select Deploy Utility.
Copyright (C) NetApp Inc.
All rights reserved.
```

```
Version: NetApp Release 9.13.1 Build:6811765 08-17-2023 03:08:09

(ONTAPdeploy) security multifactor authentication disable
MFA disabled successfully

(ONTAPdeploy)
```

- The administrator can delete the public key with:
  `security publickey delete -key`

Confirm connectivity among the ONTAP Select nodes

You can test the network connectivity among two or more ONTAP Select nodes on the internal cluster network. You typically run this test before a multi-node cluster is deployed to detect issues that might cause the operation to fail.

Before you begin

All the ONTAP Select nodes included in the test must be configured and powered on.

About this task

Each time you start a test, a new process run is created in the background and assigned a unique run identifier. Only one run can be active at a time.

The test has two modes that control its operation:

- Quick
  This mode performs a basic non-disruptive test. A PING test is performed, along with a test of the network MTU size and the vSwitch.
- Extended
  This mode performs a more comprehensive test over all the redundant network paths. If you run this on an active ONTAP Select cluster, the performance of the cluster can be impacted.



It is recommended that you always perform a quick test before creating a multi-node cluster. After the quick test completes successfully, you can optionally perform an extended test based on your production requirements.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Administration** tab at the top of the page, then select **Network Checker**.
3. Select **Start New Run**, then select the hosts and networks for the HA pair

   You can add and configure additional HA pairs as needed.
4. Select **Start** to begin the network connectivity test.

Administer the ONTAP Select Deploy mediator services

Each ONTAP Select two-node cluster is monitored by the mediator service, which assists in managing the HA capability shared by the nodes.

View the status of the mediator service

You can view the status of the mediator service with respect to each of the two-node clusters defined to the ONTAP Select Deploy utility.

About this task

You can view the configuration of each mediator, including the current status, the two ONTAP Select nodes, and the iSCSI target where the HA control information is stored. Hover over the objects on the page to display detailed information.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Administration** tab at the top of the page, then select **Mediators**.
3. Optionally, select **Filter** to customize your view of the two-node clusters monitored by the mediator service.

Clusters

Manage ONTAP Select clusters

There are several related tasks you can perform to administer an ONTAP Select cluster.

Move an ONTAP Select cluster offline and online

After you've created a cluster, you can move it offline and online as needed.

Before you begin

After a cluster is created it is initially in the online state.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Clusters** tab at the top of the page and select the desired cluster from the list.
3. Select ⋮ on the right of the cluster and select **Take Offline**.

   If the offline option is not available, the cluster is already in the offline state.
4. Select **Yes** in the popup window to confirm the request.
5. Select **Refresh** occasionally to confirm the cluster is offline.
6. To bring the cluster back online, select ⋮ and select **Bring Online**.
7. Select **Refresh** occasionally to confirm the cluster is online.

Delete an ONTAP Select cluster

You can delete an ONTAP Select cluster when it is no longer needed.

Before you begin

The cluster must be in the offline state.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Clusters** tab at the top of the page and select the desired cluster from the list.
3. Select ⋮ on the right of the cluster and select **Delete**.

   If the delete option is not available, then the cluster is not in an offline state.

4. Select **Refresh** occasionally to confirm the cluster is removed from the list.

Refresh the Deploy cluster configuration

After creating an ONTAP Select cluster, you can make changes to the cluster or the virtual machine configuration outside of the Deploy utility using the ONTAP or hypervisor administration tools. The configuration of a virtual machine can also change after it is migrated.

When these changes to the cluster or virtual machine occur, the Deploy utility configuration database is not automatically updated and can become out of sync with the state of the cluster. You should perform a cluster refresh in these and other situations to update the Deploy database based on the current state of the cluster.

Before you begin

Required information

You must have the current configuration information for the cluster, including:

- ONTAP administrator credentials
- Cluster management IP address
- Names of the nodes in the cluster

Stable cluster state

The cluster must be in a stable state. You cannot refresh a cluster when it is in the process of being created or deleted, or when it is in the *create_failed* or *delete_failed* state.

After a VM migration

After a virtual machine running ONTAP Select has been migrated, you must create a new host using the Deploy utility before performing a cluster refresh.

About this task

You can perform a cluster refresh to update the Deploy configuration database using the web user interface.

> ⓘ Instead of using the Deploy UI, you can use the cluster refresh command in the Deploy CLI shell to refresh a cluster.

Cluster and virtual machine configuration

Some of the configuration values that can change and cause the Deploy database to become out of sync include:

- Cluster and node names
- ONTAP network configuration

- ONTAP version (after an upgrade)
- Virtual machine names
- Host network names
- Storage pool names

Cluster and node states

An ONTAP Select cluster or node can be in a state that prevents it from operating properly. You should perform a cluster refresh operation to correct the following conditions:

- Node in *unknown* state
  An ONTAP Select node can be in the *unknown state* for several reasons, including the node is not found.
- Cluster in *degraded* state
  If a node is powered off, it might still appear to be online in the Deploy utility. In this situation, the cluster is in a *degraded* state.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Clusters** tab at the top left of the page and select the desired cluster from the list.
3. Select ⋮ on the right side of the page and select **Cluster Refresh**.
4. Under **Cluster Credentials**, provide the ONTAP administrator password for the cluster.
5. Select **Refresh**.

After you finish

If the operation is successful, the field *Last Refresh* is updated. You should back up the Deploy configuration data after the cluster refresh operation has completed.

Expand or contract an ONTAP Select cluster on ESXi or KVM hosts

Increase or decrease the cluster size of an existing ONTAP Select cluster for ESXi and KVM hypervisor hosts. For both host types, you can increase and decrease the cluster size in increments between four and twelve nodes.

The following cluster expansions and contractions aren't supported for ESXi and KVM hosts:

- Expansions from one or two-node clusters to six, eight, ten, or twelve-node clusters.
- Contractions from six, eight, ten, or twelve-node clusters to one or two-node clusters.

> ⓘ To change the number of nodes in a cluster to a size that isn't supported by cluster expansion or contraction, you need to perform the following tasks:
>
> 1. Deploy a new multi-node cluster by using the CLI or the web UI provided with the ONTAP Select Deploy administration utility.
> 2. If applicable, migrate the data to the new cluster by using SnapMirror replication.

You initiate the cluster expansion and contraction procedures from ONTAP Select Deploy using the CLI, API, or web interface.

Hardware and storage considerations

The cluster expansion and contraction feature is supported on the following KVM and ESXi hypervisor hosts.

ESXi

Beginning with ONTAP Select 9.15.1, cluster expansion and contraction is supported on ESXi hypervisor hosts.

Cluster expansion and contraction is supported for the following ESXi hypervisor versions:

- ESXi 9.0
- ESXi 8.0 U3
- ESXi 8.0 U2
- ESXi 8.0 U1
- ESXi 8.0 GA

KVM

Beginning with ONTAP Select 9.17.1, cluster expansion and contraction is supported on KVM hypervisor hosts.

Cluster expansion and contraction is supported for the following KVM hypervisor versions:

- Red Hat Enterprise Linux (RHEL) 64-bit 10.1, 10.0, 9.7, 9.6, 9.5, 9.4, 9.3, 9.2, 9.1, 9.0, 8.8, 8.7, and 8.6
- Rocky Linux 10.1, 10.0, 9.7, 9.6, 9.5, 9.4, 9.3, 9.2, 9.1, 9.0, 8.9, 8.8, 8.7, and 8.6

Expand the cluster

Use the cluster expansion feature to increase the size of an existing ONTAP Select cluster.

You can increase the size of an existing cluster on an ESXi or KVM host in the following increments:

- From four nodes to six, eight, ten, or twelve nodes
- From six nodes to eight, ten, or twelve nodes
- From eight nodes to ten or twelve nodes
- From ten to twelve nodes

About this task

In preparation for cluster expansion, new ESXi and KVM hosts are added to the inventory and the details of the new nodes are assigned. Before starting the cluster expansion process, a network precheck verifies the selected internal network.

Before you begin

- When deploying a multi-node cluster, you should be familiar with the network connectivity checker. You can run the network connectivity checker using the web UI or the CLI.
- Verify that you have the license details for the new nodes.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Cluster** tab at the top of the page and select the desired cluster from the list.
3. On the cluster details page, select the gear icon at the right of the page and select **Expand Cluster**.
4. Navigate to the **HA Pair 4** section.
5. Choose the following high availability (HA) pair configuration details for the fourth HA pair:
    - Instance type
    - Node names
    - Associated hypervisor hosts
    - Node IP addresses
    - Licenses
    - Networking configuration
    - Storage configuration (RAID type and storage pools)
6. Select **Save HA Pair** to save the configuration details.
7. Provide the ONTAP credentials, and select **Expand Cluster**.
8. Select **Next** and run the network pre-check by selecting **Run**.

    The network pre-check validates that the internal network selected for the ONTAP cluster traffic is functioning correctly.

9. Select **Expand Cluster** to begin the cluster expansion process, and then select **OK** in the dialog box.

    It can take up to 45 minutes for the cluster to be expanded.

10. Monitor the multi-step cluster expansion process to confirm that the cluster expanded successfully.
11. Refer to the **Events** tab for periodic updates on the operation's progress. The page is automatically refreshed at regular intervals.

After you finish

After expanding the cluster, you should back up the ONTAP Select Deploy configuration data.

Contract the cluster

Use the cluster contraction feature to decrease the size of an existing ONTAP Select cluster.

You can decrease the size of an existing cluster on an ESXi or KVM host in the following increments:

- From twelve nodes to ten, eight, six, or four nodes
- From ten nodes to eight, six, or four nodes
- From eight to six or four nodes
- From six to four nodes

About this task

The desired HA pair of nodes in the cluster are selected to prepare for cluster contraction during the procedure.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Cluster** tab at the top of the page and select the desired cluster from the list.
3. On the cluster details page, select the gear icon at the right of the page, then select **Contract Cluster**.
4. Select the HA Pair configuration details for any HA Pair you want to remove and provide the ONTAP credentials, then select **Contract Cluster**.

   It can take up to 30 minutes for the cluster to be contracted.
5. Monitor the multi-step cluster contraction process to confirm that the cluster contracted successfully.
6. Refer to the **Events** tab for periodic updates on the operation's progress. The page is automatically refreshed at regular intervals.

Nodes and hosts

Access an ONTAP Select video console

You can access the video console of the hypervisor virtual machine where ONTAP Select is running.

About this task

You might need to access the virtual machine console to troubleshoot an issue or when asked to do so by NetApp support.

Steps

1. Access the vSphere client and sign in.
2. Navigate to the appropriate location in the hierarchy to locate the ONTAP Select virtual machine.
3. Right-select the virtual machine and select **Open Console**.

Resize ONTAP Select cluster nodes

After deploying an ONTAP Select cluster, you can upgrade the hypervisor instance type of the nodes using the Deploy administration utility.

> You can perform the cluster nodes resizing operation when using the Capacity Tiers licensing model and the Capacity Pools licensing model.

> Resizing to the large instance type is only supported on ESXi.

Before you begin

The cluster must be in the online state.

About this task

This task describes how to use the Deploy web user interface. You can also use the Deploy CLI to perform the instance resizing. Regardless of which interface you use, the time needed for the resizing operation can vary significantly based on several factors and may take an extended amount of time to complete. You can only resize a node to a larger size.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Cluster** tab at the top of the page, then select the desired cluster from the list.
3. On the cluster details page, select the gear icon at the right of the page and select **Instance Resize**.
4. Select the **Instance Type** and provide the ONTAP credentials, then select **Modify**.

After you finish

You must wait for the resize operation to complete.

Replace failed software RAID drives for ONTAP Select

When a drive using software RAID fails, ONTAP Select assigns a spare drive if one is available and starts the rebuild process automatically. This is similar to how ONTAP works on FAS and AFF. However if no spare drive is available, you need to add one to the ONTAP Select node.

Both the removal of the failed drive and the addition of a new drive (marked as a spare) must be performed through ONTAP Select Deploy. Attaching a drive to the ONTAP Select VM using vSphere is not supported.

Identify the failed drive

When a drive fails you need to use the ONTAP CLI to identify the failed disk.

KVM

Before you begin

You must have the VM ID of the ONTAP Select virtual machine, as well as the ONTAP Select and ONTAP Select Deploy administrator account credentials.

About this task

You should only use this procedure when the ONTAP Select node is running on KVM and configured to use software RAID.

Steps

1. At the ONTAP Select CLI, identify the disk to be replaced:
   a. Identify the disk by serial number, UUID, or target address in the virtual machine.

      ```
      disk show -fields serial,vmdisk-target-address,uuid
      ```

   b. Optionally, display a complete list of the spare disk capacity with the partitioned disks.
      storage aggregate show-spare-disks
2. At the Linux command line interface, locate the disk.
   a. Examine the system devices, searching for the disk serial number or UUID (disk name):

      ```
      find /dev/disk/by-id/<SN|ID>
      ```

b. Examine the virtual machine configuration, searching for the target address:

```
virsh dumpxml VMID
```

ESXi

Steps

1. Sign in to the ONTAP CLI using the administrator account.
2. Identify the disk drive that failed.

```
<cluster name>::> storage disk show -container-type broken
Usable Disk Container Container
Disk Size Shelf Bay Type Type Name Owner
---------------- ---------- ----- --- ------- ----------- ---------
--------
NET-1.4 893.3GB - - SSD broken - sti-rx2540-346a'
```

Remove the failed drive

After you identify the drive that failed, remove the disk.

KVM using Deploy

You can detach a disk from a KVM host as part of replacing the disk or when it is no longer needed.

Before you begin

You must have the ONTAP Select and ONTAP Select Deploy administrator account credentials.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Clusters** tab at the top of the page and select the desired cluster from the list.
3. Select **+** next to the desired HA pair or node.

   If the option is disabled, Deploy is currently refreshing the storage information.

4. Select **Edit Storage** on the **Edit Node Storage** page.
5. Deselect the disks to be detached from the node, enter the ONTAP administrator credentials, and select **Edit Storage** to apply the changes.
6. Select **Yes** to confirm the warning in the popup window.
7. Select the **Events** tab for the cluster to monitor and confirm the detach operation.

   You can remove the physical disk from the host if it is no longer needed.

KVM using CLI

After you identify the disk, follow the steps below.

Steps

1. Detach the disk from the virtual machine:
   a. Dump the configuration.

```
virsh dumpxml VMNAME > /PATH/disk.xml
```

   b. Edit the file and remove everything except the disk to be detached from the virtual machine.

      The target address for the disk should correspond to the vmdisk-target-address field in ONTAP.

```
<disk type='block' device='lun'>
  <driver name='qemu' type='raw' cache='directsync'/>
  <source dev='/dev/disk/by-id/ata-
Micron_5100_MTFDDAK960TCC_171616D35277'/>
  <backingStore/>
  <target dev='sde' bus='scsi'/>
  <alias name='scsi0-0-0-4'/>
  <address type='drive' controller='0' bus='0' target='0'
unit='4'/>
</disk>
```

   c. Detach the disk.

```
virsh detach-disk --persistent /PATH/disk.xml
```

2. Replace the physical disk:

   You can use a utility such as `ledctl locate=` to locate the physical disk if needed.

   a. Remove the disk from the host.
   b. Select a new disk and install it in the host if necessary.
3. Edit the original disk configuration file and add the new disk.

   You should update the disk path and any other configuration information as needed.

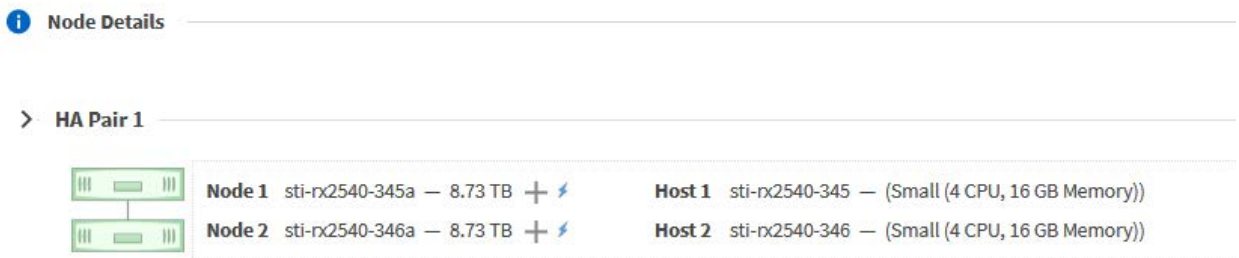
```
<disk type='block' device='lun'>
  <driver name='qemu' type='raw' cache='directsync'/>
  <source dev='/dev/disk/by-id/ata-
Micron_5100_MTFDDAK960TCC_171616D35277'/>
  <backingStore/>
  <target dev='sde' bus='scsi'/>
  <alias name='scsi0-0-0-4'/>
  <address type='drive' controller='0' bus='0' target='0' unit='4'/>
```

```
</disk>
```

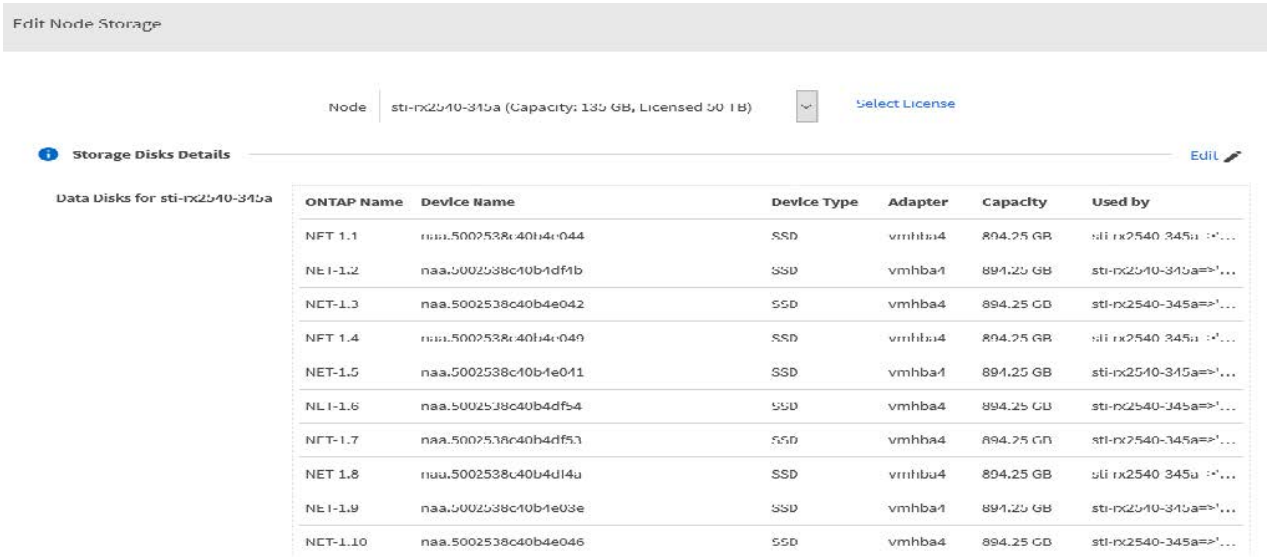
ESXi

Steps

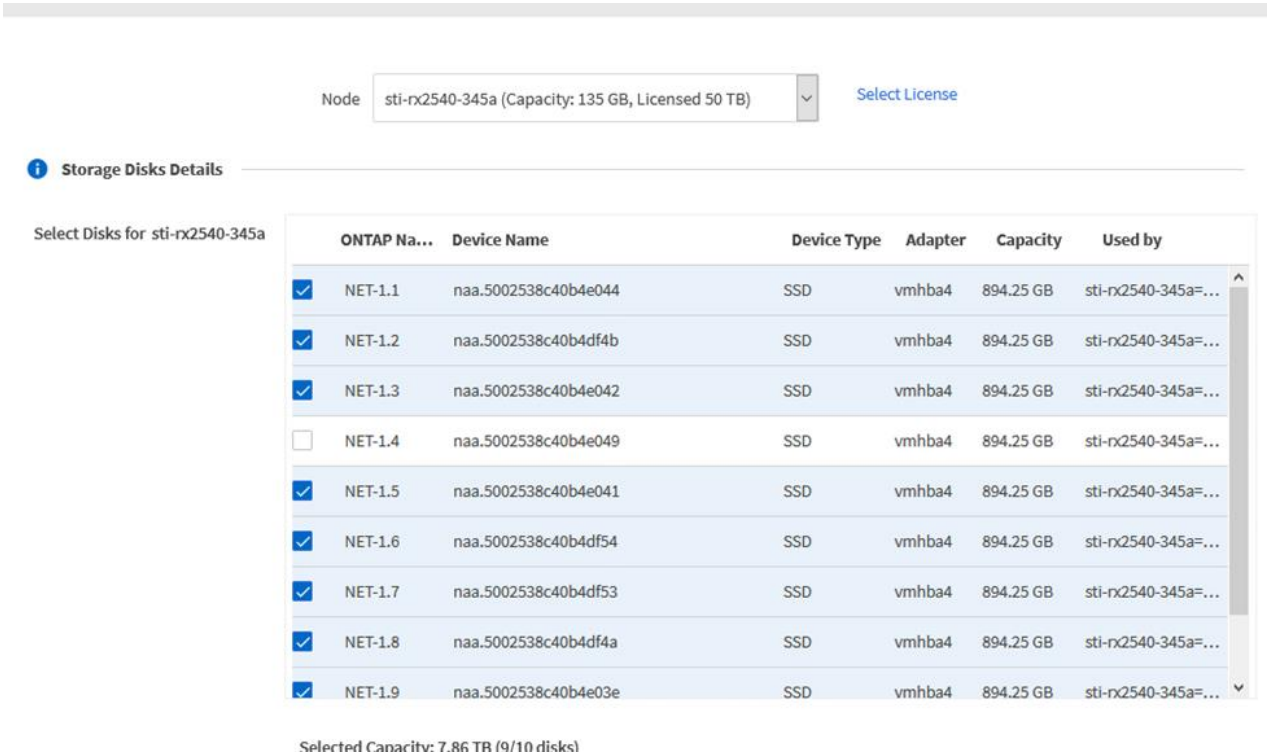
1. Sign in to the Deploy web user interface using the administrator account.
2. Select the **Clusters** tab and select the relevant cluster.



3. Select **+** to expand the storage view.



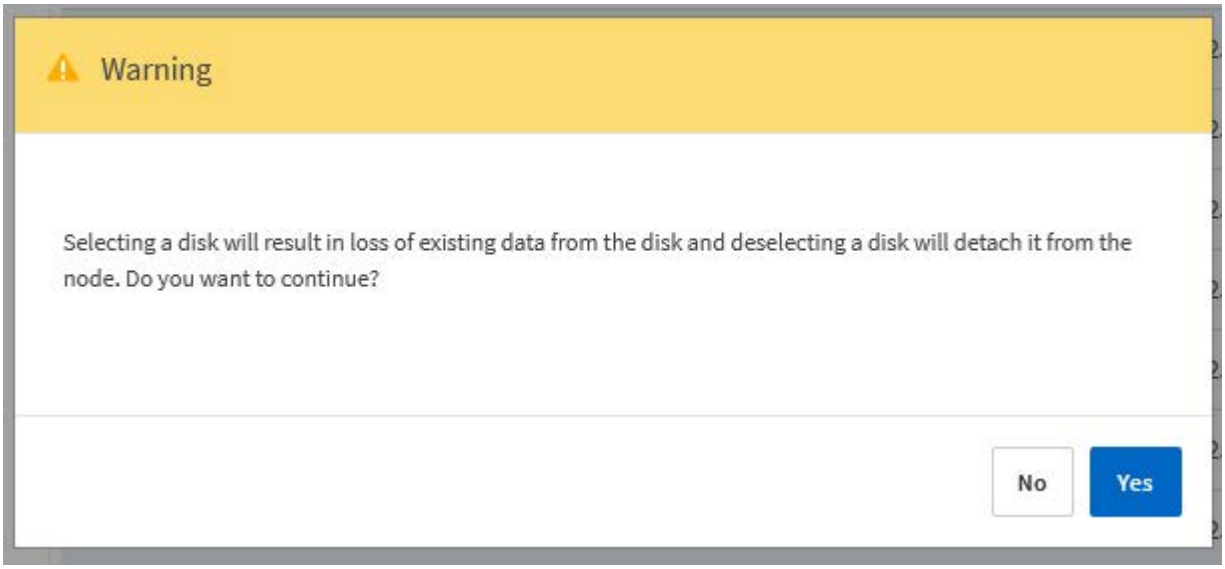
4. Select **Edit** to make changes to the attached disks and uncheck the failed drive.



5. Provide the cluster credentials and select **Edit Storage**.



6. Confirm the operation.



Add the new spare drive

After you remove the failed drive, add the spare disk.

KVM using Deploy

Attaching a disk using Deploy

You can attach a disk to a KVM host as part of replacing a disk or to add more storage capacity.

Before you begin

You must have the ONTAP Select and ONTAP Select Deploy administrator account credentials.

The new disk must be physically installed on the KVM Linux host.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Clusters** tab at the top of the page and select the desired cluster from the list.
3. Select **+** next to the desired HA pair or node.

   If the option is disabled, Deploy is currently refreshing the storage information.

4. Select **Edit Storage** on the **Edit Node Storage** page.
5. Select the disks to be attached to the node, enter the ONTAP administrator credentials, and select **Edit Storage** to apply the changes.
6. Select the **Events** tab to monitor and confirm the attach operation.
7. Examine the node storage configuration to confirm that the disk is attached.

KVM using CLI

After you identify and remove the failed drive, you can attach a new drive.

Steps

1. Attach the new disk to the virtual machine.

   ```
   virsh attach-disk --persistent /PATH/disk.xml
   ```

Results

The disk is assigned as a spare and is available to ONTAP Select. It may take a minute or longer for the disk to become available.

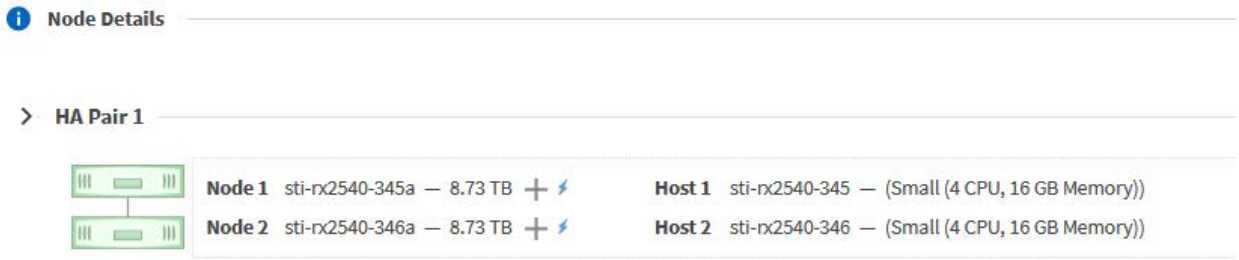
After you finish

Because the node configuration has changed, you should perform a cluster refresh operation using the Deploy administration utility.

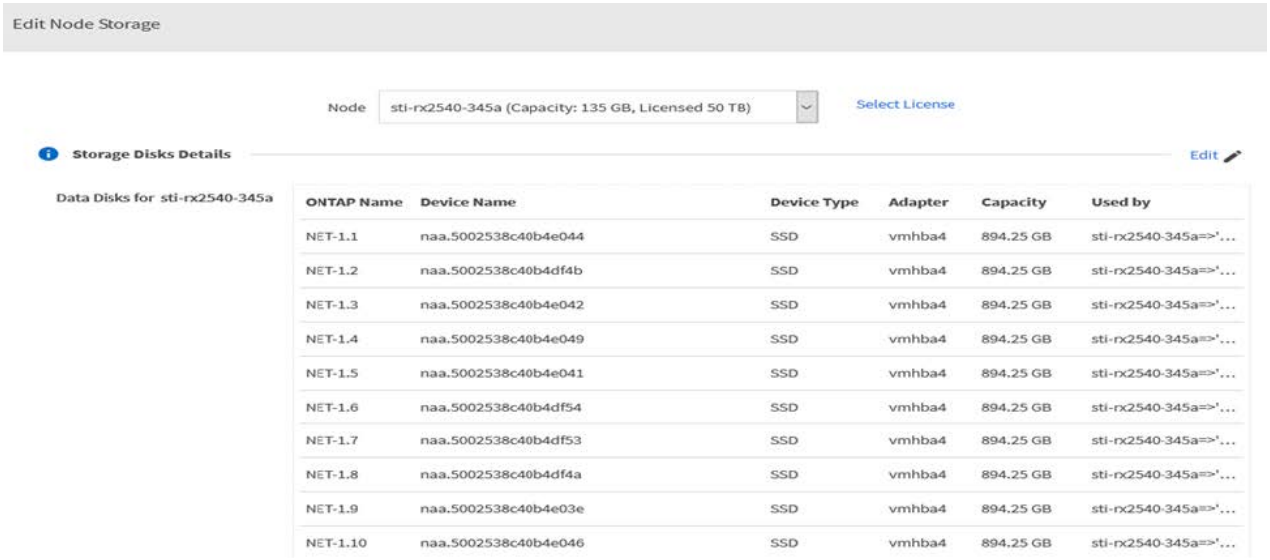
ESXi

Steps

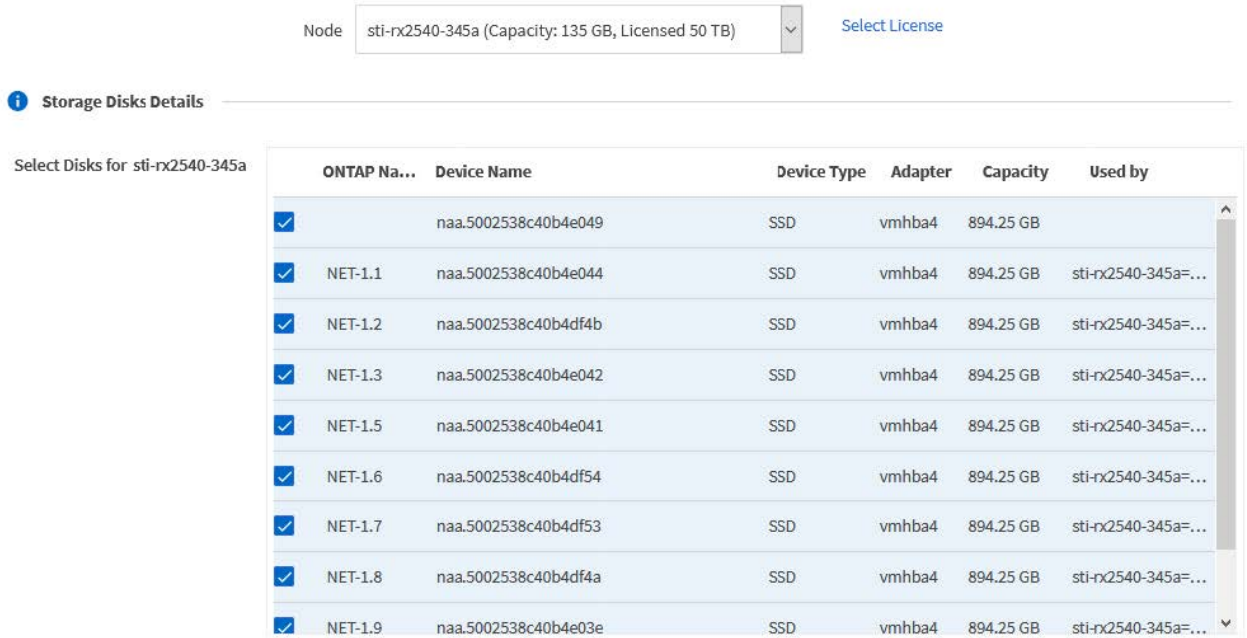
1. Sign in to the Deploy web user interface using the administrator account.
2. Select the **Clusters** tab and select the relevant cluster.



3. Select **+** to expand the storage view.



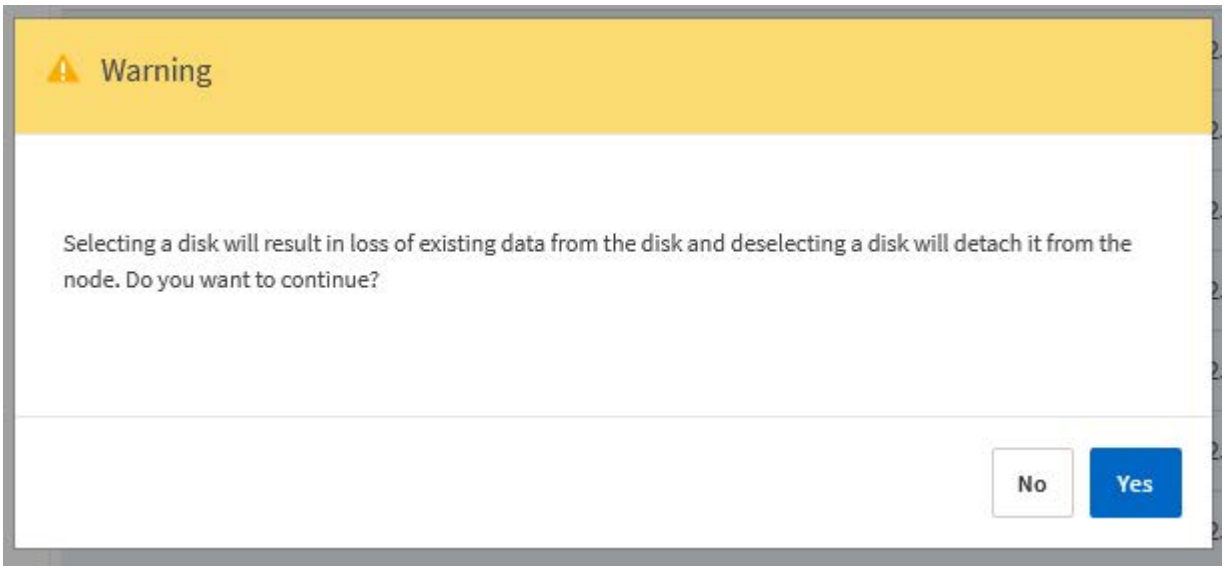
4. Select **Edit** and confirm that the new drive is available and select it.



5. Provide the cluster credentials and select **Edit Storage**.



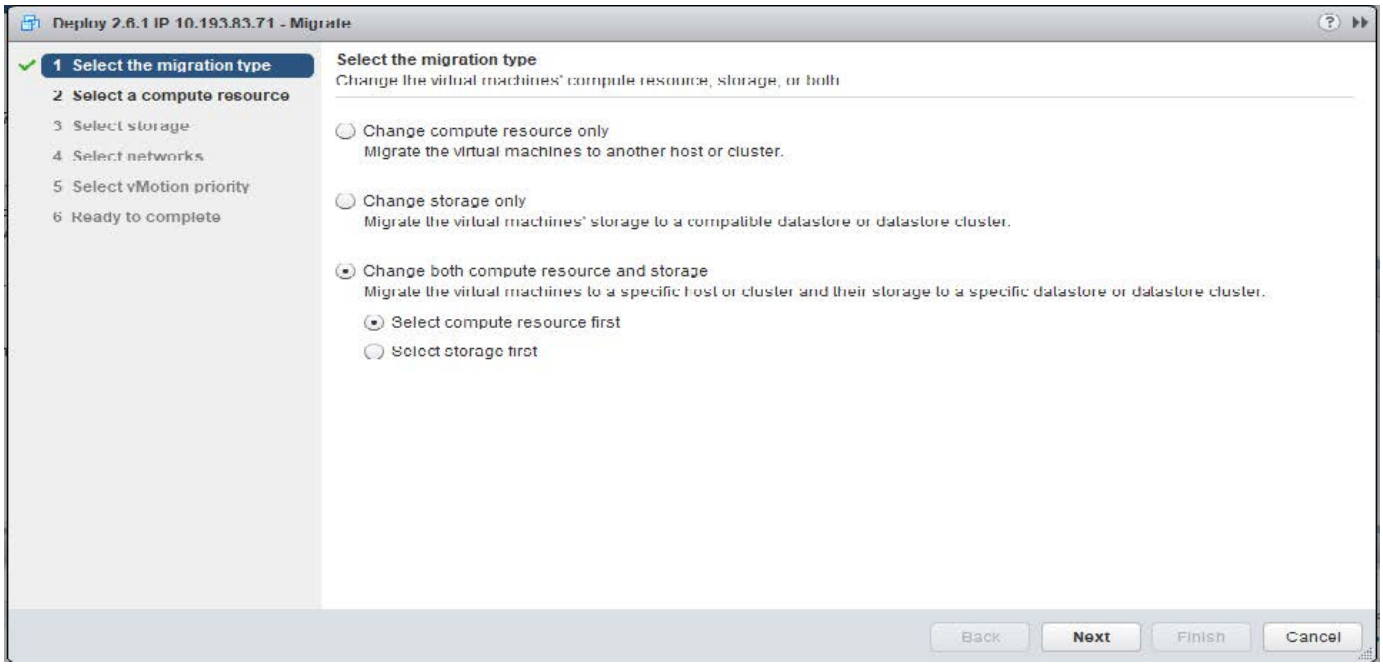
6. Confirm the operation.



Upgrade ONTAP Select node to VMFS6 using Storage vMotion

VMware does not support an in-place upgrade from VMFS 5 to VMFS 6. You can use Storage vMotion to transition from a VMFS 5 datastore to a VMFS 6 datastore for an existing ONTAP Select node.

For ONTAP Select virtual machines, Storage vMotion can be used for single-node and multi-node clusters. It can be used for both storage-only as well as compute and storage migrations.



Before you begin

Make sure the new host can support the ONTAP Select node. For example, if a RAID controller and DAS storage are used on the original host, a similar configuration should exist on the new host.

Severe performance issues can result if the ONTAP Select VM is rehosted in an unsuitable environment.

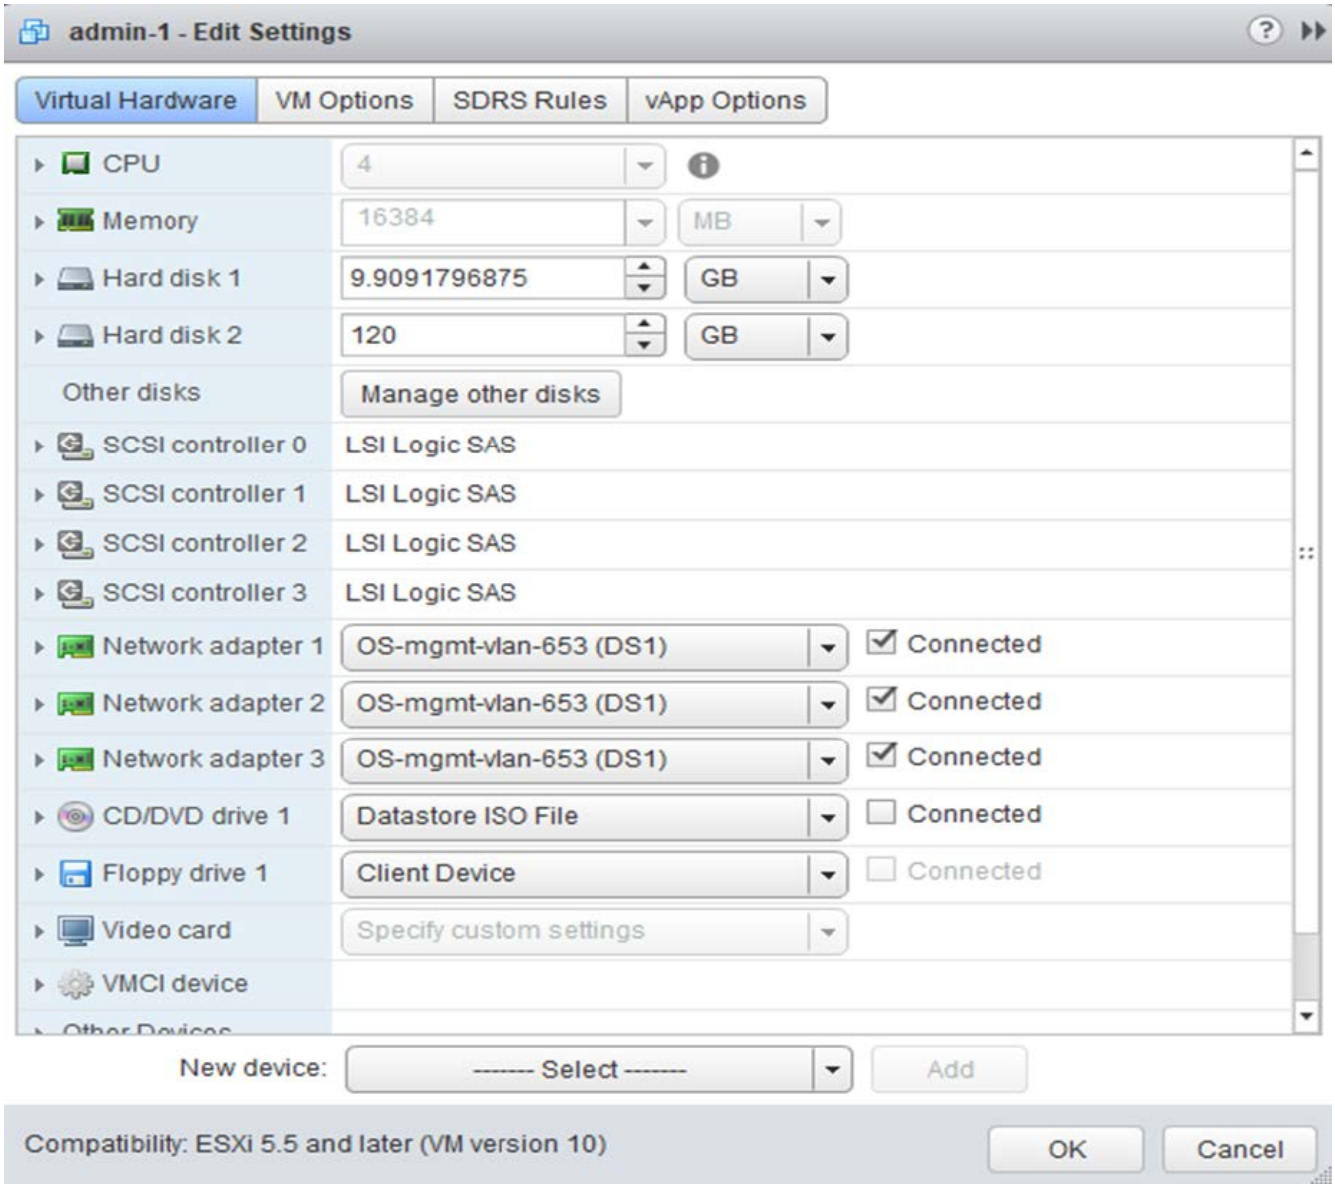
Steps

1. Shut down the ONTAP Select virtual machine.

   If the node is part of an HA pair, perform a storage failover first.

2. Clear the **CD/DVD drive** option.

   This step does not apply if you installed ONTAP Select without using ONTAP Deploy.



3. After the Storage vMotion operation completes, power on the ONTAP Select virtual machine.

   If this node is part of an HA pair, you can perform a manual giveback.

4. Perform a `cluster refresh` operation using the Deploy utility and confirm it is successful.
5. Back up the Deploy utility database.

After you finish

When the Storage vMotion operation completes, you should use the Deploy utility to perform a `cluster refresh` operation. The `cluster refresh` updates the ONTAP Deploy database with the new location of the ONTAP Select node.

Manage ONTAP Select licenses

There are several related tasks you can perform as part of administering the ONTAP Select licenses.

Manage the Capacity Tier licenses

You can add, edit, and delete ONTAP Select Capacity Tier licenses as needed.

Steps

1. Sign in to the Deploy utility through the web interface using the administrator account.
2. Select the **Administration** tab at the top of the page.
3. Select **Licenses** and select **Capacity Tier**.
4. Optionally, select **Filter** to limit the licenses displayed.
5. To replace an existing license; select a license, select ⋮, and select **Update**.
6. To add a new license, select **Add** at the top of the page and then select **Upload License(s)** and select a license file from your local workstation.

Manage the Capacity Pool licenses

You can add, edit, and delete ONTAP Select Capacity Pool licenses as needed.

Steps

1. Sign in to the Deploy utility through the web interface using the administrator account.
2. Select the **Administration** tab at the top of the page.
3. Select **Licenses** and select **Capacity Pools**.
4. Optionally, select **Filter** and to limit the licenses displayed.
5. Optionally, select a license and select ⋮ to manage an existing license.
6. Add a new license or renew an existing license:

   Add new license

   To add a new license, select **Add** at the top of the page.

   Renew existing license

   To renew an existing license:

   a. Select ⋮ on an existing license.
   b. Select **Upload License(s)**.
   c. Select a license file from your local workstation.

7. To see a list of the capacity pools:
   a. Select **Summary**.
   b. Select and expand a pool to see the clusters and nodes leasing storage from the pool.
   c. View the current status of the license under **License Information**.
   d. You can change the duration of the leases issued for the pool under Lease expiration.
8. To see a list of the clusters:
   a. Select **Details**.
   b. Select and expand the cluster to see storage utilization.

Reinstall a Capacity Pool license

Every active Capacity Pool license is locked to a specific License Manager instance, which is contained within an instance of the Deploy administration utility. If you are using a Capacity Pool license and then restore or recover the Deploy instance, the original license is no longer valid. You must generate a new capacity license file, and then install the license to the new Deploy instance.

Before you begin

- Determine all the Capacity Pool licenses used by the original Deploy instance.
- If you restore a backup as part of creating the new Deploy instance, determine if the backup is current and up-to-date.
- Locate the ONTAP Select nodes that were most recently created by the original Deploy instance (only if an up-to-date backup from the original Deploy instance is not restored to the new Deploy instance).
- Restore or recreate the Deploy instance

About this task

At a high level, this task is composed of three parts. You must regenerate and install all the Capacity Pool licenses used by the Deploy instance. After all the licenses have been reinstalled to the new Deploy instance, you can reset the serial sequence number if needed. Finally, if the Deploy IP address has changed, you must update every ONTAP Select node that uses a Capacity Pools license.

Steps

1. Contact NetApp support and have all the Capacity Pool licenses for the original Deploy instance unbound and unregistered.
2. Acquire and download a new license file for each of the Capacity Pool licenses.

   See Acquire a Capacity Pool license for more information.

3. Install the Capacity Pool licenses at the new Deploy instance:
   a. Sign in to the Deploy utility web user interface using the administrator account.
   b. Select the **Administration** tab at the top of the page.
   c. Select **Licenses** and then select **Capacity Pool**.
   d. Select **Add** and then **Upload License(s)** to select and upload the licenses.
4. If you created the new Deploy instance without restoring a backup, or you used a backup that was not current and up-to-date, you must update the serial sequence number:
   a. Sign in to the Deploy utility command line interface using the administrator account.
   b. Display the serial number for a node most recently created by the original Deploy instance:

      ```
      node show -cluster-name CLUSTER_NAME -name NODE_NAME -detailed
      ```

   c. Extract the last eight digits from the twenty-digit node serial number to obtain the last serial sequence number used by the original Deploy instance.
   d. Add 20 to the serial sequence number to create the new serial sequence number.
   e. Set the serial sequence number for the new Deploy instance:

      ```
      license-manager modify -serial-sequence SEQ_NUMBER
      ```

5. If the IP address assigned to the new Deploy instance is different than the IP address of the original Deploy

instance, you must update the IP address at every ONTAP Select node that uses a Capacity Pools license:

a. Sign in to the ONTAP command line interface of the ONTAP Select node.

b. Enter advanced privilege mode:

```
set adv
```

c. Display the current configuration:

```
system license license-manager show
```

d. Set the License Manager (Deploy) IP address used by the node:

```
system license license-manager modify -host NEW_IP_ADDRESS
```

Convert an evaluation license to a production license

You can upgrade an ONTAP Select evaluation cluster to use a production Capacity Tier license with the Deploy administration utility.

Before you begin

- Each node must have enough storage allocated to support the minimum required for a production license.
- You must have Capacity Tier licenses for each node in the evaluation cluster.

About this task

Performing a modification of the cluster license for a single-node cluster is disruptive. However, this is not the case with a multi-node cluster because the conversion process reboots each node one at a time to apply the license.

Steps

1. Sign in to the Deploy utility web user interface using the administrator account.
2. Select the **Clusters** tab a the top of the page and select the desired cluster.
3. At the top of the cluster details page, select **Click here** to modify the cluster license.

   You can also select **Modify** next to evaluation license in the **Cluster Details** section.

4. Select an available production license for each node or upload additional licenses as needed.
5. Provide the ONTAP credentials and select **Modify**.

   The license upgrade for the cluster can take several minutes. Allow the process to complete before leaving the page or making any other changes.

After you finish

The twenty-digit node serial numbers originally assigned to each node for the evaluation deployment are replaced by the nine-digit serial numbers from the production licenses used for the upgrade.

Manage an expired Capacity Pool license

Generally, when a license expires, nothing happens. However, you cannot install a different license because the nodes are associated with the expired license. Until you renew the license, you should *not* do anything that would bring the aggregate offline, such as a reboot or failover operation. The recommended action is to

expedite the license renewal.

For more information about ONTAP Select and license renewal, refer to the Licenses, installation, upgrades, and reverts section in the FAQ.

Manage add-on licenses

For the ONTAP Select product, add-on licenses are applied directly within ONTAP and are not managed through ONTAP Select Deploy. See Manage licenses overview (cluster administrators only) and Enable new features by adding license keys for more information.

Deep dive

Storage

Learn about the general storage characteristics for ONTAP Select

Discover general storage concepts that apply to the ONTAP Select environment before exploring the specific storage components.

Phases of storage configuration

The major configuration phases of the ONTAP Select host storage include the following:

- Pre-deployment prerequisites
  - Make sure that each hypervisor host is configured and ready for an ONTAP Select deployment.
  - The configuration involves the physical drives, RAID controllers and groups, LUNs, as well as related network preparation.
  - This configuration is performed outside of ONTAP Select.
- Configuration using the hypervisor administrator utility
  - You can configure certain aspects of the storage using the hypervisor administration utility (for example, vSphere in a VMware environment).
  - This configuration is performed outside of ONTAP Select.
- Configuration using the ONTAP Select Deploy administration utility
  - You can use the Deploy administration utility to configure the core logical storage constructs.
  - This is performed either explicitly through CLI commands or automatically by the utility as part of a deployment.
- Post-deployment configuration
  - After an ONTAP Select deployment completes, you can configure the cluster using the ONTAP CLI or System Manager.
  - This configuration is performed outside of ONTAP Select Deploy.

Managed versus unmanaged storage

Storage that is accessed and directly controlled by ONTAP Select is managed storage. Any other storage on the same hypervisor host is unmanaged storage.

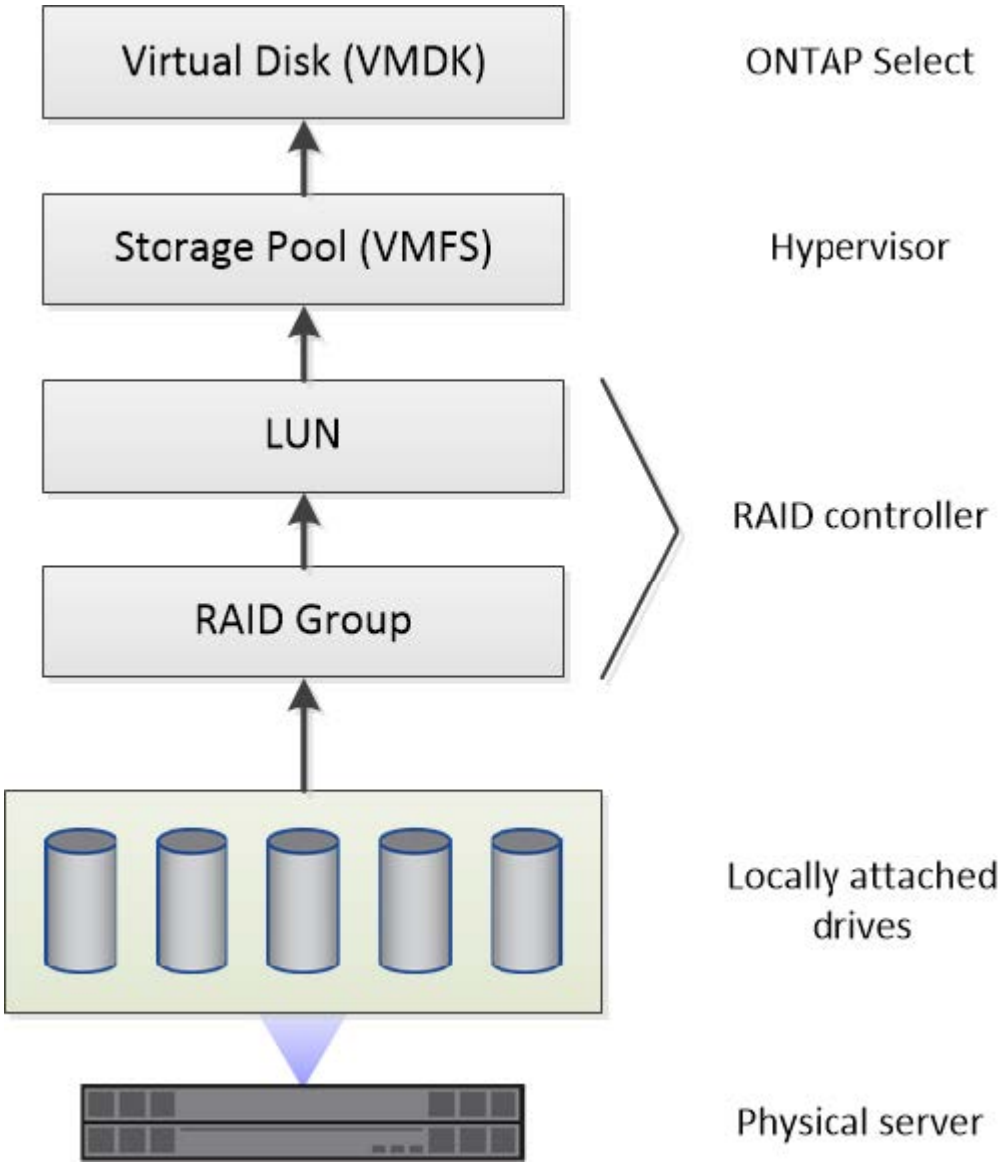
Homogeneous physical storage

All the physical drives comprising the ONTAP Select managed storage must be homogeneous. That is, all the hardware must be the same regarding the following characteristics:

- Type (SAS, NL-SAS, SATA, SSD)
- Speed (RPM)

Illustration of the local storage environment

Each hypervisor host contains local disks and other logical storage components that can be used by ONTAP Select. These storage components are arranged in a layered structure, from the physical disk.



Characteristics of the local storage components

There are several concepts that apply to the local storage components used in an ONTAP Select environment. You should be familiar with these concepts before preparing for an ONTAP Select deployment. These concepts are arranged according to category: RAID groups and LUNs, storage pools, and virtual disks.

Grouping physical drives into RAID groups and LUNs

One or more physical disks can be locally attached to the host server and available to ONTAP Select. The physical disks are assigned to RAID groups, which are then presented to the hypervisor host operating system as one or more LUNs. Each LUN is presented to the hypervisor host operating system as a physical hard drive.

When configuring an ONTAP Select host, you should be aware of the following:

- All managed storage must be accessible through a single RAID controller
- Depending on the vendor, each RAID controller supports a maximum number of drives per RAID group

One or more RAID groups

Each ONTAP Select host must have a single RAID controller. You should create a single RAID group for ONTAP Select. However, in certain situations you might consider creating more than one RAID group. Refer to Summary of best practices.

Storage pool considerations

There are several issues related to the storage pools that you should be aware of as part of preparing to deploy ONTAP Select.

In a VMware environment, a storage pool is synonymous with a VMware datastore.

Storage pools and LUNs

Each LUN is seen as a local disk on the hypervisor host and can be part of one storage pool. Each storage pool is formatted with a file system that the hypervisor host OS can use.

You must make sure that the storage pools are created properly as part of an ONTAP Select deployment. You can create a storage pool using the hypervisor administration tool. For example, with VMware you can use the vSphere client to create a storage pool. The storage pool is then passed in to the ONTAP Select Deploy administration utility.

Manage the virtual disks on ESXi

There are several issues related to the virtual disks that you should be aware of as part of preparing to deploy ONTAP Select.

Virtual disks and file systems

The ONTAP Select virtual machine is allocated multiple virtual disk drives. Each virtual disk is actually a file contained in a storage pool and is maintained by the hypervisor. There are several types of disks used by ONTAP Select, primarily system disks and data disks.

You should also be aware of the following regarding virtual disks:

- The storage pool must be available before the virtual disks can be created.
- The virtual disks cannot be created before the virtual machine is created.
- You must rely on the ONTAP Select Deploy administration utility to create all virtual disks (that is, an administrator must never create a virtual disk outside of the Deploy utility).

Configuring the virtual disks

The virtual disks are managed by ONTAP Select. They are created automatically when you create a cluster using the Deploy administration utility.

Illustration of the external storage environment on ESXi

The ONTAP Select vNAS solution enables ONTAP Select to use datastores residing on storage that is external to the hypervisor host. The datastores can be accessed through the network using VMware vSAN or directly at

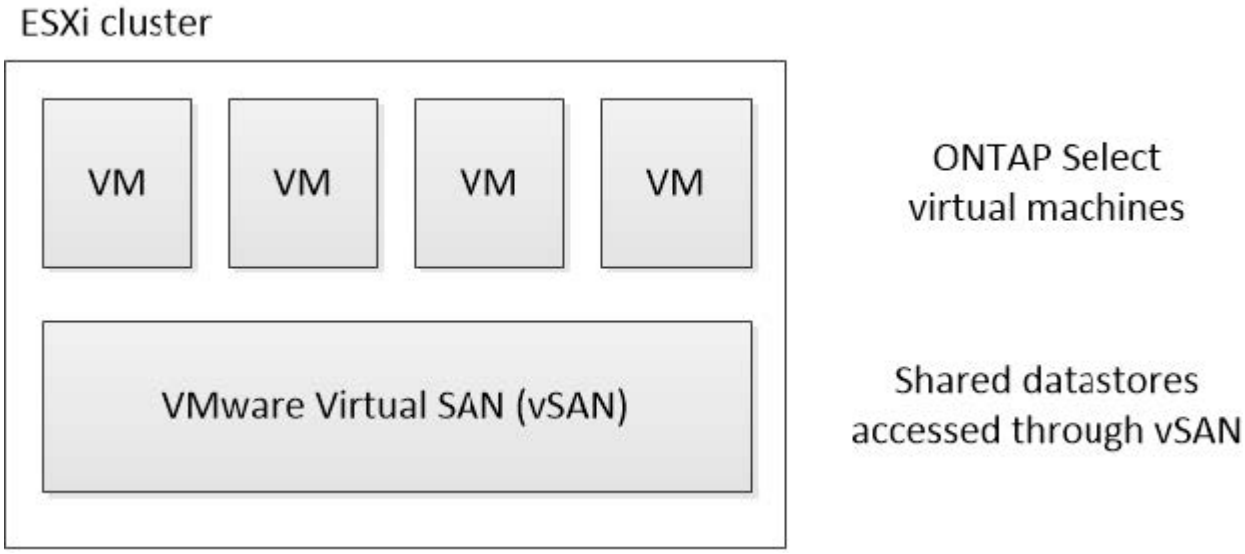
an external storage array.

ONTAP Select can be configured to use the following types of VMware ESXi network datastores which are external to the hypervisor host:

- vSAN (Virtual SAN)
- VMFS
- NFS

vSAN datastores

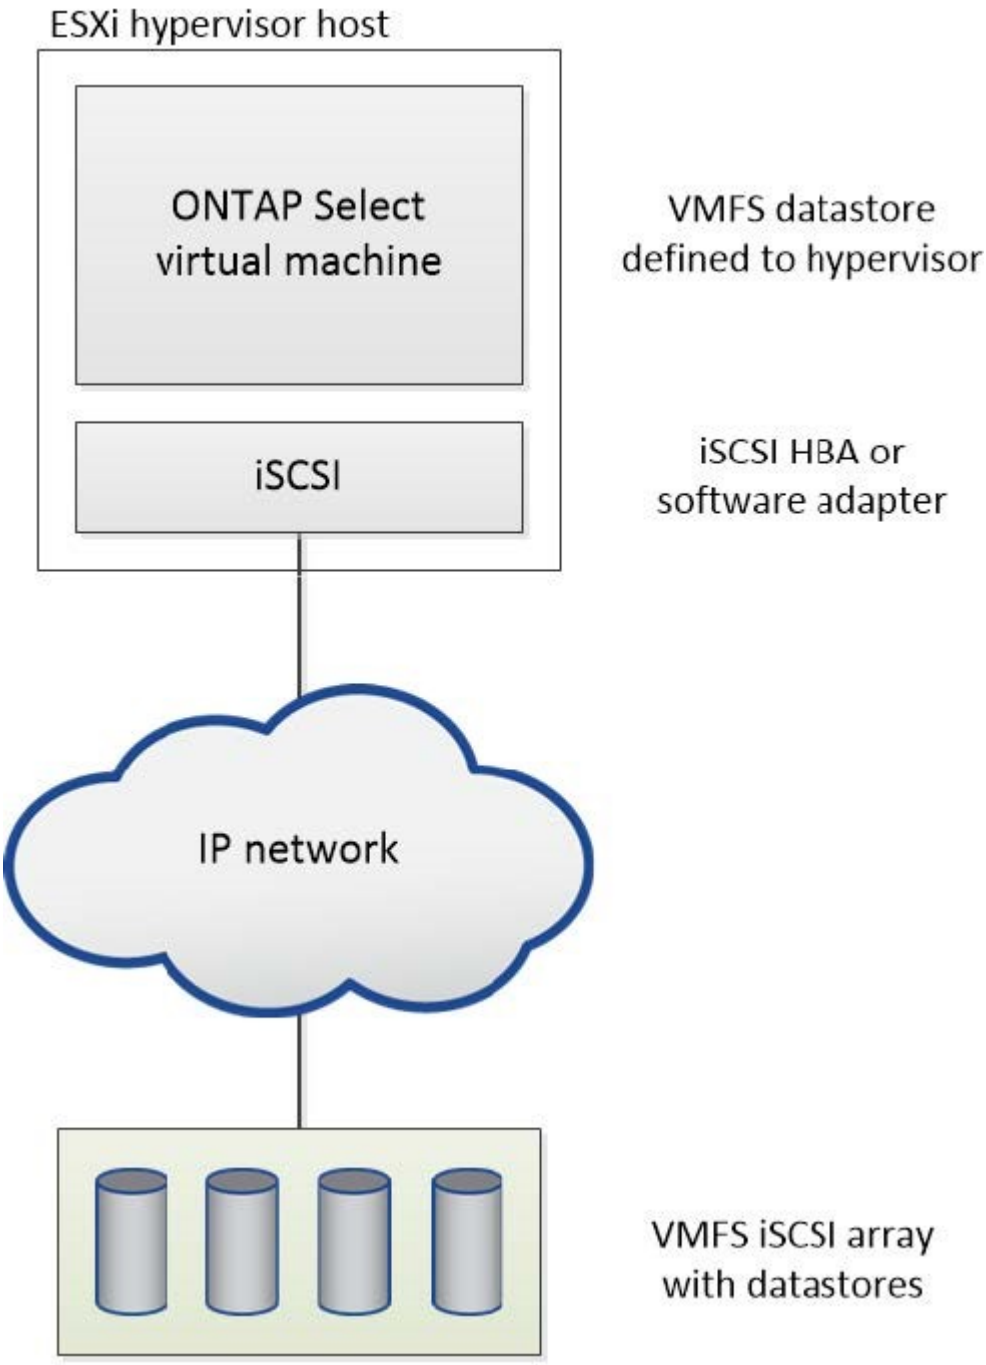
Every ESXi host can have one or more local VMFS datastores. Normally these datastores are only accessible to the local host. However, VMware vSAN allows each of the hosts in an ESXi cluster to share all of the datastores in the cluster as if they were local. The following figure illustrates how vSAN creates a pool of datastores that are shared among the hosts in the ESXi cluster.



VMFS datastore on external storage array

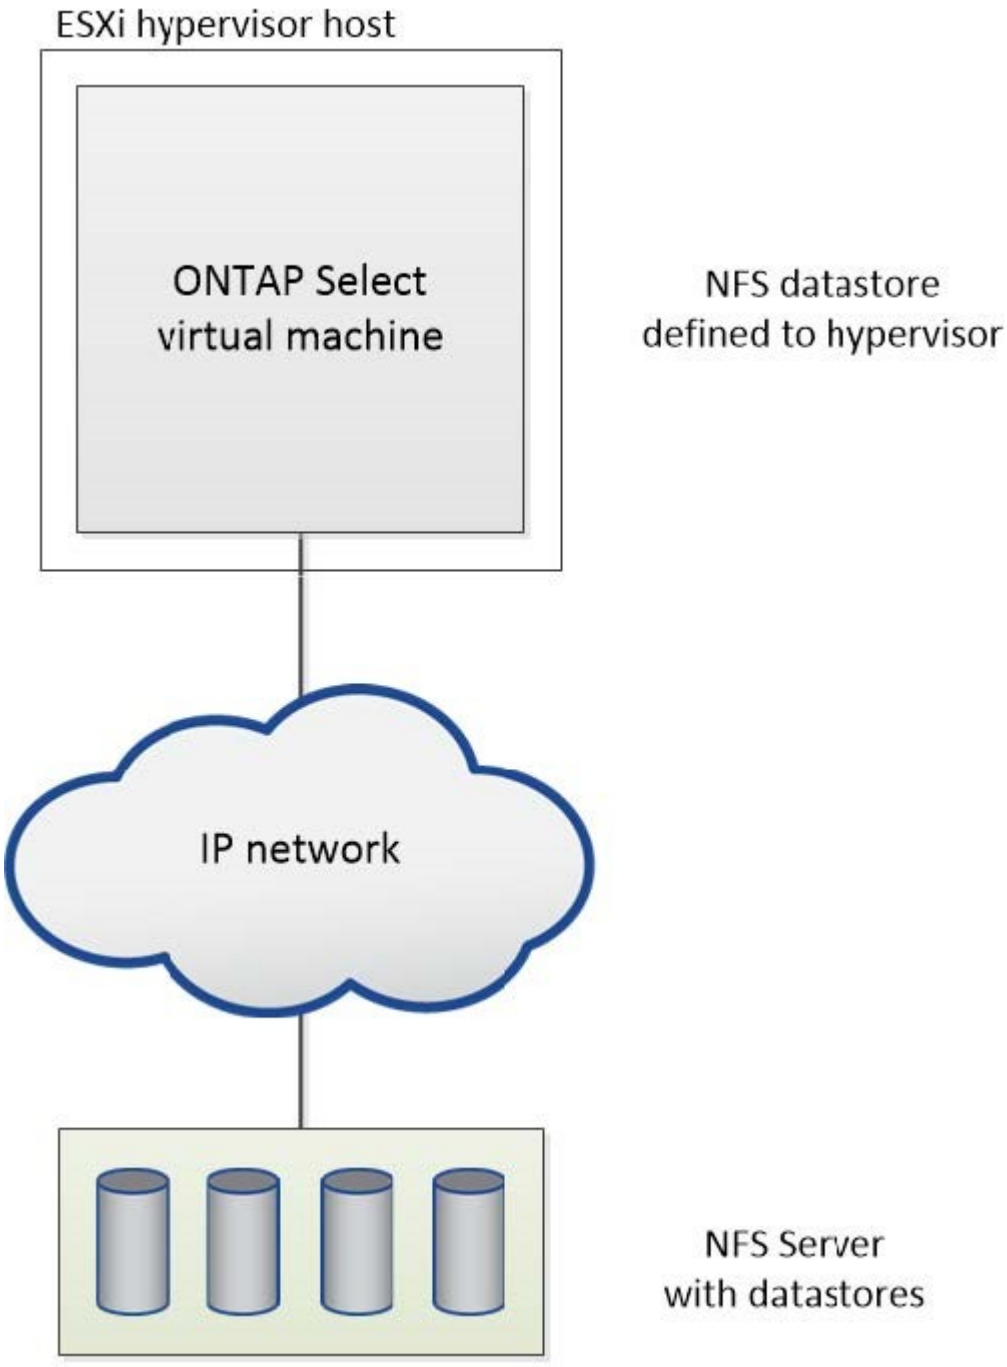
You can create a VMFS datastore residing on an external storage array. The storage is accessed using one of several different network protocols. The following figure illustrates a VMFS datastore on an external storage array accessed using the iSCSI protocol.

ONTAP Select supports all external storage arrays described in the VMware Storage/SAN Compatibility documentation, including iSCSI, Fiber Channel, and Fiber Channel over Ethernet.



NFS datastore on external storage array

You can create an NFS datastore residing on an external storage array. The storage is accessed using the NFS network protocol. The following figure illustrates an NFS datastore on external storage that is accessed through the NFS server appliance.



Hardware RAID services for ONTAP Select with local attached storage

When a hardware RAID controller is available, ONTAP Select can move RAID services to the hardware controller for both a write performance boost and protection against physical drive failures. As a result, RAID protection for all nodes within the ONTAP Select cluster is provided by the locally attached RAID controller and not through ONTAP software RAID.

ONTAP Select data aggregates are configured to use RAID 0 because the physical RAID controller is providing RAID striping to the underlying drives. No other RAID levels are supported.

RAID controller configuration for local attached storage

All locally attached disks that provide ONTAP Select with backing storage must sit behind a RAID controller. Most commodity servers come with multiple RAID controller options across multiple price points, each with varying levels of functionality. The intent is to support as many of these options as possible, providing they meet certain minimum requirements placed on the controller.

> You cannot detach virtual disks from ONTAP Select VMs that are using the hardware RAID configuration. Detaching disks is only supported for ONTAP Select VMs that are using the software RAID configuration. See Replace a failed drive in an ONTAP Select software RAID configuration for more information.

The RAID controller that manages the ONTAP Select disks must meet the following requirements:

- The hardware RAID controller must have a battery backup unit (BBU) or flash-backed write cache (FBWC) and support 12Gbps of throughput.
- The RAID controller must support a mode that can withstand at least one or two disk failures (RAID 5 and RAID 6).
- The drive cache must be set to disabled.
- The write policy must be configured for writeback mode with a fallback to write through upon BBU or flash failure.
- The I/O policy for reads must be set to cached.

All locally attached disks that provide ONTAP Select with backing storage must be placed into RAID groups running RAID 5 or RAID 6. For SAS drives and SSDs, using RAID groups of up to 24 drives allows ONTAP to reap the benefits of spreading incoming read requests across a higher number of disks. Doing so provides a significant gain in performance. With SAS/SSD configurations, performance testing was performed against single-LUN versus multi-LUN configurations. No significant differences were found, so, for simplicity's sake, NetApp recommends creating the fewest number of LUNs necessary to support your configuration needs.

NL-SAS and SATA drives require a different set of best practices. For performance reasons, the minimum number of disks is still eight, but the RAID group size should not be larger than 12 drives. NetApp also recommends using one spare per RAID group; however, global spares for all RAID groups can be used. For example, you can use two spares for every three RAID groups, with each RAID group consisting of eight to 12 drives.

> The maximum extent and datastore size for older ESXi releases is 64TB, which can affect the number of LUNs necessary to support the total raw capacity provided by these large capacity drives.

RAID mode

Many RAID controllers support up to three modes of operation, each representing a significant difference in the data path taken by write requests. These three modes are as follows:

- Writethrough. All incoming I/O requests are written to the RAID controller cache and then immediately flushed to disk before acknowledging the request back to the host.
- Writearound. All incoming I/O requests are written directly to disk, circumventing the RAID controller cache.
- Writeback. All incoming I/O requests are written directly to the controller cache and immediately acknowledged back to the host. Data blocks are flushed to disk asynchronously using the controller.

Writeback mode offers the shortest data path, with I/O acknowledgment occurring immediately after the blocks enter cache. This mode provides the lowest latency and highest throughput for mixed read/write workloads. However, without the presence of a BBU or nonvolatile flash technology, users run the risk of losing data if the system incurs a power failure when operating in this mode.

ONTAP Select requires the presence of a battery backup or flash unit; therefore, we can be confident that cached blocks are flushed to disk in the event of this type of failure. For this reason, it is a requirement that the RAID controller be configured in writeback mode.

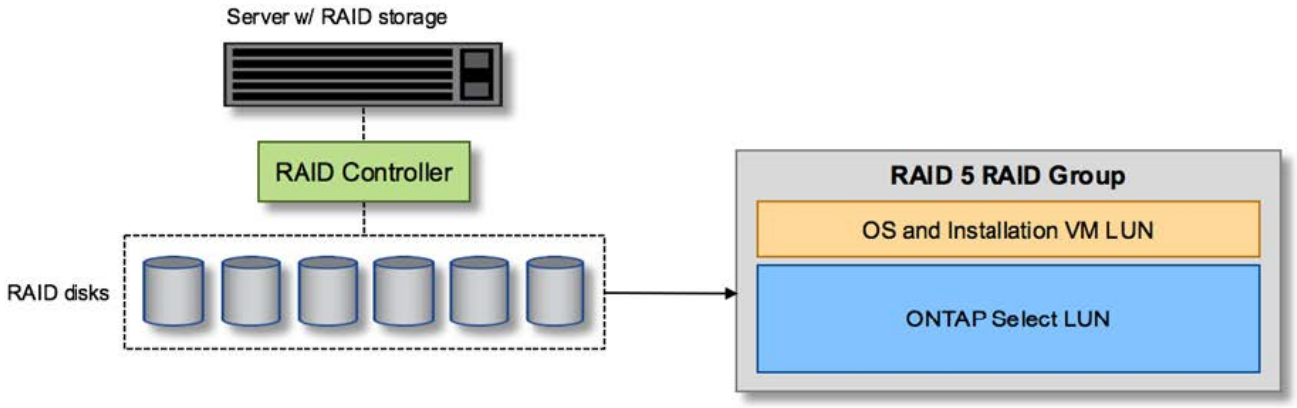
Local disks shared between ONTAP Select and OS

The most common server configuration is one in which all locally attached spindles sit behind a single RAID controller. You should provision a minimum of two LUNs: one for the hypervisor and one for the ONTAP Select VM.

For example, consider an HP DL380 g8 with six internal drives and a single Smart Array P420i RAID controller. All internal drives are managed by this RAID controller, and no other storage is present on the system.

The following figure shows this style of configuration. In this example, no other storage is present on the system; therefore, the hypervisor must share storage with the ONTAP Select node.

Server LUN configuration with only RAID-managed spindles



Provisioning the OS LUNs from the same RAID group as ONTAP Select allows the hypervisor OS (and any client VM that is also provisioned from that storage) to benefit from RAID protection. This configuration prevents a single-drive failure from bringing down the entire system.

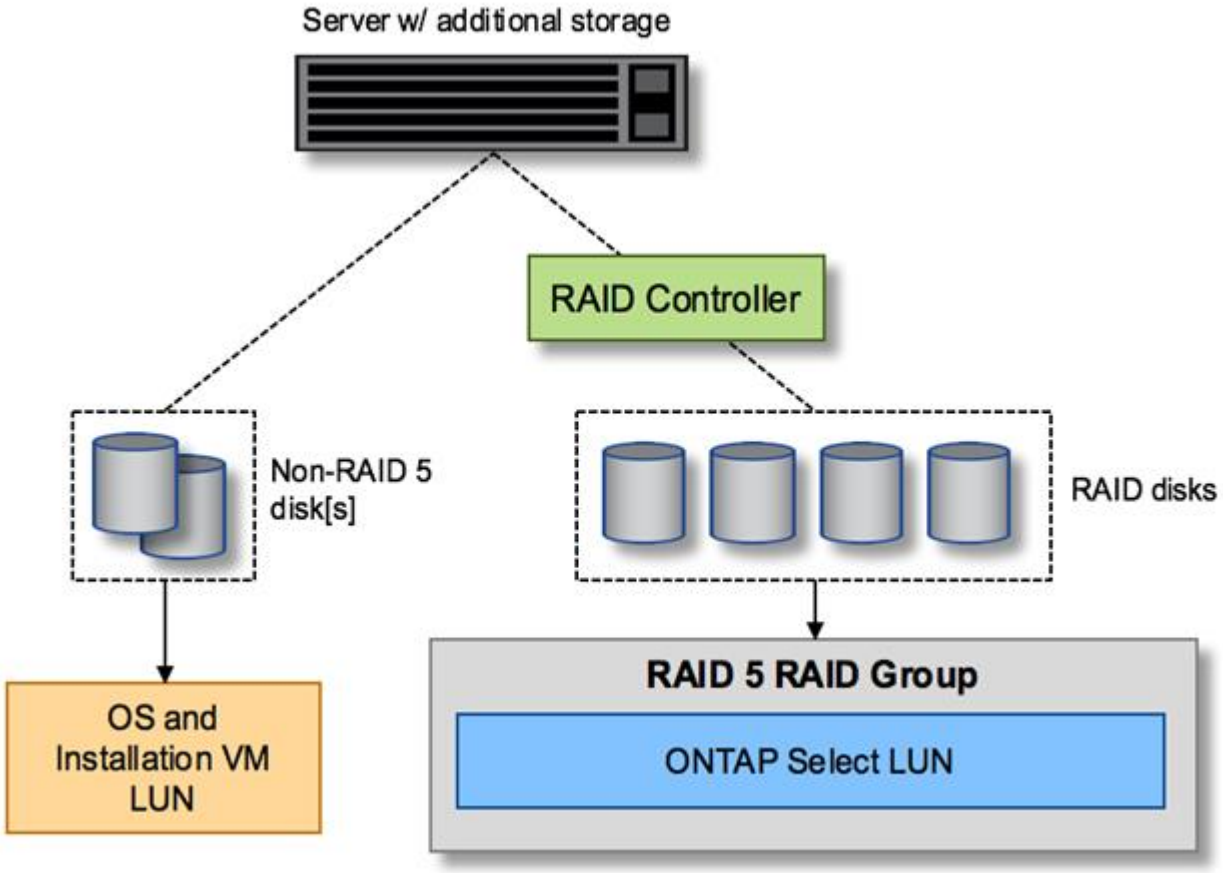
Local disks split between ONTAP Select and OS

The other possible configuration provided by server vendors involves configuring the system with multiple RAID or disk controllers. In this configuration, a set of disks is managed by one disk controller, which might or might not offer RAID services. A second set of disks is managed by a hardware RAID controller that is able to offer RAID 5/6 services.

With this style of configuration, the set of spindles that sits behind the RAID controller that can provide RAID 5/6 services should be used exclusively by the ONTAP Select VM. Depending on the total storage capacity under management, you should configure the disk spindles into one or more RAID groups and one or more LUNs. These LUNs would then be used to create one or more datastores, with all datastores being protected by the RAID controller.

The first set of disks is reserved for the hypervisor OS and any client VM that is not using ONTAP storage, as shown in the following figure.

Server LUN configuration on mixed RAID/non-RAID system



Multiple LUNs

There are two cases for which single-RAID group/single-LUN configurations must change. When using NL-SAS or SATA drives, the RAID group size must not exceed 12 drives. In addition, a single LUN can become larger than the underlying hypervisor storage limits either individual file system extent maximum size or total storage pool maximum size. Then the underlying physical storage must be broken up into multiple LUNs to enable successful file system creation.

VMware vSphere virtual machine file system limits

The maximum size of a datastore on some versions of ESXi is 64TB.

If a server has more than 64TB of storage attached, multiple LUNs might need to be provisioned, each smaller than 64TB. Creating multiple RAID groups to improve the RAID rebuild time for SATA/NL-SAS drives also results in multiple LUNs being provisioned.

When multiple LUNs are required, a major point of consideration is making sure that these LUNs have similar and consistent performance. This is especially important if all the LUNs are to be used in a single ONTAP aggregate. Alternatively, if a subset of one or more LUNs has a distinctly different performance profile, we strongly recommend isolating these LUNs in a separate ONTAP aggregate.

Multiple file system extents can be used to create a single datastore up to the maximum size of the datastore. To restrict the amount of capacity that requires an ONTAP Select license, make sure to specify a capacity cap during the cluster installation. This functionality allows ONTAP Select to use (and therefore require a license for) only a subset of the space in a datastore.

Alternatively, one can start by creating a single datastore on a single LUN. When additional space requiring a larger ONTAP Select capacity license is needed, then that space can be added to the same datastore as an extent, up to the maximum size of the datastore. After the maximum size is reached, new datastores can be created and added to ONTAP Select. Both types of capacity extension operations are supported and can be achieved by using the ONTAP Deploy storage-add functionality. Each ONTAP Select node can be configured to support up to 400TB of storage. Provisioning capacity from multiple datastores requires a two-step process.

The initial cluster create can be used to create an ONTAP Select cluster consuming part or all of the space in the initial datastore. A second step is to perform one or more capacity addition operations using additional datastores until the desired total capacity is reached. This functionality is detailed in the section Increase storage capacity.

VMFS overhead is non-zero (see VMware KB 1001618), and attempting to use the entire space reported as free by a datastore has resulted in spurious errors during cluster create operations.

A 2% buffer is left unused in each datastore. This space does not require a capacity license because it is not used by ONTAP Select. ONTAP Deploy automatically calculates the exact number of gigabytes for the buffer, as long as a capacity cap is not specified. If a capacity cap is specified, that size is enforced first. If the capacity cap size falls within the buffer size, the cluster create fails with an error message specifying the correct maximum size parameter that can be used as a capacity cap:

```
“InvalidPoolCapacitySize: Invalid capacity specified for storage pool
“ontap-select-storage-pool”, Specified value: 34334204 GB. Available
(after leaving 2% overhead space): 30948”
```

VMFS 6 is supported for both new installations and as the target of a Storage vMotion operation of an existing ONTAP Deploy or ONTAP Select VM.

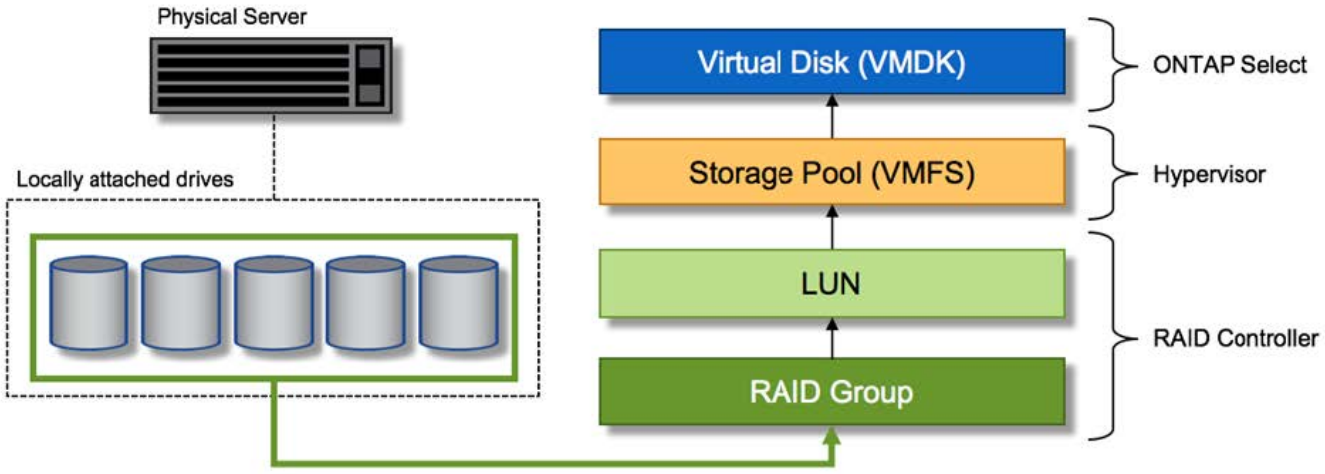
VMware does not support in-place upgrades from VMFS 5 to VMFS 6. Therefore, Storage vMotion is the only mechanism that allows any VM to transition from a VMFS 5 datastore to a VMFS 6 datastore. However, support for Storage vMotion with ONTAP Select and ONTAP Deploy was expanded to cover other scenarios besides the specific purpose of transitioning from VMFS 5 to VMFS 6.

ONTAP Select virtual disks

At its core, ONTAP Select presents ONTAP with a set of virtual disks provisioned from one or more storage pools. ONTAP is presented with a set of virtual disks that it treats as physical, and the remaining portion of the storage stack is abstracted by the hypervisor. The following figure shows this relationship in more detail, highlighting the relationship between the physical RAID controller, the hypervisor, and the ONTAP Select VM.

- RAID group and LUN configuration occur from within the server’s RAID controller software. This configuration is not required when using VSAN or external arrays.
- Storage pool configuration occurs from within the hypervisor.
- Virtual disks are created and owned by individual VMs; in this example, by ONTAP Select.

Virtual disk to physical disk mapping



Virtual disk provisioning

To provide for a more streamlined user experience, the ONTAP Select management tool, ONTAP Deploy, automatically provisions virtual disks from the associated storage pool and attaches them to the ONTAP Select VM. This operation occurs automatically during both initial setup and during storage-add operations. If the ONTAP Select node is part of an HA pair, the virtual disks are automatically assigned to a local and mirror storage pool.

ONTAP Select breaks up the underlying attached storage into equal-sized virtual disks, each not exceeding 16TB. If the ONTAP Select node is part of an HA pair, a minimum of two virtual disks are created on each cluster node and assigned to the local and mirror plex to be used within a mirrored aggregate.

For example, an ONTAP Select can assigned a datastore or LUN that is 31TB (the space remaining after the VM is deployed and the system and root disks are provisioned). Then four ~7.75TB virtual disks are created and assigned to the appropriate ONTAP local and mirror plex.



Adding capacity to an ONTAP Select VM likely results in VMDKs of different sizes. For details, refer to the section Increase storage capacity. Unlike FAS systems, different sized VMDKs can exist in the same aggregate. ONTAP Select uses a RAID 0 stripe across these VMDKs, which results in the ability to fully use all the space in each VMDK regardless of its size.

Virtualized NVRAM

NetApp FAS systems are traditionally fitted with a physical NVRAM PCI card, a high-performing card containing nonvolatile flash memory. This card provides a significant boost in write performance by granting ONTAP with the ability to immediately acknowledge incoming writes back to the client. It can also schedule the movement of modified data blocks back to the slower storage media in a process known as destaging.

Commodity systems are not typically fitted with this type of equipment. Therefore, the functionality of this NVRAM card has been virtualized and placed into a partition on the ONTAP Select system boot disk. It is for this reason that placement of the system virtual disk of the instance is extremely important. This is also why the product requires the presence of a physical RAID controller with a resilient cache for local attached storage configurations.

NVRAM is placed on its own VMDK. Splitting the NVRAM in its own VMDK allows the ONTAP Select VM to use the vNVMe driver to communicate with its NVRAM VMDK. It also requires that the ONTAP Select VM uses hardware version 13, which is compatible with ESXi 8.0 and later.

Data path explained: NVRAM and RAID controller

The interaction between the virtualized NVRAM system partition and the RAID controller can be best highlighted by walking through the data path taken by a write request as it enters the system.

Incoming write requests to the ONTAP Select VM are targeted at the VM's NVRAM partition. At the virtualization layer, this partition exists within an ONTAP Select system disk, a VMDK attached to the ONTAP Select VM. At the physical layer, these requests are cached in the local RAID controller, like all block changes targeted at the underlying spindles. From here, the write is acknowledged back to the host.

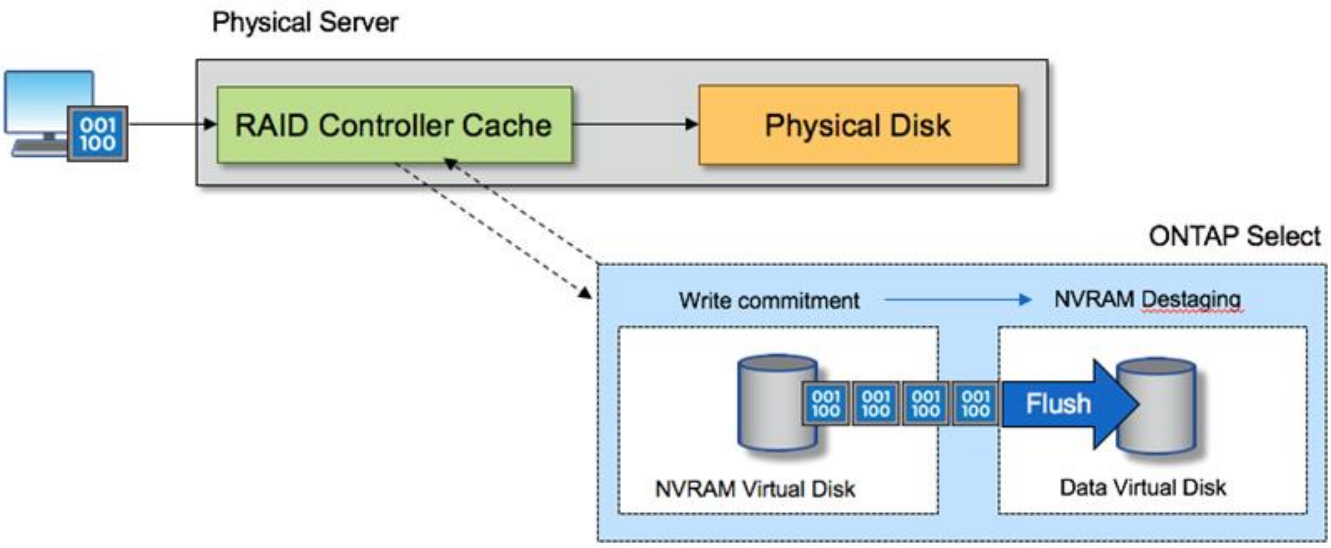
At this point, physically, the block resides in the RAID controller cache, waiting to be flushed to disk. Logically, the block resides in NVRAM waiting for destaging to the appropriate user data disks.

Because changed blocks are automatically stored within the RAID controller's local cache, incoming writes to the NVRAM partition are automatically cached and periodically flushed to physical storage media. This should not be confused with the periodic flushing of NVRAM contents back to ONTAP data disks. These two events are unrelated and occur at different times and frequencies.

The following figure shows the I/O path an incoming write takes. It highlights the difference between the physical layer (represented by the RAID controller cache and disks) and the virtual layer (represented by the VM's NVRAM and data virtual disks).

Although blocks changed on the NVRAM VMDK are cached in the local RAID controller cache, the cache is not aware of the VM construct or its virtual disks. It stores all changed blocks on the system, of which NVRAM is only a part. This includes write requests bound for the hypervisor, if it is provisioned from the same backing spindles.

Incoming writes to ONTAP Select VM



The NVRAM partition is separated on its own VMDK. That VMDK is attached using the vNVME driver available in ESXi versions of 8.0 or later. This change is most significant for ONTAP Select installations with software RAID, which do not benefit from the RAID controller cache.

ONTAP Select software RAID configuration services for local attached storage

Software RAID is a RAID abstraction layer implemented within the ONTAP software stack. It provides the same functionality as the RAID layer within a traditional ONTAP platform such as FAS. The RAID layer performs drive parity calculations and provides protection against individual drive failures within an ONTAP Select node.

Independent of the hardware RAID configurations, ONTAP Select also provides a software RAID option. A hardware RAID controller might not be available or might be undesirable in certain environments, such as when ONTAP Select is deployed on a small form-factor commodity hardware. Software RAID expands the available deployment options to include such environments. To enable software RAID in your environment, here are some points to remember:

- It is available with a Premium or Premium XL license.
- It only supports SSD or NVMe (requires Premium XL license) drives for ONTAP root and data disks.
- It requires a separate system disk for the ONTAP Select VM boot partition.
  - Choose a separate disk, either an SSD or an NVMe drive, to create a datastore for the system disks (NVRAM, Boot/CF card, Coredump, and Mediator in a multi-node setup).



- The terms service disk and system disk are used interchangeably.
  - Service disks are the virtual disks (VMDKs) that are used within the ONTAP Select VM to service various items such as clustering, booting, and so on.
  - Service disks are physically located on a single physical disk (collectively called the service/system physical disk) as seen from the host. That physical disk must contain a DAS datastore. ONTAP Deploy creates these service disks for the ONTAP Select VM during cluster deployment.
- It is not possible to further separate the ONTAP Select system disks across multiple datastores or across multiple physical drives.
- Hardware RAID is not deprecated.

Software RAID configuration for local attached storage

When using software RAID, the absence of a hardware RAID controller is ideal, but, if a system does have an existing RAID controller, it must adhere to the following requirements:

- You must disable the hardware RAID controller so that disks can be presented directly to the system (a JBOD). You can usually make this change in the RAID controller BIOS.
- Or the hardware RAID controller should be in the SAS HBA mode. For example, some BIOS configurations allow an “AHCI” mode in addition to RAID, which you could choose to enable the JBOD mode. This enables a passthrough, so that the physical drives can be seen as is on the host.

Depending on the maximum number of drives supported by the controller, an additional controller might be required. With the SAS HBA mode, ensure that the I/O controller (SAS HBA) is supported with a minimum of 6Gbps speed. However, NetApp recommends a 12Gbps speed.

No other hardware RAID controller modes or configurations is supported. For example, some controllers allow a RAID 0 support that can artificially enable disks to pass-through but the implications can be undesirable. The supported size of physical disks (SSD only) is between 200GB and 16TB.

Administrators need to keep track of which drives are in use by the ONTAP Select VM and prevent inadvertent use of those drives on the host.

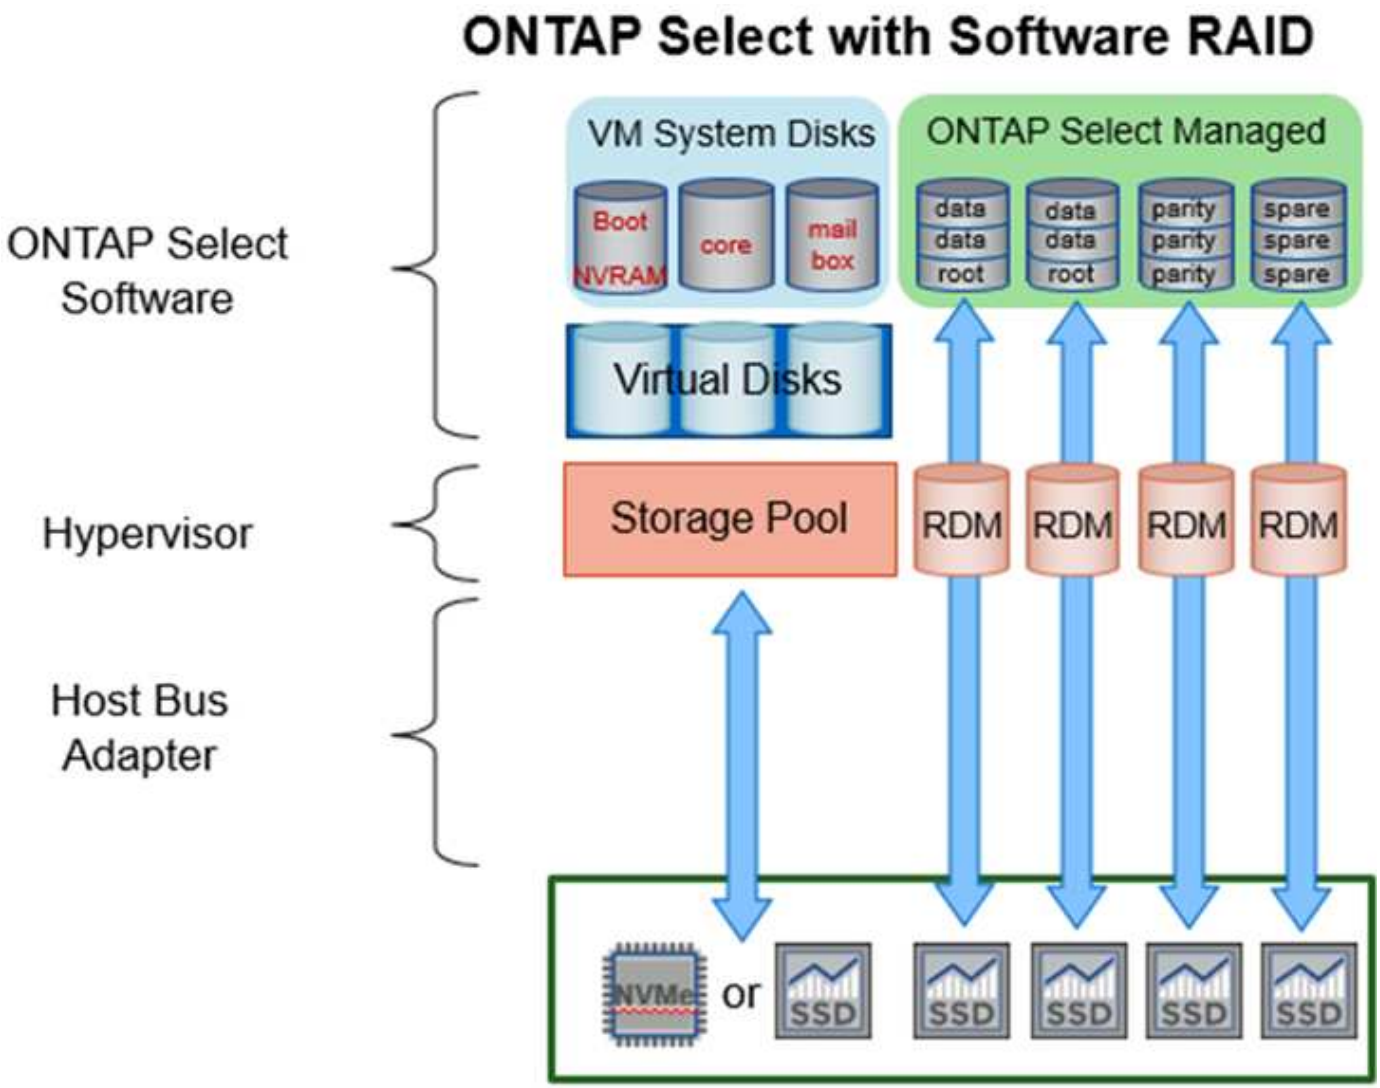
ONTAP Select virtual and physical disks

For configurations with hardware RAID controllers, physical disk redundancy is provided by the RAID controller. ONTAP Select is presented with one or more VMDKs from which the ONTAP admin can configure data aggregates. These VMDKs are striped in a RAID 0 format because using ONTAP software RAID is redundant, inefficient, and ineffective due to resiliency provided at the hardware level. Furthermore, the VMDKs used for system disks are in the same datastore as the VMDKs used to store user data.

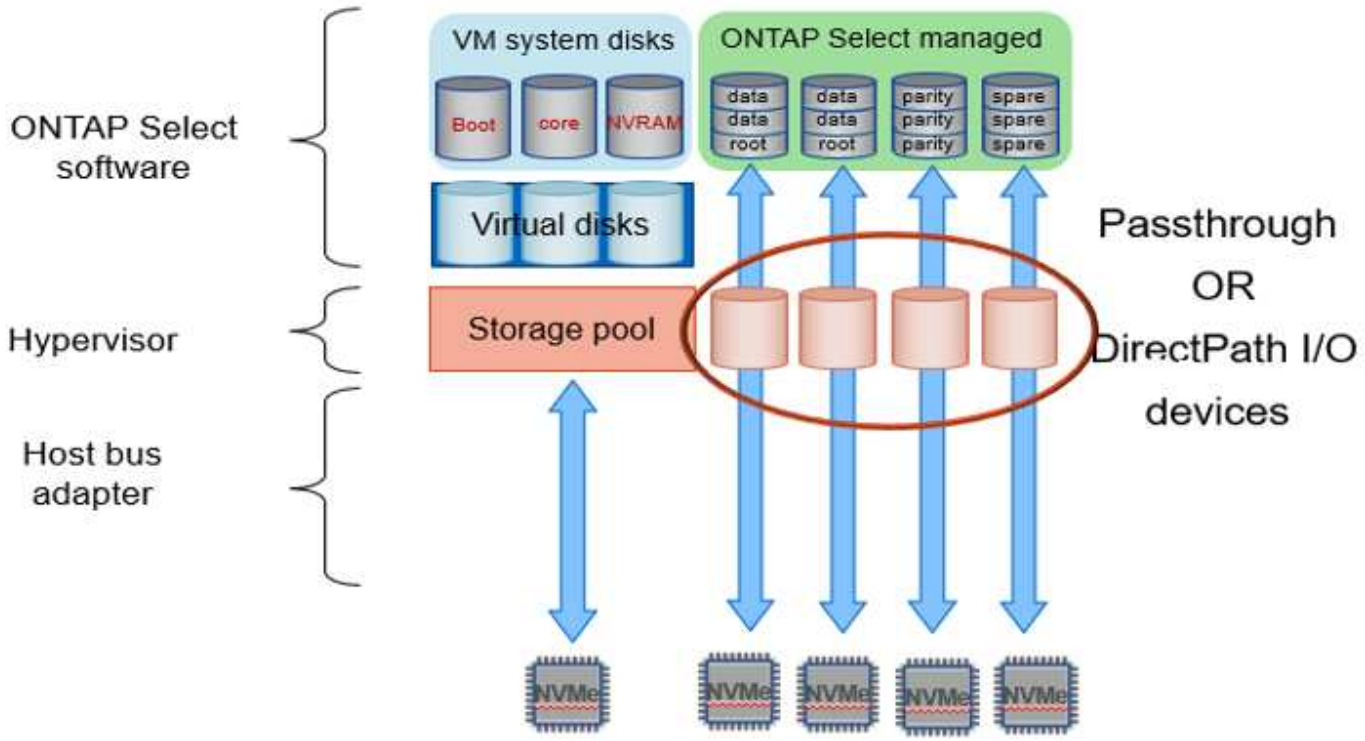
When using software RAID, ONTAP Deploy presents ONTAP Select with a set of VMDKs and physical disks Raw Device Mappings [RDMs] for SSDs and passthrough or DirectPath IO devices for NVMe.

The following figures show this relationship in more detail, highlighting the difference between the virtualized disks used for the ONTAP Select VM internals and the physical disks used to store user data.

ONTAP Select software RAID: use of virtualized disks and RDMs



The system disks (VMDKs) reside in the same datastore and on the same physical disk. The virtual NVRAM disk requires a fast and durable media. Therefore, only NVMe and SSD-type datastores are supported.



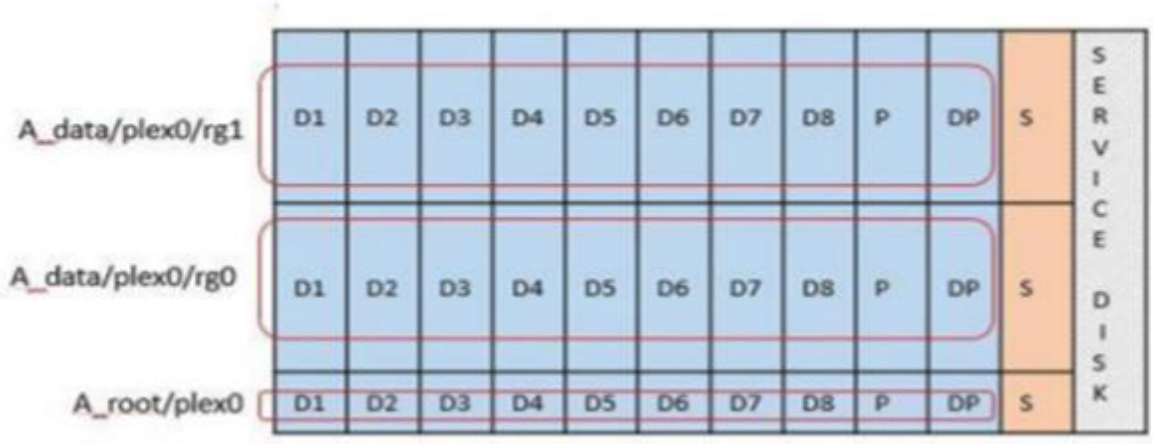
The system disks (VMDKs) reside in the same datastore and on the same physical disk. The virtual NVRAM disk requires a fast and durable media. Therefore, only NVMe and SSD-type datastores are supported. When using NVMe drives for data, the system disk should also be an NVMe device for performance reasons. A good candidate for the system disk in an all NVMe configuration is an INTEL Optane card.

With the current release, it is not possible to further separate the ONTAP Select system disks across multiple datastores or multiple physical drives.

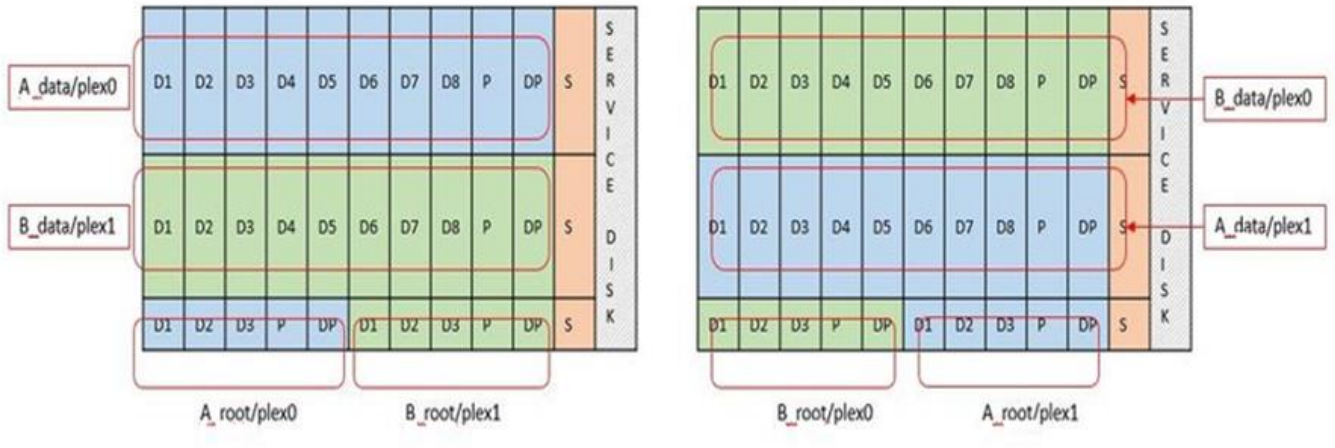
Each data disk is divided into three parts: a small root partition (stripe) and two equal-sized partitions to create two data disks seen within the ONTAP Select VM. Partitions use the Root Data Data (RD2) schema as shown in the following figures for a single node cluster and for a node in an high-availability (HA) pair.

`P` denotes a parity drive, `DP` denotes a dual parity drive, and `S` denotes a spare drive.

RDD disk partitioning for single-node clusters



RDD disk partitioning for multi-node clusters (HA pairs)



ONTAP software RAID supports the following RAID types: RAID 4, RAID-DP, and RAID-TEC. These are the same RAID constructs used by FAS and AFF platforms. For root provisioning, ONTAP Select supports only RAID 4 and RAID-DP. When using RAID-TEC for the data aggregate, the overall protection is RAID-DP. ONTAP Select HA uses a shared-nothing architecture that replicates each node's configuration to the other node. That means each node must store its root partition and a copy of its peer's root partition. A data disk has a single root partition. This means that the minimum number of data disks varies depending on whether the ONTAP Select node is part of an HA pair.

For single node clusters, all data partitions are used to store local (active) data. For nodes that are part of an HA pair, one data partition is used to store local (active) data for that node and the second data partition is used to mirror active data from the HA peer.

Passthrough (DirectPath IO) devices vs. Raw Device Maps (RDMs)

ESXi and KVM hypervisors don't support NVMe disks as Raw Device Maps (RDMs). To allow ONTAP Select to take direct control of NVMe disks, you must configure these drives as passthrough devices within ESXi or KVM. When configuring an NVMe device as a passthrough device, it requires support from the server BIOS and you might need to reboot the host. Additionally, there are limits to the number of passthrough devices that can be assigned per host, which might vary depending on the platform. However, ONTAP Deploy limits this to 14 NVMe devices per ONTAP Select node. This means that the NVMe configuration provides a very high IOPS density (IOPS/TB) at the expense of total capacity. Alternatively, if you desire a high performance configuration with larger storage capacity, the recommended configuration is a large ONTAP Select VM size, an INTEL Optane card for the system disk, and a nominal number of SSD drives for data storage.

To take full advantage of NVMe performance, consider the large ONTAP Select VM size.

There is an additional difference between passthrough devices and RDMs. RDMs can be mapped to a running VM. Passthrough devices require a VM reboot. This means that any NVMe drive replacement or capacity expansion (drive addition) procedure will require an ONTAP Select VM reboot. The drive replacement and capacity expansion (drive addition) operation is driven by a workflow in ONTAP Deploy. ONTAP Deploy manages the ONTAP Select reboot for single node clusters and failover / failback for HA pairs. However, it is important to note the difference between working with SSD data drives (no ONTAP Select reboot / failovers are required) and working with NVMe data drives (ONTAP Select reboot / failover is required).

Physical and virtual disk provisioning

To provide a more streamlined user experience, ONTAP Deploy automatically provisions the system (virtual) disks from the specified datastore (physical system disk) and attaches them to the ONTAP Select VM. This operation occurs automatically during the initial setup so that the ONTAP Select VM can boot. The RDMs are partitioned and the root aggregate is automatically built. If the ONTAP Select node is part of an HA pair, the data partitions are automatically assigned to a local storage pool and a mirror storage pool. This assignment occurs automatically during both cluster-creation operations and storage-add operations.

Because the data disks on the ONTAP Select VM are associated with the underlying physical disks, there are performance implications for creating configurations with a larger number of physical disks.



The root aggregate's RAID group type depends on the number of disks available. ONTAP Deploy picks the appropriate RAID group type. If it has sufficient disks allocated to the node, it uses RAID-DP, otherwise it creates a RAID-4 root aggregate.

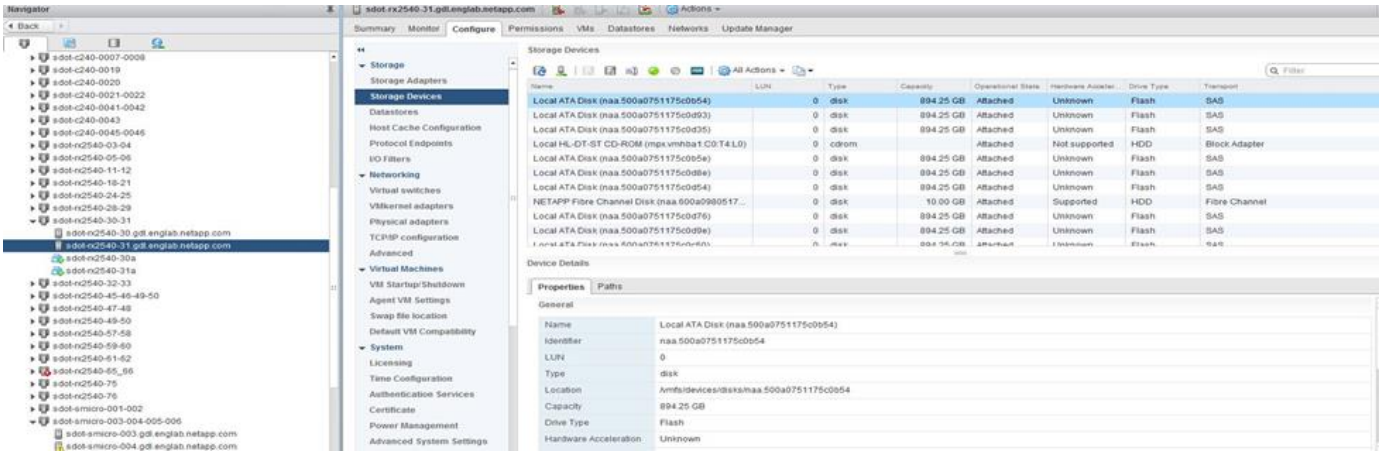
When adding capacity to an ONTAP Select VM using software RAID, the administrator must consider the physical drive size and the number of drives required. For details, see Increase storage capacity.

Similar to FAS and AFF systems, you can only add drives with equal or larger capacities to an existing RAID group. Larger capacity drives are right sized. If you are creating new RAID groups, the new RAID group size should match the existing RAID group size to make sure that the overall aggregate performance does not deteriorate.

Match an ONTAP Select disk to the corresponding ESXi or KVM disk

ONTAP Select disks are usually labeled NET x.y. You can use the following ONTAP command to obtain the disk UUID:

```
<system name>::> disk show NET-1.1
Disk: NET-1.1
Model: Micron_5100_MTFD
Serial Number: 1723175C0B5E
UID:
*500A0751:175C0B5E*:00000000:00000000:00000000:00000000:00000000:00000000:
00000000:00000000
BPS: 512
Physical Size: 894.3GB
Position: shared
Checksum Compatibility: advanced_zoned
Aggregate: -
Plex: -This UID can be matched with the device UID displayed in the
'storage devices' tab for the ESX host
```



In the ESXi or KVM shell, you can enter the following command to blink the LED for a given physical disk (identified by its naa.unique-id).

ESXi

```
esxcli storage core device set -d <naa_id> -l=locator -L=<seconds>
```

KVM

```
cat /sys/block/<block_device_id>/device/wwid
```

Multiple drive failures when using software RAID

It is possible for a system to encounter a situation in which multiple drives are in a failed state at the same time. The behavior of the system depends on the aggregate RAID protection and the number of failed drives.

A RAID4 aggregate can survive one disk failure, a RAID-DP aggregate can survive two disk failures, and a RAID-TEC aggregate can survive three disks failures.

If the number of failed disks is less than the maximum number of failures that RAID type supports, and if a spare disk is available, the reconstruction process starts automatically. If spare disks are not available, the aggregate serves data in a degraded state until spare disks are added.

If the number of failed disks is more than the maximum number of failures that the RAID type supports, then the local plex is marked as failed, and the aggregate state is degraded. Data is served from the second plex residing on the HA partner. This means that any I/O requests for node 1 are sent through cluster interconnect port e0e (iSCSI) to the disks physically located on node 2. If the second plex also fails, then the aggregate is marked as failed and data is unavailable.

A failed plex must be deleted and recreated for the correct mirroring of data to resume. Note that a multi-disk failure resulting in a data aggregate being degraded also results in a root aggregate being degraded. ONTAP Select uses the root-data-data (RDD) partitioning schema to split each physical drive into a root partition and two data partitions. Therefore, losing one or more disks might impact multiple aggregates, including the local root or the copy of the remote root aggregate, as well as the local data aggregate and the copy of the remote data aggregate.

A failed plex is deleted and recreated in the following example output:

```
C3111E67::> storage aggregate plex delete -aggregate aggr1 -plex plex1
Warning: Deleting plex "plex1" of mirrored aggregate "aggr1" in a non-
shared HA configuration will disable its synchronous mirror protection and
disable
         negotiated takeover of node "sti-rx2540-335a" when aggregate
"aggr1" is online.
Do you want to continue? {y|n}: y
[Job 78] Job succeeded: DONE

C3111E67::> storage aggregate mirror -aggregate aggr1
Info: Disks would be added to aggregate "aggr1" on node "sti-rx2540-335a"
in the following manner:
      Second Plex
        RAID Group rg0, 5 disks (advanced_zoned checksum, raid_dp)
                                                                  Usable
Physical
          Position   Disk                      Type           Size
Size
          ---------- ------------------------- ---------- --------
--------
          shared     NET-3.2                   SSD               -
-
          shared     NET-3.3                   SSD               -
-
          shared     NET-3.4                   SSD         208.4GB
208.4GB
          shared     NET-3.5                   SSD         208.4GB
208.4GB
          shared     NET-3.12                  SSD         208.4GB
208.4GB

      Aggregate capacity available for volume use would be 526.1GB.
      625.2GB would be used from capacity license.
Do you want to continue? {y|n}: y

C3111E67::> storage aggregate show-status -aggregate aggr1
Owner Node: sti-rx2540-335a
 Aggregate: aggr1 (online, raid_dp, mirrored) (advanced_zoned checksums)
  Plex: /aggr1/plex0 (online, normal, active, pool0)
   RAID Group /aggr1/plex0/rg0 (normal, advanced_zoned checksums)
                                                                  Usable
Physical
     Position Disk                      Pool Type     RPM     Size
Size Status
     -------- ------------------------- ---- ----- ------ --------
```

```
-------- ----------
     shared   NET-1.1                       0   SSD        -  205.1GB
447.1GB (normal)
     shared   NET-1.2                       0   SSD        -  205.1GB
447.1GB (normal)
     shared   NET-1.3                       0   SSD        -  205.1GB
447.1GB (normal)
     shared   NET-1.10                      0   SSD        -  205.1GB
447.1GB (normal)
     shared   NET-1.11                      0   SSD        -  205.1GB
447.1GB (normal)
  Plex: /aggr1/plex3 (online, normal, active, pool1)
   RAID Group /aggr1/plex3/rg0 (normal, advanced_zoned checksums)
                                                                 Usable
Physical
     Position Disk                      Pool Type     RPM     Size
Size Status
     -------- --------------------------- ---- ----- ------ --------
-------- ----------
     shared   NET-3.2                       1   SSD        -  205.1GB
447.1GB (normal)
     shared   NET-3.3                       1   SSD        -  205.1GB
447.1GB (normal)
     shared   NET-3.4                       1   SSD        -  205.1GB
447.1GB (normal)
     shared   NET-3.5                       1   SSD        -  205.1GB
447.1GB (normal)
     shared   NET-3.12                      1   SSD        -  205.1GB
447.1GB (normal)
10 entries were displayed..
```

In order to test or simulate one or multiple drive failures, use the `storage disk fail -disk NET-x.y -immediate` command. If there is a spare in the system, the aggregate will begin to reconstruct. You can check the status of the reconstruction using the command `storage aggregate show`. You can remove the simulated failed drive using ONTAP Deploy. Note that ONTAP has marked the drive as `Broken`. The drive is not actually broken and can be added back using ONTAP Deploy. In order to erase the Broken label, enter the following commands in the ONTAP Select CLI:

```
set advanced
disk unfail -disk NET-x.y -spare true
disk show -broken
```

The output for the last command should be empty.

Virtualized NVRAM

NetApp FAS systems are traditionally fitted with a physical NVRAM PCI card. This card is a high-performing card containing nonvolatile flash memory that provides a significant boost in write performance. It does this by granting ONTAP the ability to immediately acknowledge incoming writes back to the client. It can also schedule the movement of modified data blocks back to slower storage media in a process known as destaging.

Commodity systems are not typically fitted with this type of equipment. Therefore, the functionality of the NVRAM card has been virtualized and placed into a partition on the ONTAP Select system boot disk. It is for this reason that placement of the system virtual disk of the instance is extremely important.

ONTAP Select vSAN and external array configurations

Virtual NAS (vNAS) deployments support ONTAP Select clusters on virtual SAN (vSAN), some HCI products, and external array types of datastores. The underlying infrastructure of these configurations provide datastore resiliency.

The minimum requirement is that the hypervisor that you are using (VMware ESXi or KVM on a supported Linux host) supports the underlying configuration. If the hypervisor is ESXi, it should be listed on the respective VMware HCLs.

vNAS architecture

The vNAS nomenclature is used for all setups that do not use DAS. For multi-node ONTAP Select clusters, this includes architectures for which the two ONTAP Select nodes in the same HA pair share a single datastore (including vSAN datastores). The nodes can also be installed on separate datastores from the same shared external array. This allows for array-side storage efficiencies to reduce the overall footprint of the entire ONTAP Select HA pair. The architecture of ONTAP Select vNAS solutions is very similar to that of ONTAP Select on DAS with a local RAID controller. That is to say that each ONTAP Select node continues to have a copy of its HA partner's data. ONTAP storage efficiency policies are node scoped. Therefore, array side storage efficiencies are preferable because they can potentially be applied across data sets from both ONTAP Select nodes.

It is also possible that each ONTAP Select node in an HA pair uses a separate external array. This is a common choice when using ONTAP Select Metrocluster SDS with external storage.

When using separate external arrays for each ONTAP Select node, it is very important that the two arrays provide similar performance characteristics to the ONTAP Select VM.

vNAS architectures versus local DAS with hardware RAID controllers

The vNAS architecture is logically most similar to the architecture of a server with DAS and a RAID controller. In both cases, ONTAP Select consumes datastore space. That datastore space is carved into VMDKs, and these VMDKs form the traditional ONTAP data aggregates. ONTAP Deploy makes sure that the VMDKs are properly sized and assigned to the correct plex (in the case of HA pairs) during cluster -create and storage-add operations.

There are two major differences between vNAS and DAS with a RAID controller. The most immediate difference is that vNAS does not require a RAID controller. vNAS assumes that the underlying external array provides the data persistence and resiliency that a DAS with a RAID controller setup would provide. The second and more subtle difference has to do with NVRAM performance.

vNAS NVRAM

The ONTAP Select NVRAM is a VMDK. This means ONTAP Select emulates a byte addressable space (traditional NVRAM) on top of a block addressable device (VMDK). However, the performance of the NVRAM is absolutely critical to the overall performance of the ONTAP Select node.

For DAS setups with a hardware RAID controller, the hardware RAID controller cache acts as the NVRAM cache, because all writes to the NVRAM VMDK are first hosted in the RAID controller cache.

For VNAS architectures, ONTAP Deploy automatically configures ONTAP Select nodes with a boot argument called Single Instance Data Logging (SIDL). When this boot argument is present, ONTAP Select bypasses the NVRAM and writes the data payload directly to the data aggregate. The NVRAM is only used to record the address of the blocks changed by the WRITE operation. The benefit of this feature is that it avoids a double write: one write to NVRAM and a second write when the NVRAM is destaged. This feature is only enabled for vNAS because local writes to the RAID controller cache have a negligible additional latency.

The SIDL feature is not compatible with all ONTAP Select storage efficiency features. The SIDL feature can be disabled at the aggregate level using the following command:

```
storage aggregate modify -aggregate aggr-name -single-instance-data
-logging off
```

Write performance is affected if the SIDL feature is turned off. It is possible to re-enable the SIDL feature after all the storage efficiency policies on all the volumes in that aggregate are disabled:

```
volume efficiency stop -all true -vserver * -volume * (all volumes in the
affected aggregate)
```

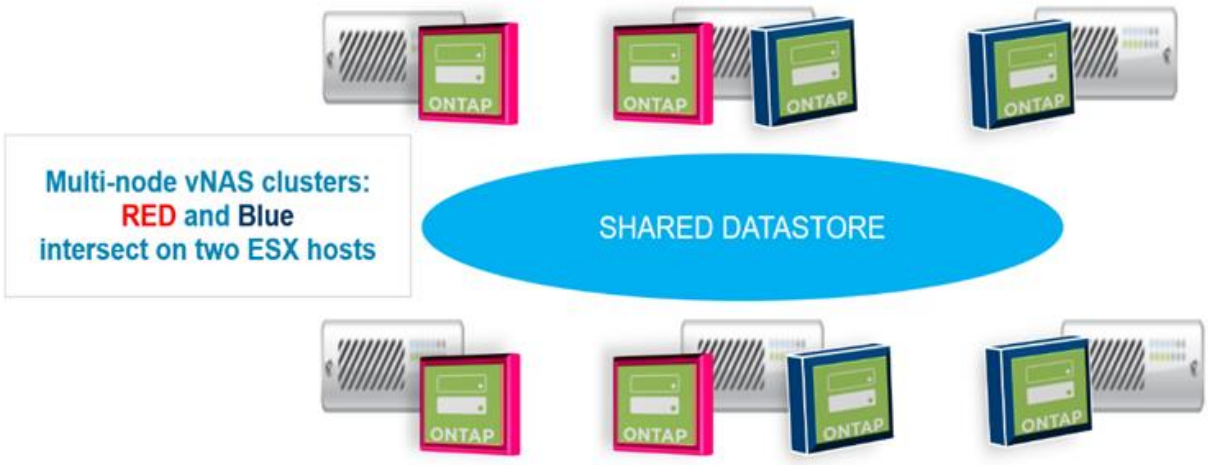
Collocate ONTAP Select Nodes When Using vNAS on ESXi

ONTAP Select includes support for multi-node ONTAP Select clusters on shared storage. ONTAP Deploy enables the configuration of multiple ONTAP Select nodes on the same ESXi host as long as these nodes are not part of the same cluster.

This configuration is only valid for VNAS environments (shared datastores). Multiple ONTAP Select instances per host are not supported when using DAS storage because these instances compete for the same hardware RAID controller.

ONTAP Deploy makes sure that the initial deployment of the multi-node VNAS cluster does not place multiple ONTAP Select instances from the same cluster on the same host. The following figure illustrates an example of a correct deployment of two four-node clusters that intersect on two hosts.

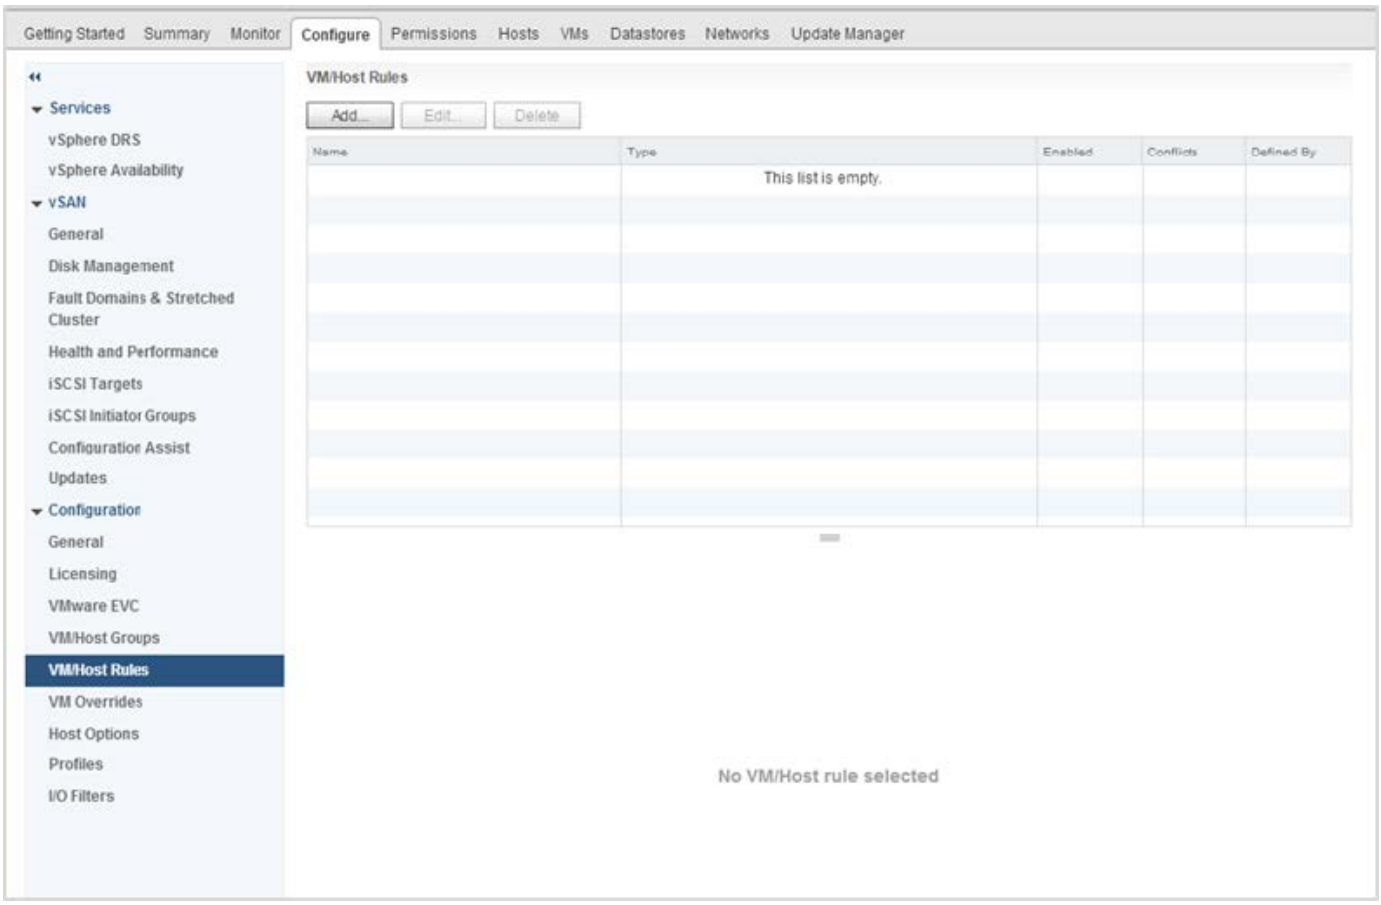
Initial deployment of multi-node VNAS clusters

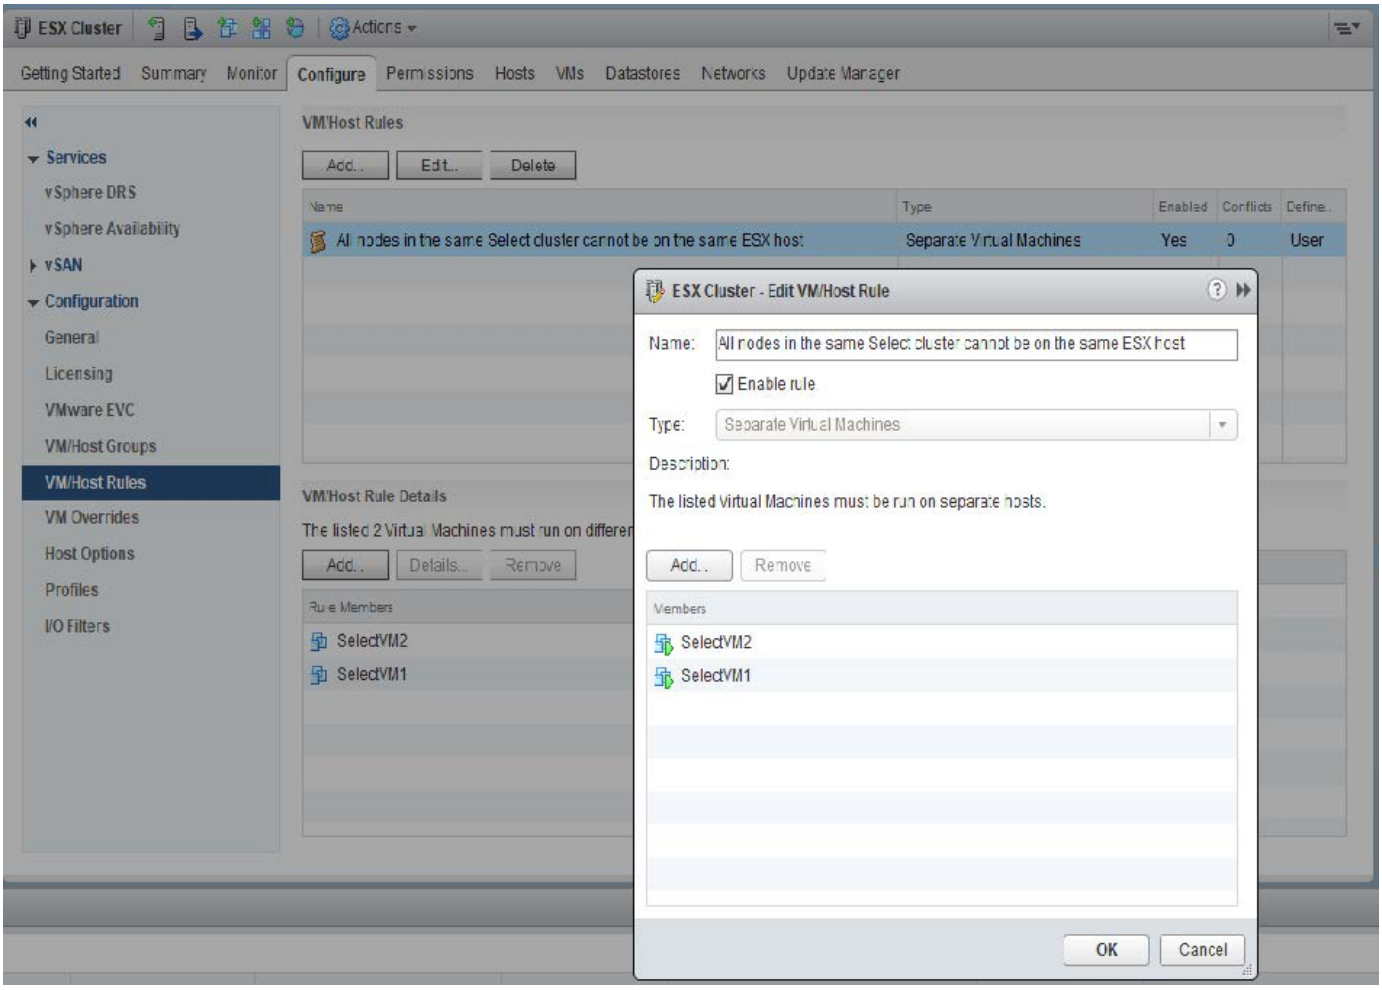


After deployment, the ONTAP Select nodes can be migrated between hosts. This could result in nonoptimal and unsupported configurations for which two or more ONTAP Select nodes from the same cluster share the same underlying host. NetApp recommends the manual creation of VM anti-affinity rules so that VMware automatically maintains physical separation between the nodes of the same cluster, not just the nodes from the same HA pair.

Anti-affinity rules require that DRS is enabled on the ESXi cluster.

Refer to the following example on how to create an anti-affinity rule for the ONTAP Select VMs. If the ONTAP Select cluster contains more than one HA pair, all nodes in the cluster must be included in this rule.





Two or more ONTAP Select nodes from the same ONTAP Select cluster could potentially be found on the same ESXi host for one of the following reasons:

- DRS is not present due to VMware vSphere license limitations or if DRS is not enabled.
- The DRS anti-affinity rule is bypassed because a VMware HA operation or administrator-initiated VM migration takes precedence.

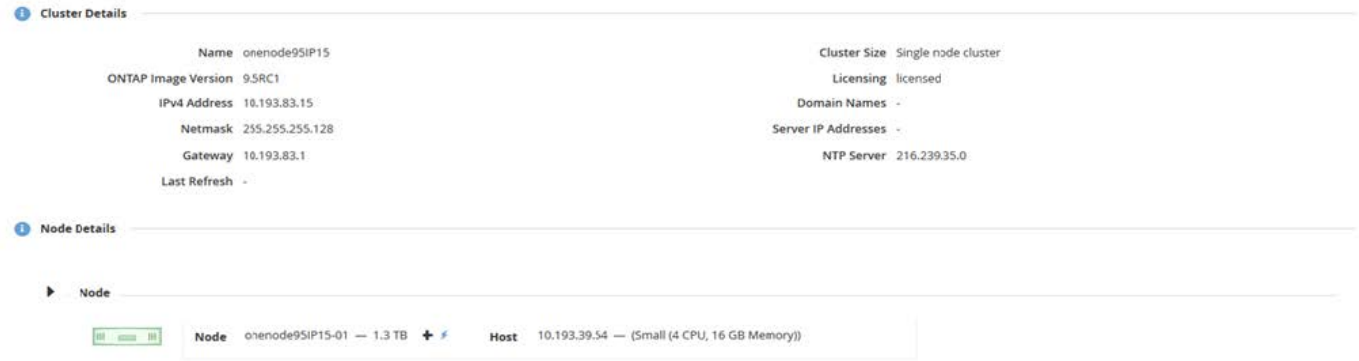
ONTAP Deploy does not proactively monitor the ONTAP Select VM locations. However, a cluster refresh operation reflects this unsupported configuration in the ONTAP Deploy logs:



Considerations for increasing ONTAP Select storage capacity

ONTAP Deploy can be used to add and license additional storage for each node in an ONTAP Select cluster.

The storage-add functionality in ONTAP Deploy is the only way to increase the storage under management, and directly modifying the ONTAP Select VM is not supported. The following figure shows the “+” icon that initiates the storage-add wizard.



The following considerations are important for the success of the capacity-expansion operation. Adding capacity requires the existing license to cover the total amount of space (existing plus new). A storage-add operation that results in the node exceeding its licensed capacity fails. A new license with sufficient capacity should be installed first.

If the extra capacity is added to an existing ONTAP Select aggregate, then the new storage pool (datastore) should have a performance profile similar to that of the existing storage pool (datastore). Note that it is not possible to add non-SSD storage to an ONTAP Select node installed with an AFF-like personality (flash enabled). Mixing DAS and external storage is also not supported.

If locally attached storage is added to a system to provide for additional local (DAS) storage pools, you must build an additional RAID group and LUN (or LUNs). Just as with FAS systems, care should be taken to make sure that the new RAID group performance is similar to that of the original RAID group if you are adding new space to the same aggregate. If you are creating a new aggregate, the new RAID group layout could be different if the performance implications for the new aggregate are well understood.

The new space can be added to that same data store as an extent if the total size of the data store does not exceed the supported maximum data store size. Adding a data store extent to the data store in which ONTAP Select is already installed can be done dynamically and does not affect the operations of the ONTAP Select node.

If the ONTAP Select node is part of an HA pair, some additional issues should be considered.

In an HA pair, each node contains a mirror copy of the data from its partner. Adding space to node 1 requires that an identical amount of space is added to its partner, node 2, so that all the data from node 1 is replicated to node 2. In other words, the space added to node 2 as part of the capacity-add operation for node 1 is not visible or accessible on node 2. The space is added to node 2 so that node 1 data is fully protected during an HA event.

There is an additional consideration with regard to performance. The data on node 1 is synchronously replicated to node 2. Therefore, the performance of the new space (datastore) on node 1 must match the performance of the new space (datastore) on node 2. In other words, adding space on both nodes, but using different drive technologies or different RAID group sizes, can lead to performance issues. This is due to the RAID SyncMirror operation used to maintain a copy of the data on the partner node.

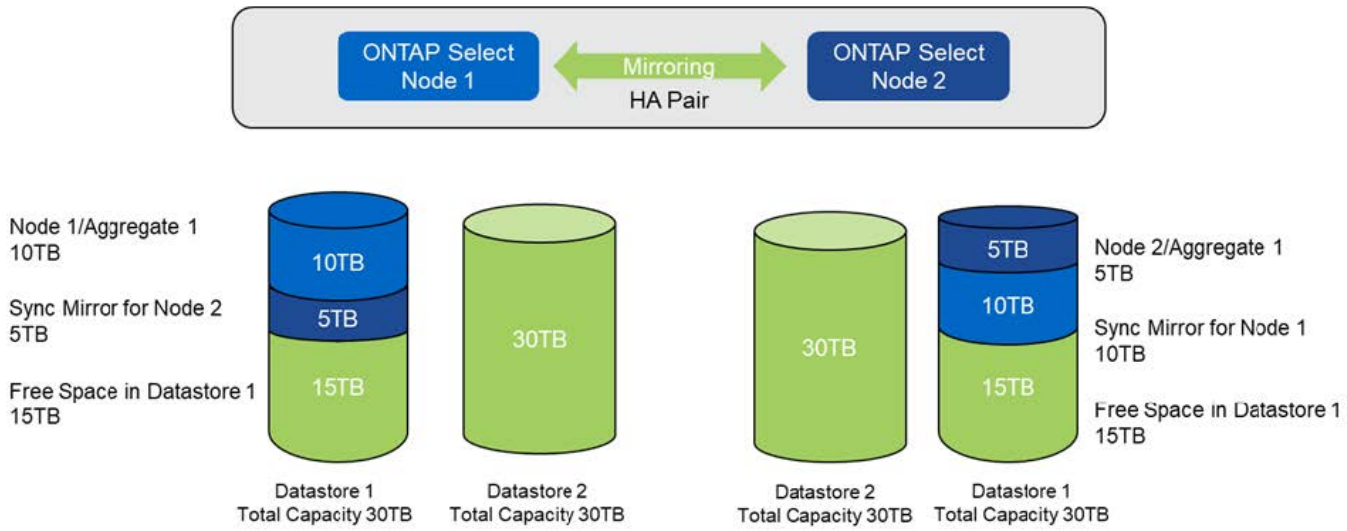
To increase user-accessible capacity on both nodes in an HA pair, two storage-add operations must be performed, one for each node. Each storage-add operation requires additional space on both nodes. The total space required on each node is equal to the space required on node 1 plus the space required on node 2.

Initial setup is with two nodes, each node having two datastores with 30TB of space in each datastore. ONTAP Deploy creates a two-node cluster, with each node consuming 10TB of space from datastore 1. ONTAP Deploy

configures each node with 5TB of active space per node.

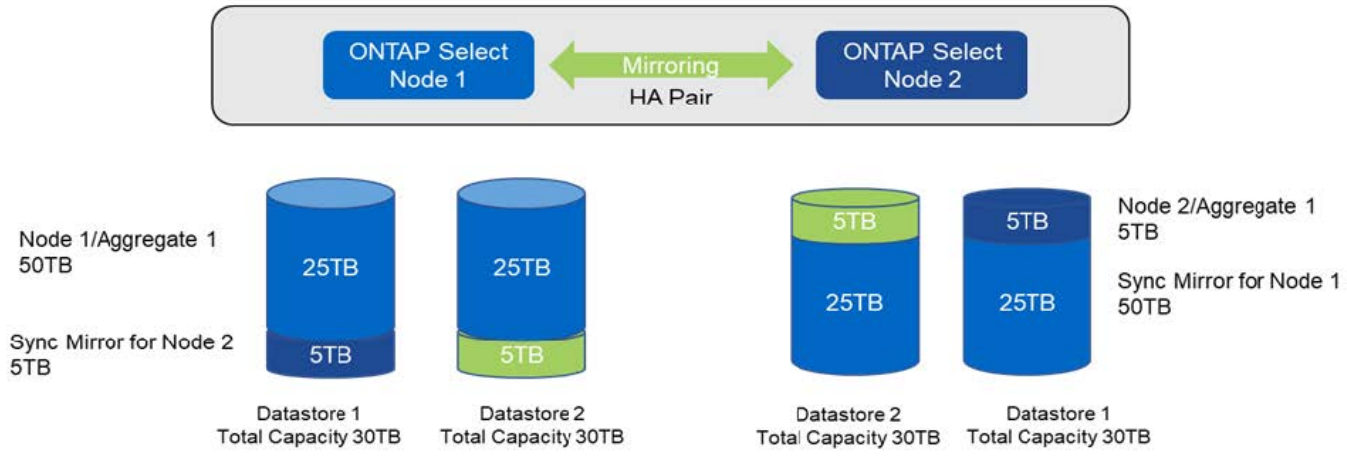
The following figure shows the results of a single storage-add operation for node 1. ONTAP Select still uses an equal amount of storage (15TB) on each node. However, node 1 has more active storage (10TB) than node 2 (5TB). Both nodes are fully protected as each node hosts a copy of the other node's data. There is additional free space left in datastore 1, and datastore 2 is still completely free.

Capacity distribution: allocation and free space after a single storage-add operation



Two additional storage-add operations on node 1 consume the rest of datastore 1 and a part of datastore 2 (using the capacity cap). The first storage-add operation consumes the 15TB of free space left in datastore 1. The following figure shows the result of the second storage-add operation. At this point, node 1 has 50TB of active data under management, while node 2 has the original 5TB.

Capacity distribution: allocation and free space after two additional storage-add operation for node 1



The maximum VMDK size used during capacity add operations is 16TB. The maximum VMDK size used during cluster create operations continues to be 8TB. ONTAP Deploy creates correctly sized VMDKs depending on your configuration (a single-node or multi-node cluster) and the amount of capacity being added. However, the maximum size of each VMDK should not exceed 8TB during the cluster create operations and 16TB during the storage-add operations.

Increase capacity for ONTAP Select with Software RAID

The storage-add wizard can similarly be used to increase capacity under management for ONTAP Select nodes using software RAID. The wizard only presents those DAS SDD drives that are available and can be mapped as RDMs to the ONTAP Select VM.

Though it is possible to increase the capacity license by a single TB, when working with software RAID, it is not possible to physically increase the capacity by a single TB. Similar to adding disks to a FAS or AFF array, certain factors dictate the minimum amount of storage that can be added in a single operation.

> In an HA pair, adding storage to node 1 requires that an identical number of drives is also available on the node's HA pair (node 2). Both the local drives and the remote disks are used by one storage-add operation on node 1. That is to say, the remote drives are used to make sure that the new storage on node 1 is replicated and protected on node 2. In order to add locally usable storage on node 2, a separate storage-add operation and a separate and equal number of drives must be available on both nodes.

ONTAP Select partitions any new drives into the same root, data, and data partitions as the existing drives. The partitioning operation takes place during the creation of a new aggregate or during the expansion on an existing aggregate. The size of the root partition stripe on each disk is set to match the existing root partition size on the existing disks. Therefore, each one of the two equal data partition sizes can be calculated as the disk total capacity minus the root partition size divided by two. The root partition stripe size is variable, and it is computed during the initial cluster setup as follows. Total root space required (68GB for a single-node cluster and 136GB for HA pairs) is divided across the initial number of disks minus any spare and parity drives. The root partition stripe size is maintained to be constant on all the drives being added to the system.

If you are creating a new aggregate, the minimum number of drives required varies depending on the RAID type and whether the ONTAP Select node is part of an HA pair.

If adding storage to an existing aggregate, some additional considerations are necessary. It is possible to add drives to an existing RAID group, assuming that the RAID group is not at the maximum limit already. Traditional FAS and AFF best practices for adding spindles to existing RAID groups also apply here, and creating a hot spot on the new spindle is a potential concern. In addition, only drives of equal or larger data partition sizes can be added to an existing RAID group. As explained above, the data partition size is not the same as drive raw size. If the data partitions being added are larger than the existing partitions, the new drives is right-sized. In other words, a portion of capacity of each new drive remains unutilized.

It is also possible to use the new drives to create a new RAID group as part of an existing aggregate. In this case, the RAID group size should match the existing RAID group size.

ONTAP Select storage efficiency support

ONTAP Select provides storage efficiency options that are similar to the storage efficiency options present on FAS and AFF arrays.

ONTAP Select virtual NAS (vNAS) deployments using all-flash VSAN or generic flash arrays should follow the best practices for ONTAP Select with non-SSD direct-attached storage (DAS).

An AFF-like personality is automatically enabled on new installations as long as you have DAS storage with SSD drives and a premium license.

With an AFF-like personality, the following inline SE features are automatically enabled during installation:

- Inline zero pattern detection

- Volume inline deduplication
- Volume background deduplication
- Adaptive inline compression
- Inline data compaction
- Aggregate inline deduplication
- Aggregate background deduplication

To verify that ONTAP Select has enabled all the default storage efficiency policies, run the following command on a newly created volume:

```
<system name>::> set diag
Warning: These diagnostic commands are for use by NetApp personnel only.
Do you want to continue? {y|n}: y
twonode95IP15::*> sis config
Vserver:                                SVM1
Volume:                                 _export1_NFS_volume
Schedule                                -
Policy:                                 auto
Compression:                            true
Inline Compression:                     true
Compression Type:                       adaptive
Application IO Si                       8K
Compression Algorithm:                  lzopro
Inline Dedupe:                          true
Data Compaction:                        true
Cross Volume Inline Deduplication:      true
Cross Volume Background Deduplication:  true
```

For ONTAP Select upgrades from 9.6 and later, you must install ONTAP Select on DAS SSD storage with a premium license. In addition, you must check the **Enable Storage Efficiencies** check box during the initial cluster installation with ONTAP Deploy. Enabling an AFF-like personality post-ONTAP upgrade when prior conditions have not been met requires the manual creation of a boot argument and a node reboot. Contact technical support for further details.

ONTAP Select storage efficiency configurations

The following table summarizes the various storage efficiency options available, enabled by default, or not enabled by default but recommended, depending on the media type and software license.

ONTAP Select features	DAS SSD (premium or premium XL[1])	DAS HDD (all licenses)	vNAS (all licenses)
Inline zero detection	Yes (default)	Yes Enabled by user on a per-volume basis	Yes Enabled by user on a per-volume basis
Volume inline deduplication	Yes (default)	Not available	Not supported

ONTAP Select features	DAS SSD (premium or premium XL[1])	DAS HDD (all licenses)	vNAS (all licenses)
32K inline compression (secondary compression)	Yes Enabled by user on a per volume basis.	Yes Enabled by user on a per-volume basis	Not supported
8K inline compression (adaptive compression)	Yes (default)	Yes Enabled by user on a per volume basis	Not supported
Background compression	Not supported	Yes Enabled by user on a per volume basis	Yes Enabled by user on a per-volume basis
Compression scanner	Yes	Yes	Yes Enabled by user on a per-volume basis
Inline data compaction	Yes (default)	Yes Enabled by user on a per volume basis	Not supported
Compaction scanner	Yes	Yes	Not supported
Aggregate inline deduplication	Yes (default)	N/A	Not supported
Volume background deduplication	Yes (default)	Yes Enabled by user on a per volume basis	Yes Enabled by user on a per-volume basis
Aggregate background deduplication	Yes (default)	N/A	Not supported

[1]ONTAP Select 9.6 supports a new license (premium XL) and a new VM size (large). However, the large VM is only supported for DAS configurations using software RAID. Hardware RAID and vNAS configurations are not supported with the large ONTAP Select VM in the 9.6 release.

Notes on upgrade behavior for DAS SSD configurations

After upgrading to ONTAP Select 9.6 or later, wait for the `system node upgrade-revert show` command to indicate that the upgrade has completed before verifying the storage efficiency values for existing volumes.

On a system upgraded to ONTAP Select 9.6 or later, a new volume created on an existing aggregate or a newly created aggregate has the same behavior as a volume created on a fresh deployment. Existing volumes that undergo the ONTAP Select code upgrade have most of the same storage efficiency policies as a newly created volume with some variations:

Scenario 1

If no storage efficiency policies were enabled on a volume prior to the upgrade, then:

- Volumes with `space guarantee = volume` do not have inline data-compaction, aggregate inline deduplication, and aggregate background deduplication enabled. These options can be enabled post-upgrade.
- Volumes with `space guarantee = none` do not have background compression enabled. This option can be enabled post upgrade.
- Storage efficiency policy on the existing volumes is set to auto after upgrade.

Scenario 2

If some storage efficiencies are already enabled on a volume prior to the upgrade, then:

- Volumes with `space guarantee = volume` do not see any difference after upgrade.

- Volumes with `space guarantee = none` have aggregate background deduplication turned on.
- Volumes with `storage policy inline-only` have their policy set to auto.
- Volumes with user defined storage efficiency policies have no change in policy, with the exception of volumes with `space guarantee = none`. These volumes have aggregate background deduplication enabled.

Networking

ONTAP Select networking concepts and characteristics

First become familiar with general networking concepts that apply to the ONTAP Select environment. Then explore the specific characteristics and options available with the single-node and multi-node clusters.

Physical networking

The physical network supports an ONTAP Select cluster deployment primarily by providing the underlying layer two switching infrastructure. The configuration related to the physical network includes both the hypervisor host and the broader switched network environment.

Host NIC options

Each ONTAP Select hypervisor host must be configured with either two or four physical ports. The exact configuration you choose depends on several factors, including:

- Whether the cluster contains one or multiple ONTAP Select hosts
- What hypervisor operating system is used
- How the virtual switch is configured
- Whether LACP is used with the links or not

Physical switch configuration

You must make sure that the configuration of the physical switches supports the ONTAP Select deployment. The physical switches are integrated with the hypervisor-based virtual switches. The exact configuration you choose depends on several factors. The primary considerations include the following:

- How will you maintain separation between the internal and external networks?
- Will you maintain a separation between the data and management networks?
- How will the layer two VLANs be configured?

Logical networking

ONTAP Select uses two different logical networks, separating the traffic according to type. Specifically, traffic can flow among the hosts within the cluster as well as to the storage clients and other machines outside of the cluster. The virtual switches managed by the hypervisors help support the logical network.

Internal network

With a multi-node cluster deployment, the individual ONTAP Select nodes communicate using an isolated “internal” network. This network is not exposed or available outside of the nodes in the ONTAP Select cluster.

The internal network is only present with a multi-node cluster.

The internal network has the following characteristics:

- Used to process ONTAP intra-cluster traffic including:
  - Cluster
  - High Availability Interconnect (HA-IC)
  - RAID Synch Mirror (RSM)
- Single layer-two network based on a VLAN
- Static IP addresses are assigned by ONTAP Select:
  - IPv4 only
  - DHCP not used
  - Link-local address
- The MTU size is 9000 bytes by default and can be adjusted within 7500-9000 range (inclusive)

External network

The external network processes traffic between the nodes of an ONTAP Select cluster and the external storage clients as well as the other machines. The external network is a part of every cluster deployment and has the following characteristics:

- Used to process ONTAP traffic including:
  - Data (NFS, CIFS, iSCSI)
  - Management (cluster and node; optionally SVM)
  - Intercluster (optional)
- Optionally supports VLANs:
  - Data port group
  - Management port group
- IP addresses that are assigned based on the configuration choices of the administrator:
  - IPv4 or IPv6
- MTU size is 1500 bytes by default (can be adjusted)

The external network is present with clusters of all sizes.

Virtual machine networking environment

The hypervisor host provides several networking features.

ONTAP Select relies on the following capabilities exposed through the virtual machine:

Virtual machine ports

There are several ports available for use by ONTAP Select. They are assigned and used based on several factors, including the size of the cluster.

Virtual switch

The virtual switch software within the hypervisor environment, whether vSwitch (VMware) or Open vSwitch (KVM), joins the ports exposed by the virtual machine with the physical Ethernet NIC ports. You must configure a vSwitch for every ONTAP Select host, as appropriate for your environment.

ONTAP Select single and multiple node network configurations

ONTAP Select supports both single-node and multi-node network configurations.

Single node network configuration

Single-node ONTAP Select configurations do not require the ONTAP internal network, because there is no cluster, HA, or mirror traffic.

Unlike the multi-node version of the ONTAP Select product, each ONTAP Select VM contains three virtual network adapters, presented to ONTAP network ports e0a, e0b, and e0c.

These ports are used to provide the following services: management, data, and intercluster LIFs.

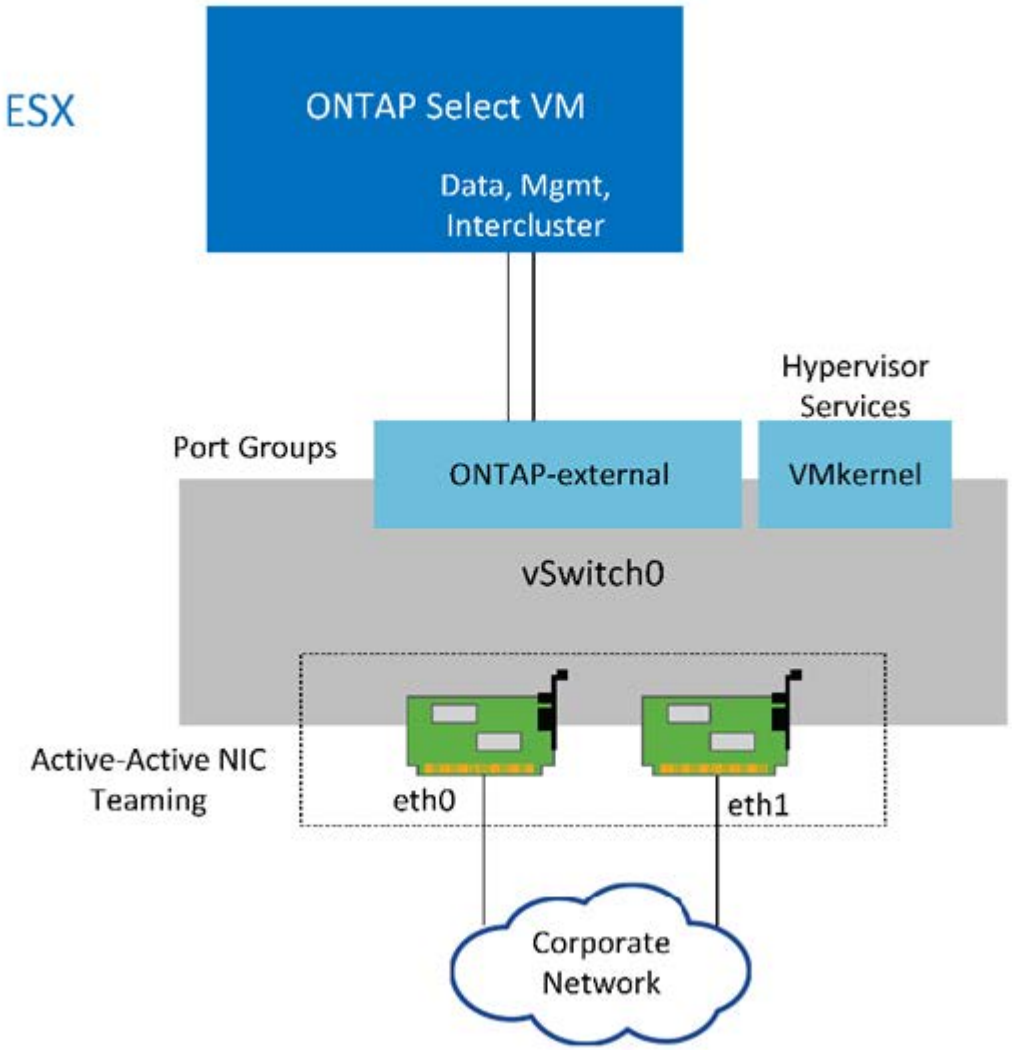
KVM

You can deploy ONTAP Select as a single-node cluster. The hypervisor host includes a virtual switch that provides access to the external network.

ESXi

The following figure shows the relationship between these ports and the underlying physical adapters. The figure depicts one ONTAP Select cluster node on the ESXi hypervisor.

Network configuration of single-node ONTAP Select cluster



Even though two adapters are sufficient for a single-node cluster, NIC teaming is still required.

LIF assignment

As explained in the multi-node LIF assignment section of this document, ONTAP uses IPspaces to keep cluster network traffic separate from data and management traffic. The single-node variant of this platform does not contain a cluster network. Therefore, there are no ports present in the cluster IPspace.

Cluster and node management LIFs are automatically created during ONTAP Select cluster setup. You can create the remaining LIFs post deployment.

Management and data LIFs (e0a, e0b, and e0c)

ONTAP ports e0a, e0b, and e0c are delegated as candidate ports for LIFs that carry the following types of traffic:

- SAN/NAS protocol traffic (CIFS, NFS, and iSCSI)
- Cluster, node, and SVM management traffic
- Intercluster traffic (SnapMirror and SnapVault)

Multi-node network configuration

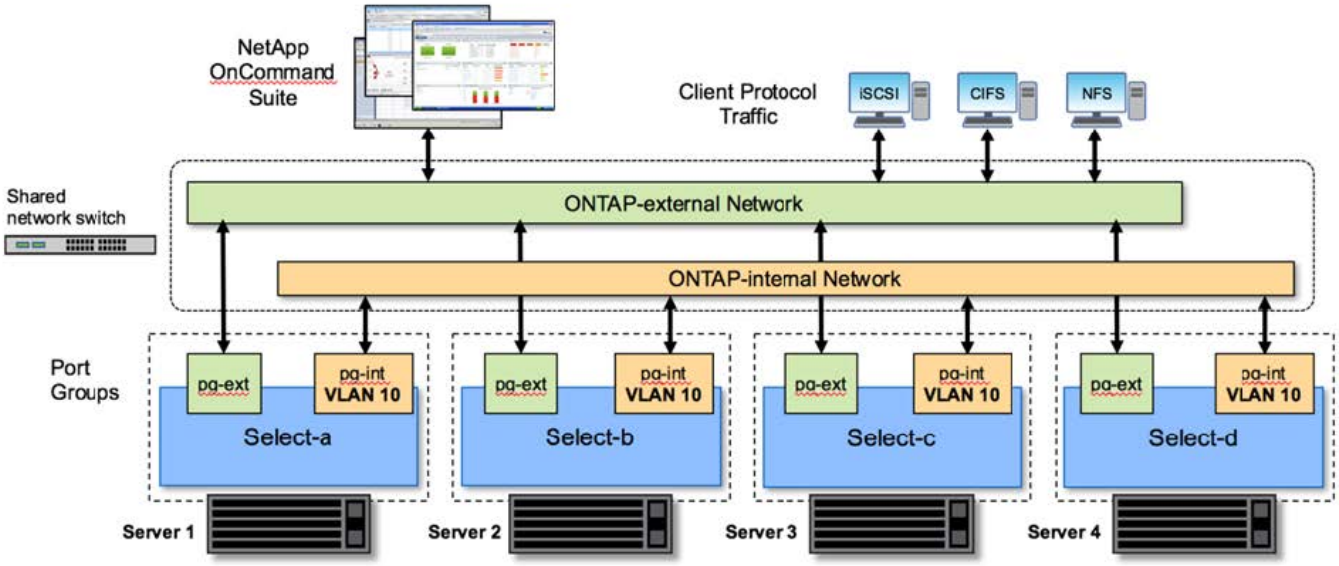
The multi-node ONTAP Select network configuration consists of two networks.

These are an internal network, responsible for providing cluster and internal replication services, and an external network, responsible for providing data access and management services. End-to-end isolation of traffic that flows within these two networks is extremely important for building an environment that is suitable for cluster resiliency.

These networks are represented in the following figure, which shows a four-node ONTAP Select cluster running on a VMware vSphere platform. Six, eight, ten, and twelve-node clusters have a similar network layout.

Each ONTAP Select instance resides on a separate physical server. Internal and external traffic is isolated using separate network port groups, which are assigned to each virtual network interface and allow the cluster nodes to share the same physical switch infrastructure.

Overview of an ONTAP Select multi-node cluster network configuration



Each ONTAP Select VM contains seven virtual network adapters presented to ONTAP as a set of seven network ports, e0a through e0g. Although ONTAP treats these adapters as physical NICs, they are in fact virtual and map to a set of physical interfaces through a virtualized network layer. As a result, each hosting server does not require six physical network ports.

Adding virtual network adapters to the ONTAP Select VM is not supported.

These ports are preconfigured to provide the following services:

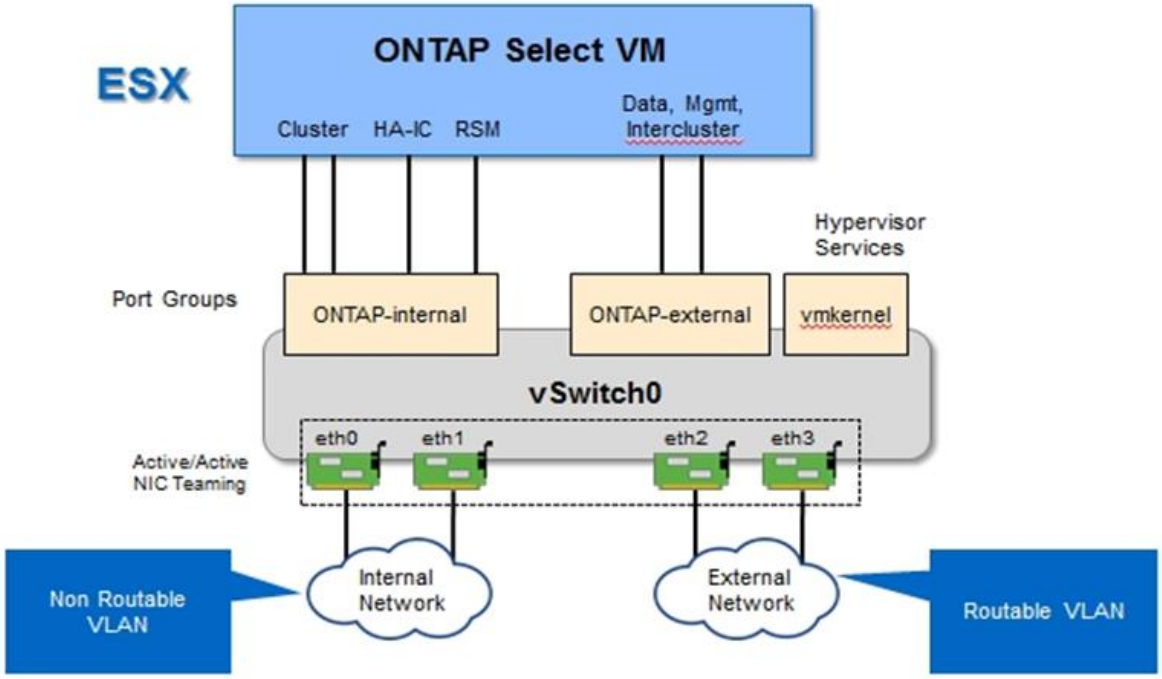
- e0a, e0b, and e0g. Management and data LIFs
- e0c, e0d. Cluster network LIFs
- e0e. RSM
- e0f. HA interconnect

Ports e0a, e0b, and e0g reside on the external network. Although ports e0c through e0f perform several different functions, collectively they compose the internal Select network. When making network design

decisions, you should place these ports on a single layer-2 network. There is no need to separate these virtual adapters across different networks.

The relationship between these ports and the underlying physical adapters is illustrated in the following figure, which depicts one ONTAP Select cluster node on the ESXi hypervisor.

Network configuration of a single node that is part of a multi-node ONTAP Select cluster



Segregating internal and external traffic across different physical NICs prevents insufficient access to network resources, introducing latencies into the system. Additionally, aggregation through NIC teaming allows the ONTAP Select cluster node to continue accessing the network if a single network adapter fails.

Both the external network and internal network port groups contain all four NIC adapters in a symmetrical manner. The active ports in the external network port group are the standby ports in the internal network. Conversely, the active ports in the internal network port group are the standby ports in the external network port group.

LIF assignment

With the introduction of IPspaces, ONTAP port roles have been deprecated. Like FAS arrays, ONTAP Select clusters contain both a default IPspace and a cluster IPspace. By placing network ports e0a, e0b, and e0g into the default IPspace and ports e0c and e0d into the cluster IPspace, those ports have essentially been walled off from hosting LIFs that do not belong. The remaining ports within the ONTAP Select cluster are consumed through the automatic assignment of interfaces providing internal services. They are not exposed through the ONTAP shell, as is the case with the RSM and HA interconnect interfaces.

Not all LIFs are visible through the ONTAP command shell. The HA interconnect and RSM interfaces are hidden from ONTAP and are used internally to provide their respective services.

The network ports and LIFs are explained in detail in the following sections.

Management and data LIFs (e0a, e0b, and e0g)

ONTAP ports e0a, e0b, and e0g are delegated as candidate ports for LIFs that carry the following types of traffic:

- SAN/NAS protocol traffic (CIFS, NFS, and iSCSI)
- Cluster, node, and SVM management traffic
- Intercluster traffic (SnapMirror and SnapVault)

> Cluster and node management LIFs are automatically created during ONTAP Select cluster setup. You can create the remaining LIFs post deployment.

Cluster network LIFs (e0c, e0d)

ONTAP ports e0c and e0d are delegated as home ports for cluster interfaces. Within each ONTAP Select cluster node, two cluster interfaces are automatically generated during ONTAP setup using link local IP addresses (169.254.x.x).

> You cannot assign a static IP address to these interfaces, and you should not create additional cluster interfaces.

Cluster network traffic must flow through a low-latency, non-routed layer-2 network. Due to cluster throughput and latency requirements, you should physically locate the ONTAP Select cluster within proximity (for example, multipack, single data center). Building four, six, eight, ten, or twelve-node stretch cluster configurations by separating HA nodes across a WAN or significant geographical distances is not supported. A stretched two-node configuration with a mediator is supported.

For details, refer to the section Two-node stretched HA (MetroCluster SDS) best practices.

> To make sure of maximum throughput for cluster network traffic, this network port is configured to use jumbo frames (7500 to 9000 MTU). For correct cluster operation, verify that jumbo frames are enabled on all upstream virtual and physical switches providing internal network services to ONTAP Select cluster nodes.

RAID SyncMirror traffic (e0e)

Synchronous replication of blocks across HA partner nodes occurs using an internal network interface residing on network port e0e. This functionality occurs automatically, using network interfaces configured by ONTAP during cluster setup, and requires no configuration by the administrator.

> Port e0e is reserved by ONTAP for internal replication traffic. Therefore, neither the port nor the hosted LIF is visible in the ONTAP CLI or in System Manager. This interface is configured to use an automatically generated link local IP address, and you cannot assign an alternate IP address. This network port requires the use of jumbo frames (7500 to 9000 MTU).

HA interconnect (e0f)

NetApp FAS arrays use specialized hardware to pass information between HA pairs in an ONTAP cluster. Software-defined environments, however, do not tend to have this type of equipment available (such as InfiniBand or iWARP devices), so an alternate solution is needed. Although several possibilities were considered, ONTAP requirements placed on the interconnect transport required that this functionality be emulated in software. As a result, within an ONTAP Select cluster, the functionality of the HA interconnect

(traditionally provided by hardware) has been designed into the OS, using Ethernet as a transport mechanism.

Each ONTAP Select node is configured with an HA interconnect port, e0f. This port hosts the HA interconnect network interface, which is responsible for two primary functions:

- Mirroring the contents of NVRAM between HA pairs
- Sending/receiving HA status information and network heartbeat messages between HA pairs

HA interconnect traffic flows through this network port using a single network interface by layering remote direct memory access (RDMA) frames within Ethernet packets.

> Like the RSM port (e0e), neither the physical port nor the hosted network interface is visible to users from either the ONTAP CLI or from System Manager. As a result, you cannot modify the IP address of this interface, and you cannot change the state of the port. This network port requires the use of jumbo frames (7500 to 9000 MTU).

ONTAP Select internal and external networks

Characteristics of ONTAP Select internal and external networks.

ONTAP Select internal network

The internal ONTAP Select network, which is only present in the multi-node variant of the product, is responsible for providing the ONTAP Select cluster with cluster communication, HA interconnect, and synchronous replication services. This network includes the following ports and interfaces:

- **e0c, e0d.** Hosting cluster network LIFs
- **e0e.** Hosting the RSM LIF
- **e0f.** Hosting the HA interconnect LIF

The throughput and latency of this network are critical in determining the performance and resiliency of the ONTAP Select cluster. Network isolation is required for cluster security and to make sure that system interfaces are kept separate from other network traffic. Therefore, this network must be used exclusively by the ONTAP Select cluster.

> Using the Select internal network for traffic other than Select cluster traffic, such as application or management traffic, is not supported. There can be no other VMs or hosts on the ONTAP internal VLAN.

Network packets traversing the internal network must be on a dedicated VLAN-tagged layer-2 network. This can be accomplished by completing one of the following tasks:

- Assigning a VLAN-tagged port group to the internal virtual NICs (e0c through e0f) (VST mode)
- Using the native VLAN provided by the upstream switch where the native VLAN is not used for any other traffic (assign a port group with no VLAN ID, that is, EST mode)

In all cases, VLAN tagging for internal network traffic is done outside of the ONTAP Select VM.

> Only ESXi standard and distributed vSwitches are supported. Other virtual switches or direct connectivity between ESXi hosts are not supported. The internal network must be fully opened; NAT or firewalls are not supported.

Within an ONTAP Select cluster, internal traffic and external traffic are separated using virtual layer-2 network objects known as port groups. Correct vSwitch assignment of these port groups is extremely important, especially for the internal network, which is responsible for providing cluster, HA interconnect, and mirror replication services. Insufficient network bandwidth to these network ports can cause performance degradation and even affect the stability of the cluster node. Therefore, four, six, eight, ten, and twelve-node clusters require that the internal ONTAP Select network use 10Gb connectivity; 1Gb NICs are not supported. Trade-offs can be made to the external network, however, because limiting the flow of incoming data to an ONTAP Select cluster does not affect its ability to operate reliably.

A two-node cluster can use either four 1Gb ports for internal traffic or a single 10Gb port instead of the two 10Gb ports required by the four-node cluster. In an environment in which conditions prevent the server from being fit with four 10Gb NIC cards, two 10Gb NIC cards can be used for the internal network and two 1Gb NICs can be used for the external ONTAP network.

Internal network validation and troubleshooting

The internal network in a multi-node cluster can be validated by using the network connectivity checker functionality. This function can be invoked from the Deploy CLI running the `network connectivity-check start` command.

Run the following command to view the output of the test:

```
network connectivity-check show --run-id X (X is a number)
```

This tool is only useful for troubleshooting the internal network in a multi-node Select cluster. The tool should not be used to troubleshoot single-node clusters (including vNAS configurations), ONTAP Deploy to ONTAP Select connectivity, or client-side connectivity issues.

The cluster create wizard (part of the ONTAP Deploy UI) includes the internal network checker as an optional step available during the creation of multi-node clusters. Given the important role that the internal network plays in multi-node clusters, making this step part of the cluster create workflow improves the success rate of cluster create operations.

Starting with ONTAP Deploy 2.10, the MTU size used by the internal network can be set between 7,500 and 9,000. The network connectivity checker can also be used to test MTU size between 7,500 and 9,000. The default MTU value is set to the value of the virtual network switch. That default would have to be replaced with a smaller value if a network overlay like VXLAN is present in the environment.

ONTAP Select external network

The ONTAP Select external network is responsible for all outbound communications by the cluster and, therefore, is present on both the single-node and multi-node configurations. Although this network does not have the tightly defined throughput requirements of the internal network, the administrator should be careful not to create network bottlenecks between the client and ONTAP VM, because performance issues could be mischaracterized as ONTAP Select problems.



In a manner similar to internal traffic, external traffic can be tagged at the vSwitch layer (VST) and at the external switch layer (EST). In addition, the external traffic can be tagged by the ONTAP Select VM itself in a process known as VGT. Refer to the section Data and management traffic separation for further details.

The following table highlights the major differences between the ONTAP Select internal and external networks.

Internal versus external network quick reference

Description	Internal Network	External Network
Network services	Cluster HA/IC RAID SyncMirror (RSM)	Data management Intercluster (SnapMirror and SnapVault)
Network isolation	Required	Optional
Frame size (MTU)	7,500 to 9,000	1,500 (default) 9,000 (supported)
IP address assignment	Autogenerated	User-defined
DHCP support	No	No

NIC teaming

To make sure that the internal and external networks have both the necessary bandwidth and resiliency characteristics required to provide high performance and fault tolerance, physical network adapter teaming is recommended. Two-node cluster configurations with a single 10Gb link are supported. However, the NetApp recommended best practice is to make use of NIC teaming on both the internal and the external networks of the ONTAP Select cluster.

MAC address generation

The MAC addresses assigned to all ONTAP Select network ports are generated automatically by the included deployment utility. The utility uses a platform-specific, organizationally unique identifier (OUI) specific to NetApp to make sure there is no conflict with FAS systems. A copy of this address is then stored in an internal database within the ONTAP Select installation VM (ONTAP Deploy), to prevent accidental reassignment during future node deployments. At no point should the administrator modify the assigned MAC address of a network port.

Supported ONTAP Select network configurations

Select the best hardware and configure your network to optimize performance and resiliency.

Server vendors understand that customers have different needs and choice is critical. As a result, when purchasing a physical server, there are numerous options available when making network connectivity decisions. Most commodity systems ship with various NIC choices that provide single-port and multiport options with varying permutations of speed and throughput. This includes support for 25Gb/s and 40Gb/s NIC adapters with VMware ESX.

Because the performance of the ONTAP Select VM is tied directly to the characteristics of the underlying hardware, increasing the throughput to the VM by selecting higher-speed NICs results in a higher-performing cluster and a better overall user experience. Four 10Gb NICs or two higher-speed NICs (25/40 Gb/s) can be used to achieve a high performance network layout. There are a number of other configurations that are also supported. For two-node clusters, 4 x 1Gb ports or 1 x 10Gb ports are supported. For single node clusters, 2 x 1Gb ports are supported.

Network minimum and recommended configurations

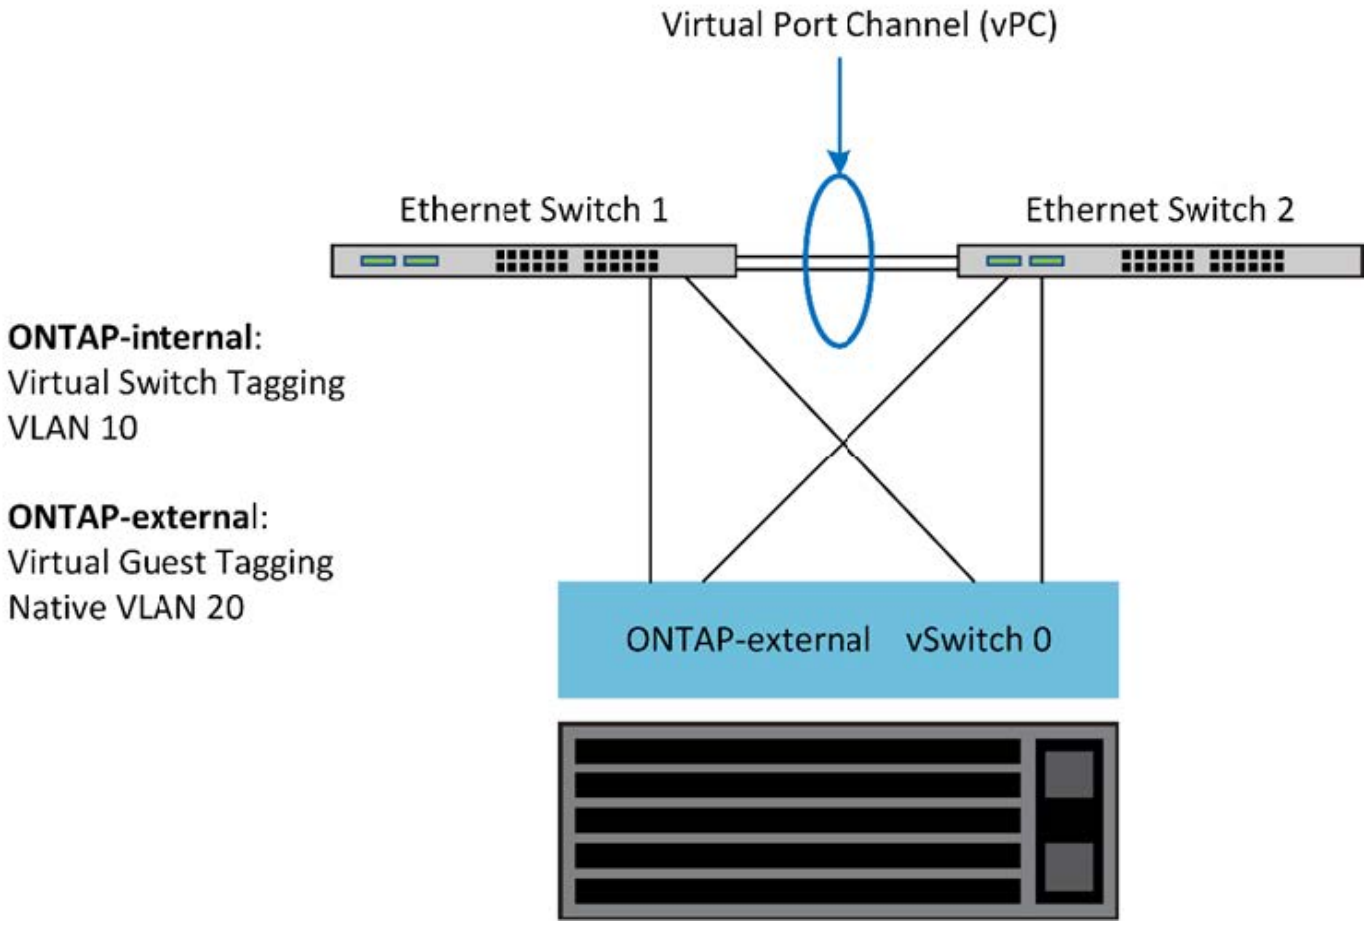
There are several supported Ethernet configurations based on the cluster size.

Cluster size	Minimum Requirements	Recommendation
Single-node cluster	2 x 1GbE	2 x 10GbE
Two-node cluster or MetroCluster SDS	4 x 1GbE or 1 x 10GbE	2 x 10GbE
Four, six, eight, ten, or twelve-node cluster	2 x 10GbE	4 x 10GbE or 2 x 25/40GbE

Conversion between single link and multiple link topologies on a running cluster is not supported because of the possible need to convert between different NIC teaming configurations required for each topology.

Network configuration using multiple physical switches

When sufficient hardware is available, NetApp recommends using the multiswitch configuration shown in the following figure, due to the added protection against physical switch failures.



ONTAP Select VMware vSphere vSwitch configuration on ESXi

ONTAP Select vSwitch configuration and load-balancing policies for two-NIC and four-NIC configurations.

ONTAP Select supports the use of both standard and distributed vSwitch configurations. Distributed vSwitches support link aggregation constructs (LACP). Link aggregation is a common network construct used to

aggregate bandwidth across multiple physical adapters. LACP is a vendor-neutral standard. It provides an open protocol for network endpoints that bundle groups of physical network ports into a single logical channel. ONTAP Select can work with port groups that are configured as a Link Aggregation Group (LAG). However, NetApp recommends using the individual physical ports as simple uplink (trunk) ports to avoid the LAG configuration. In these cases, the best practices for standard and distributed vSwitches are identical.

This section describes the vSwitch configuration and load-balancing policies that should be used in both two-NIC and four-NIC configurations.

When you configure the port groups for ONTAP Select, follow these best practices; the load-balancing policy at the port-group level is Route Based on Originating Virtual Port ID. VMware recommends that STP be set to Portfast on the switch ports connected to the ESXi hosts.

All vSwitch configurations require a minimum of two physical network adapters bundled into a single NIC team. ONTAP Select supports a single 10Gb link for two-node clusters. However, NetApp recommends using NIC aggregation to ensure hardware redundancy.

On a vSphere server, NIC teams are the aggregation construct used to bundle multiple physical network adapters into a single logical channel, allowing the network load to be shared across all member ports. It's important to remember that NIC teams can be created without support from the physical switch. Load-balancing and failover policies can be applied directly to a NIC team, which is unaware of the upstream switch configuration. In this case, policies are only applied to outbound traffic.



Static port channels are not supported with ONTAP Select. LACP-enabled channels are supported with distributed vSwitches but using LACP LAGs might result in un-even load distribution across the LAG members.

For single node clusters, ONTAP Deploy configures the ONTAP Select VM to use a port group for the external network and either the same port group or, optionally, a different port group for the cluster and node management traffic. For single node clusters, you can add the desired number of physical ports to the external port group as active adapters.

For multi-node clusters, ONTAP Deploy configures each ONTAP Select VM to use one or two port groups for the internal network and separately, one or two port groups for the external network. The cluster and node management traffic can either use the same port group as the external traffic, or optionally a separate port group. The cluster and node management traffic cannot share the same port group with internal traffic.

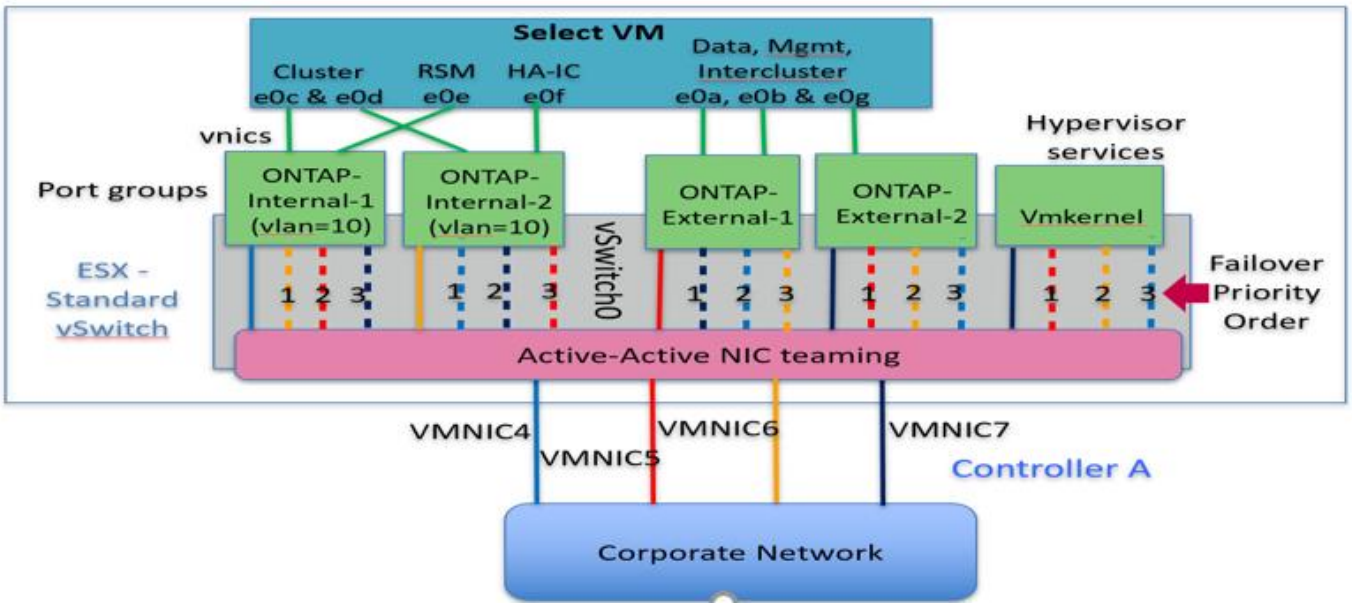


ONTAP Select supports a maximum of four VMNICs.

Standard or distributed vSwitch and four physical ports per Node

You can assign four port groups to each node in a multi-node cluster. Each port group has a single active physical port and three standby physical ports as in the following figure.

vSwitch with four physical ports per node



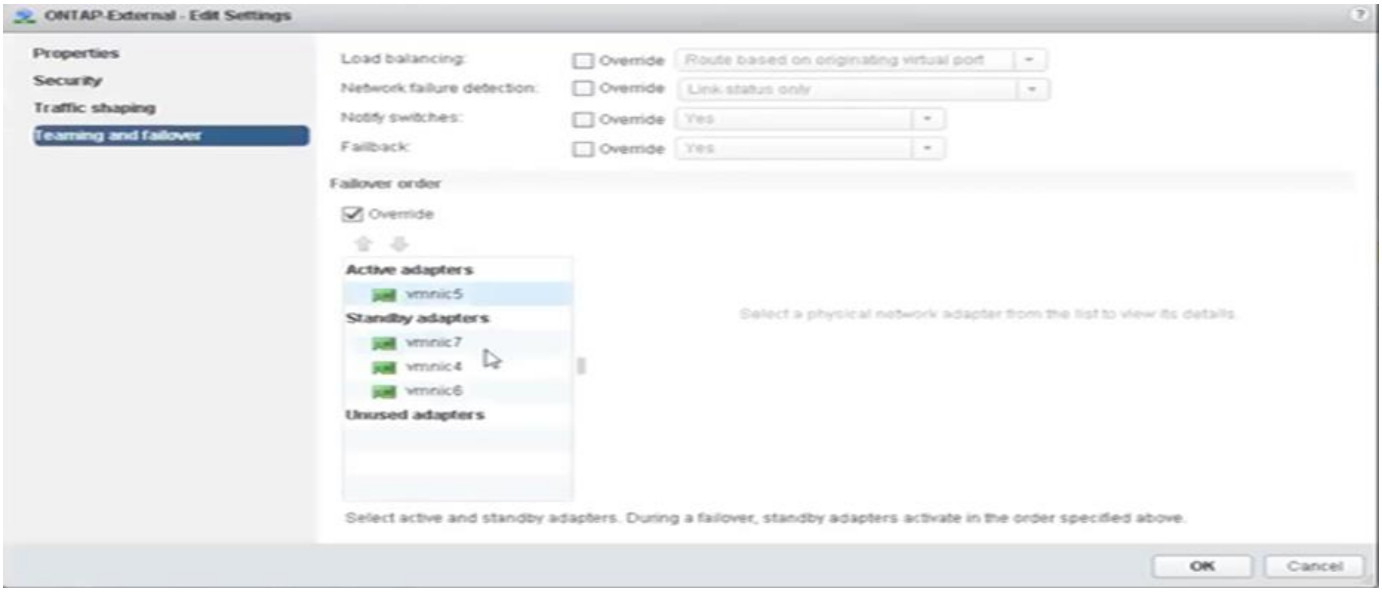
The order of the ports in the standby list is important. The following table provides an example of the physical port distribution across the four port groups.

Network minimum and recommended configurations

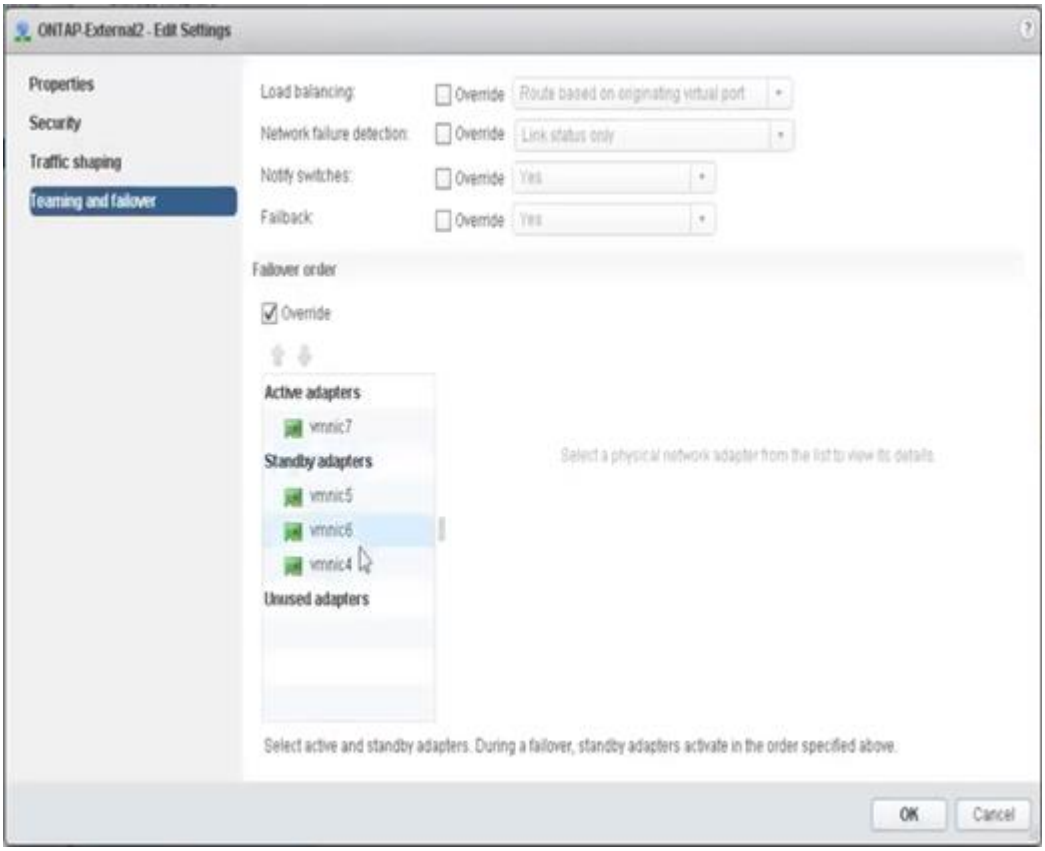
Port Group	External 1	External 2	Internal 1	Internal 2
Active	vmnic0	vmnic1	vmnic2	vmnic3
Standby 1	vmnic1	vmnic0	vmnic3	vmnic2
Standby 2	vmnic2	vmnic3	vmnic0	vmnic1
Standby 3	vmnic3	vmnic2	vmnic1	vmnic0

The following figures show the configurations of the external network port groups from the vCenter UI (ONTAP-External and ONTAP-External2). Note that the active adapters are from different network cards. In this setup, vmnic 4 and vmnic 5 are dual ports on the same physical NIC, while vmnic 6 and vminc 7 are similarly dual ports on a separate NIC (vnmics 0 through 3 are not used in this example). The order of the standby adapters provides a hierarchical fail over with the ports from the internal network being last. The order of internal ports in the standby list is similarly swapped between the two external port groups.

Part 1: ONTAP Select external port group configurations



Part 2: ONTAP Select external port group configurations



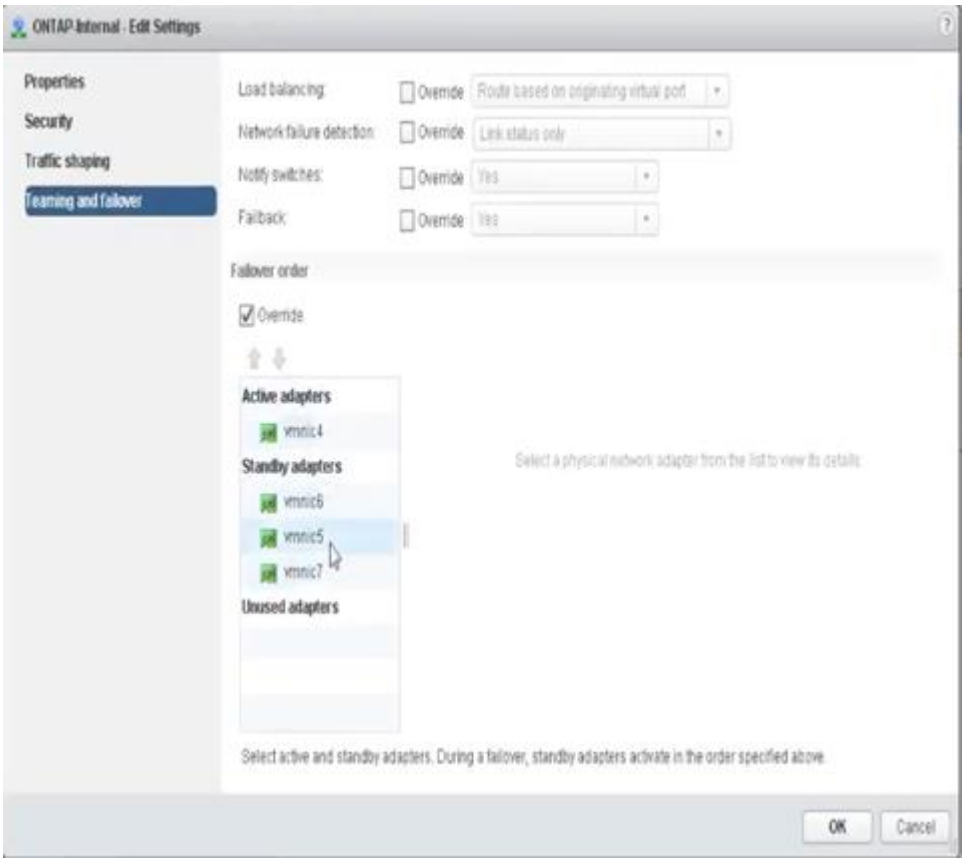
For readability, the assignments are as follows:

ONTAP-External	ONTAP-External2
Active adapters: vmnic5 Standby adapters: vmnic7, vmnic4, vmnic6	Active adapters: vmnic7 Standby adapters: vmnic5, vmnic6, vmnic4

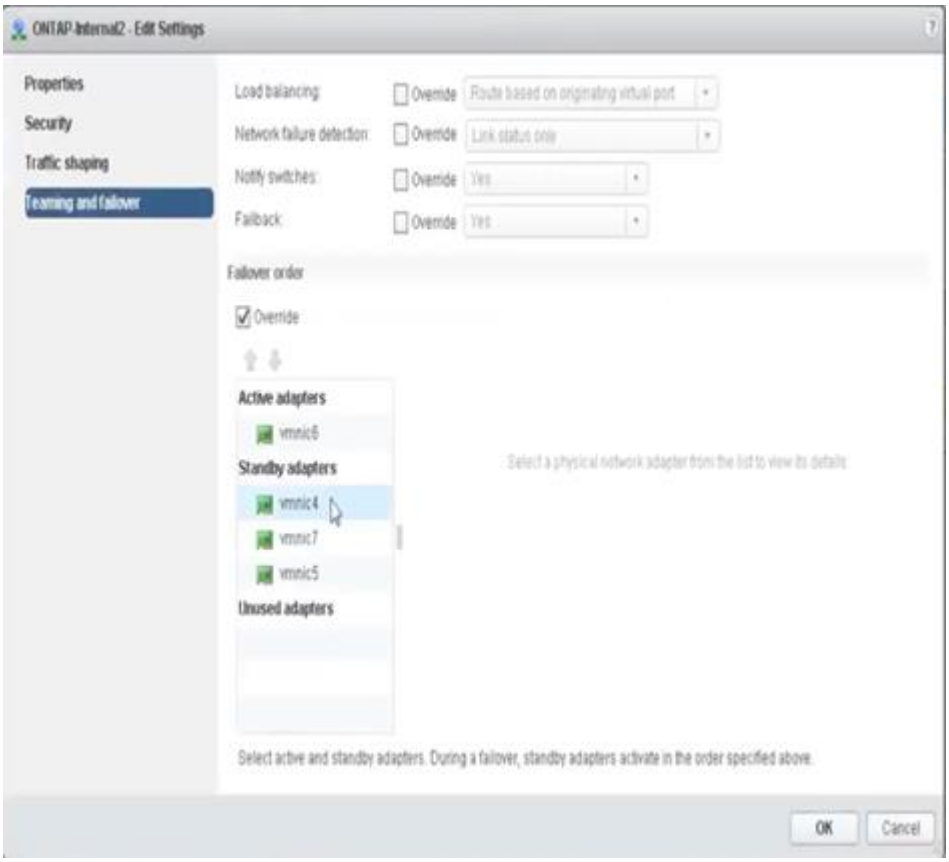
The following figures show the configurations of the internal network port groups (ONTAP-Internal and ONTAP-Internal2). Note that the active adapters are from different network cards. In this setup, vmnic 4 and vmnic 5

are dual ports on the same physical ASIC, whereas vmnic 6 and vmnic 7 are similarly dual ports on a separate ASIC. The order of the standby adapters provides a hierarchical fail over with the ports from the external network being last. The order of external ports in the standby list is similarly swapped between the two internal port groups.

Part 1: ONTAP Select internal port group configurations



Part 2: ONTAP Select internal port groups



For readability, the assignments are as follows:

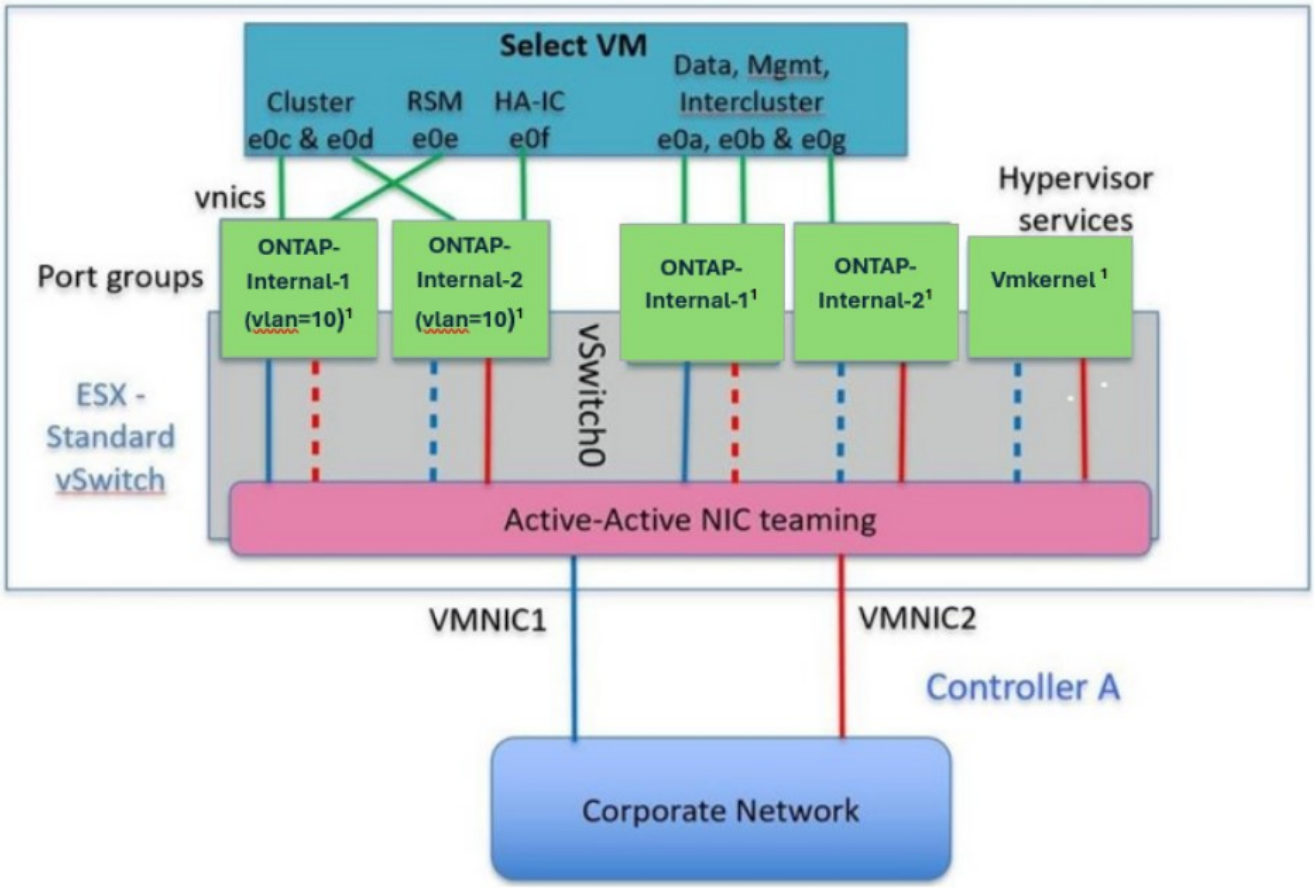
ONTAP-Internal	ONTAP-Internal2
Active adapters: vmnic4 Standby adapters: vmnic6, vmnic5, vmnic7	Active adapters: vmnic6 Standby adapters: vmnic4, vmnic7, vmnic5

Standard or distributed vSwitch and two physical ports per node

When using two high speed (25/40Gb) NICs, the recommended port group configuration is conceptually very similar to the configuration with four 10Gb adapters. You should use four port groups even when using only two physical adapters. The port group assignments are as follows:

Port Group	External 1 (e0a,e0b)	Internal 1 (e0c,e0e)	Internal 2 (e0d,e0f)	External 2 (e0g)
Active	vmnic0	vmnic0	vmnic1	vmnic1
Standby	vmnic1	vmnic1	vmnic0	vmnic0

vSwitch with two high speed (25/40Gb) physical ports per node

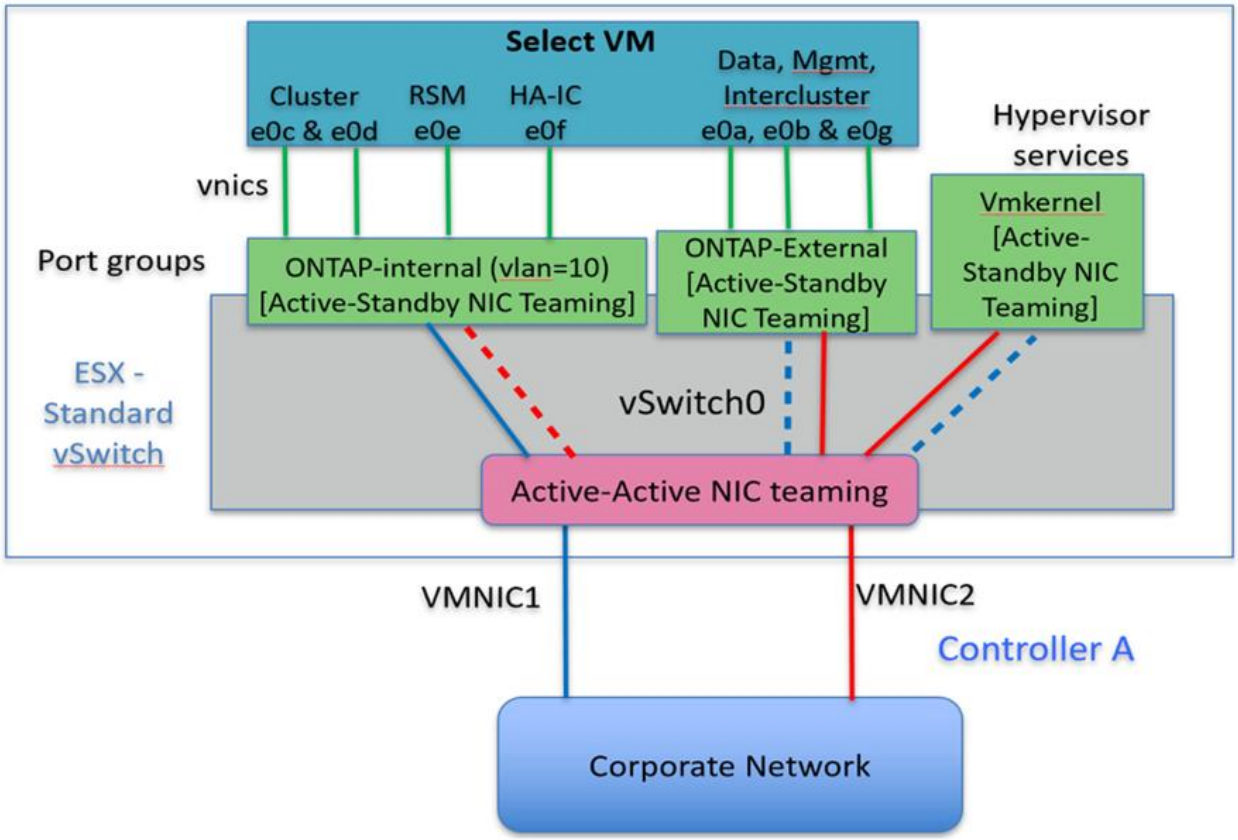


[1] The port groups attached to the virtual NICs are configured to use one NIC as active and the rest as standby.

When using two physical ports (10Gb or less), each port group should have an active adapter and a standby adapter configured opposite to each other. The internal network is only present for multi-node ONTAP Select clusters. For single-node clusters, both adapters can be configured as active in the external port group.

The following example shows the configuration of a vSwitch and the two port groups responsible for handling internal and external communication services for a multi-node ONTAP Select cluster. The external network can use the internal network VMNIC in the event of a network outage because the internal network VMNICs are part of this port group and configured in standby mode. The opposite is the case for the external network. Alternating the active and standby VMNICs between the two port groups is critical for the correct failover of the ONTAP Select VMs during network outages.

vSwitch with two physical ports (10Gb or less) per node

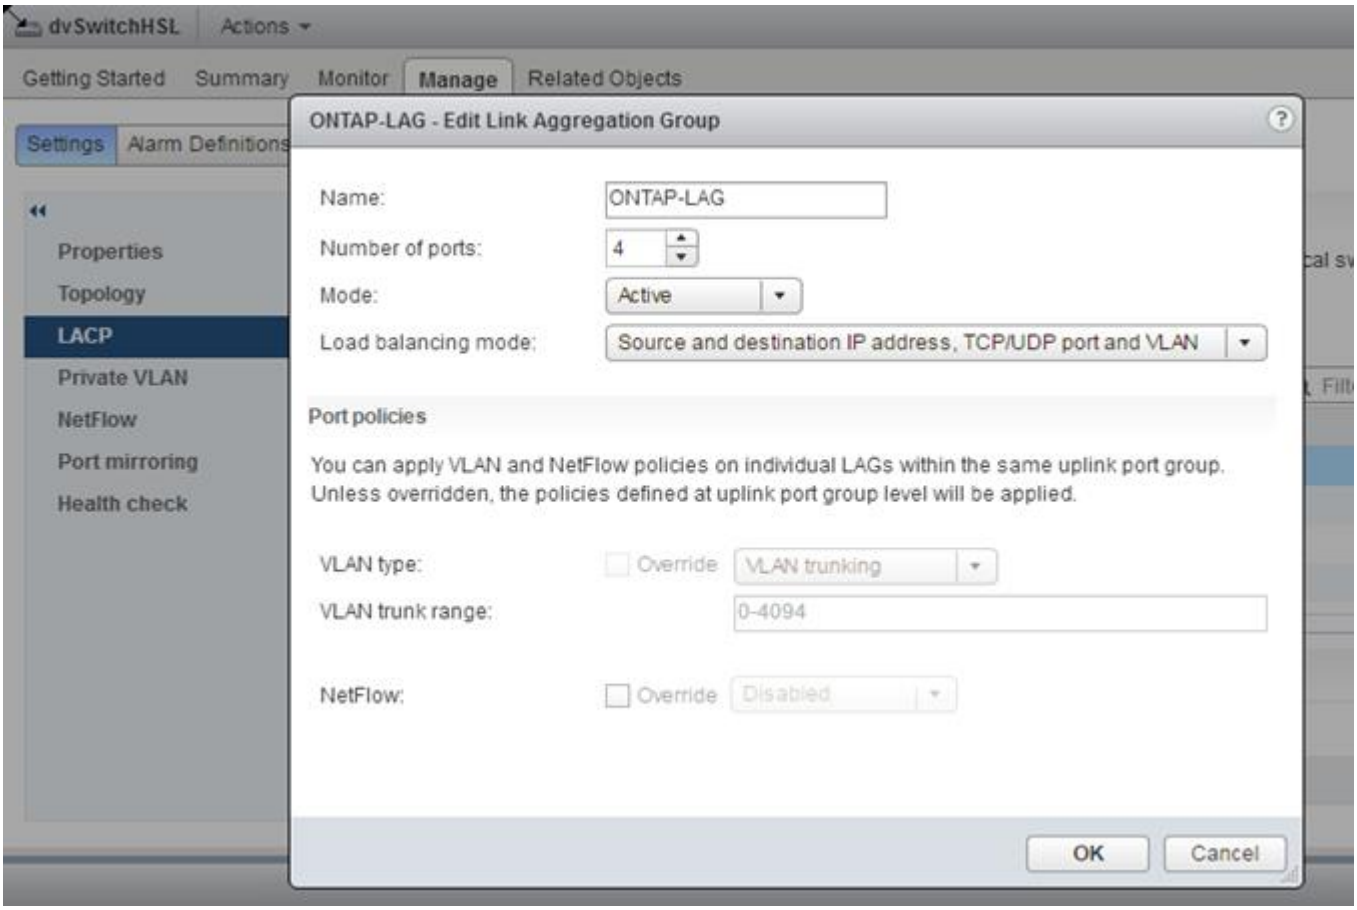


Distributed vSwitch with LACP

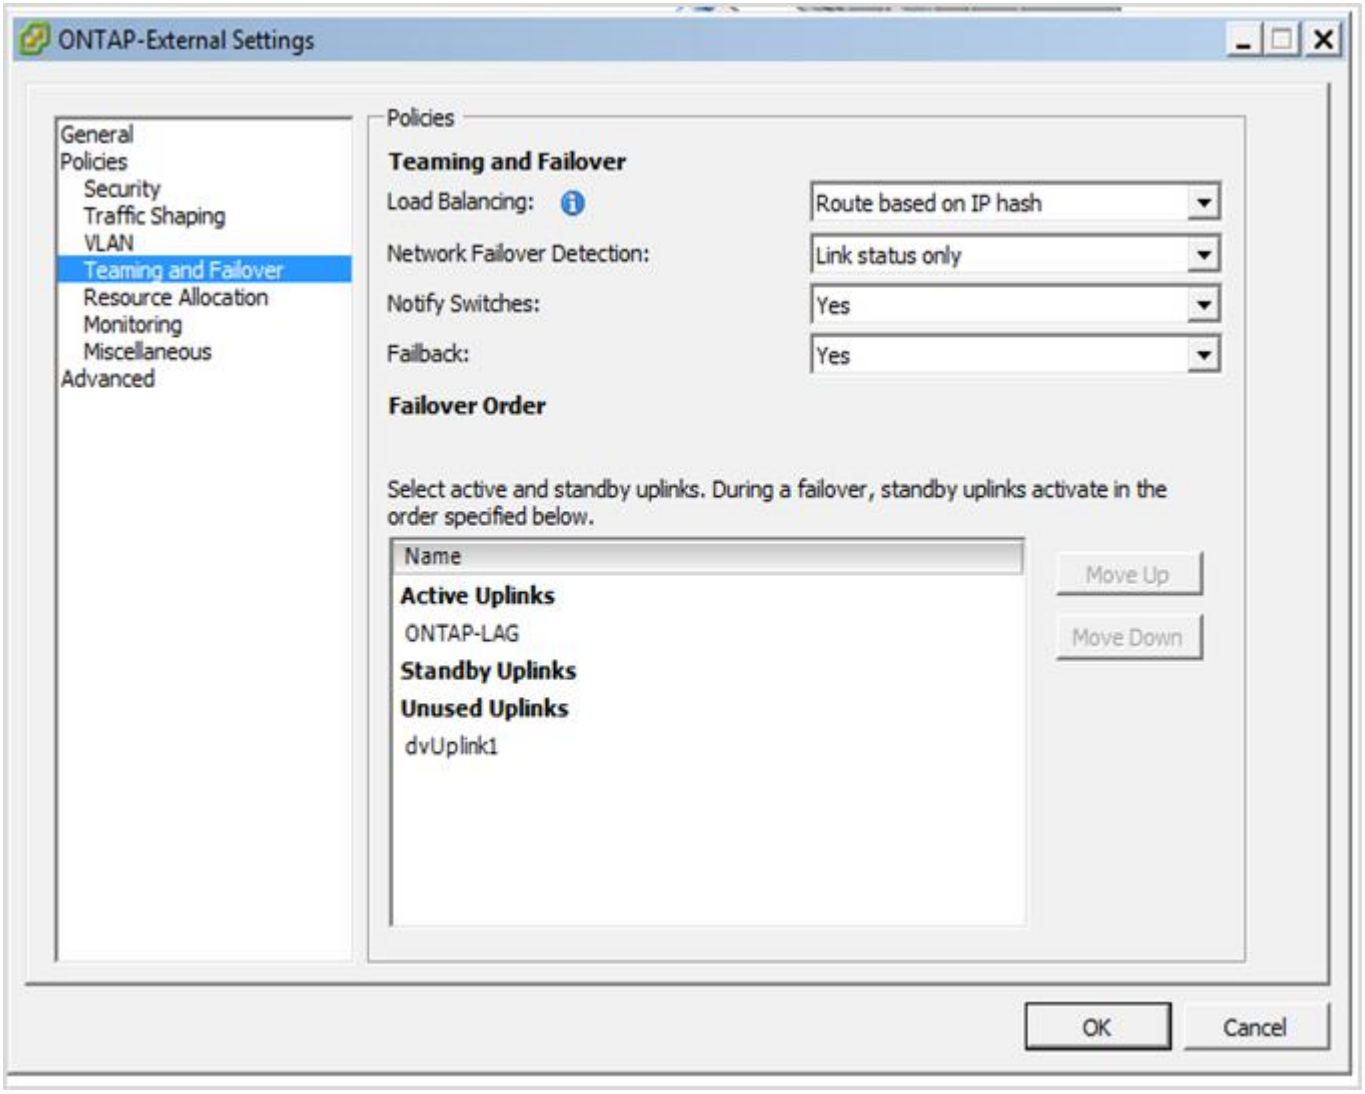
When using distributed vSwitches in your configuration, LACP can be used (though it is not a best practice) in order to simplify the network configuration. The only supported LACP configuration requires that all the VMNICs are in a single LAG. The uplink physical switch must support an MTU size between 7,500 to 9,000 on all the ports in the channel. The internal and external ONTAP Select networks should be isolated at the port group level. The internal network should use a nonroutable (isolated) VLAN. The external network can use either VST, EST, or VGT.

The following examples show the distributed vSwitch configuration using LACP.

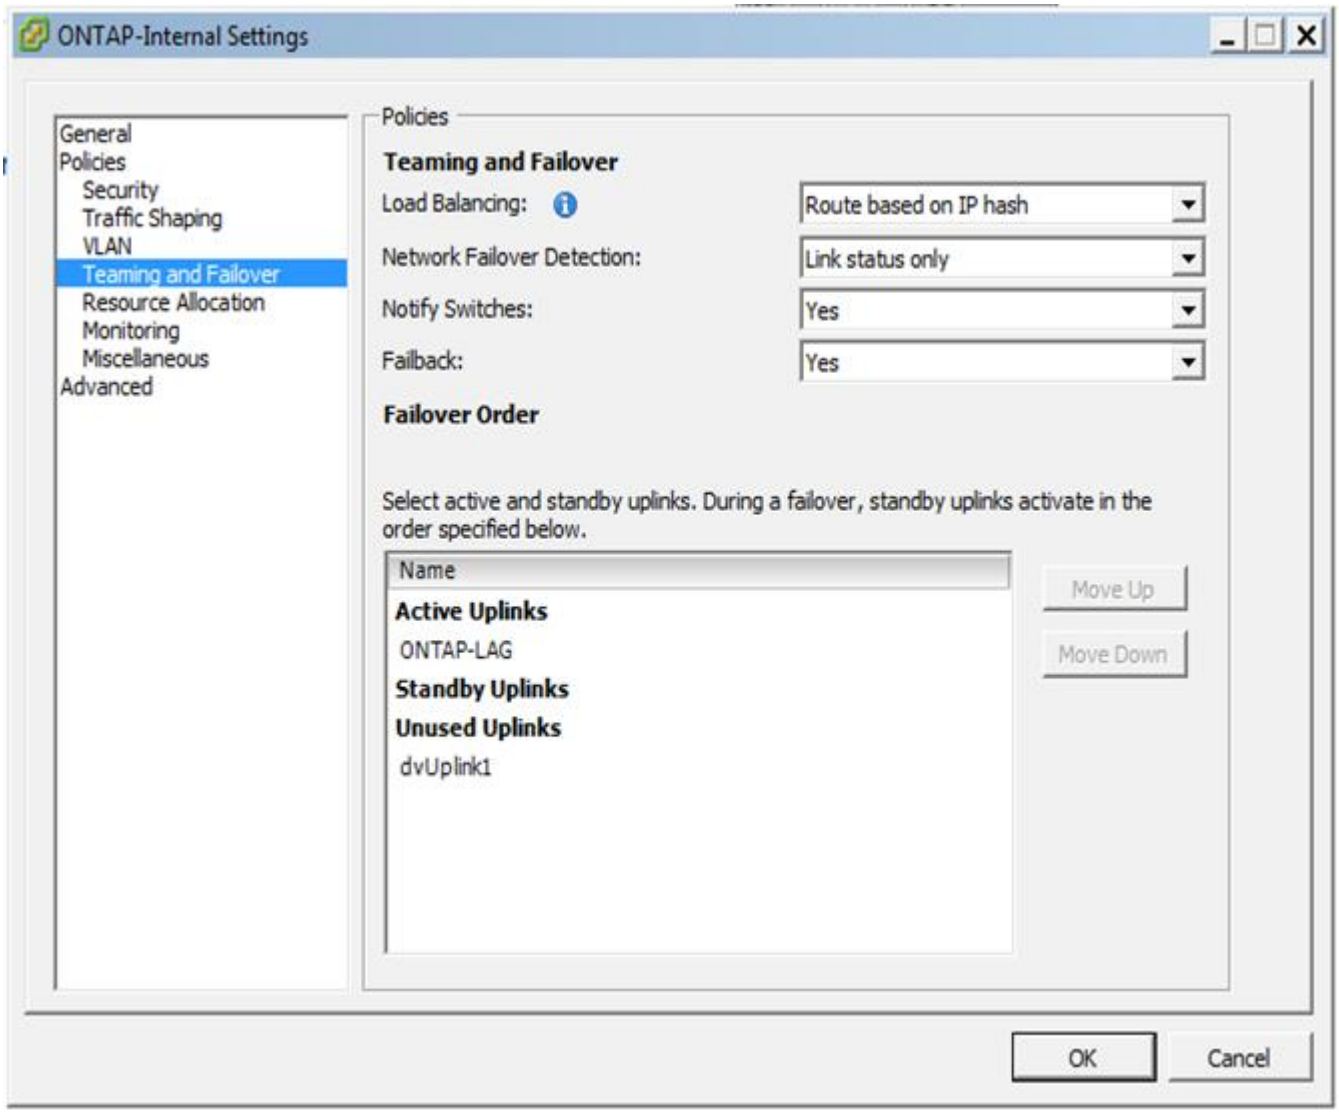
LAG properties when using LACP



External port group configurations using a distributed vSwitch with LACP enabled



Internal port group configurations using a distributed vSwitch with LACP enabled



LACP requires you to configure the upstream switch ports as a port channel. Before you enable this configuration on the distributed vSwitch, ensure that an LACP-enabled port channel is correctly configured.

ONTAP Select physical switch configuration

Upstream physical switch configuration details based on single-switch and multi-switch environments.

Careful consideration should be taken when making connectivity decisions from the virtual switch layer to physical switches. Separation of internal cluster traffic from external data services should extend to the upstream physical networking layer through isolation provided by layer-2 VLANs.

Physical switch ports should be configured as trunkports. ONTAP Select external traffic can be separated across multiple layer-2 networks in one of two ways. One method is by using ONTAP VLAN-tagged virtual ports with a single port group. The other method is by assigning separate port groups in VST mode to management port e0a. You must also assign data ports to e0b and e0c/e0g depending on the ONTAP Select release and the single-node or multi-node configuration. If the external traffic is separated across multiple layer-2 networks, the uplink physical switch ports should have those VLANs in its allowed VLAN list.

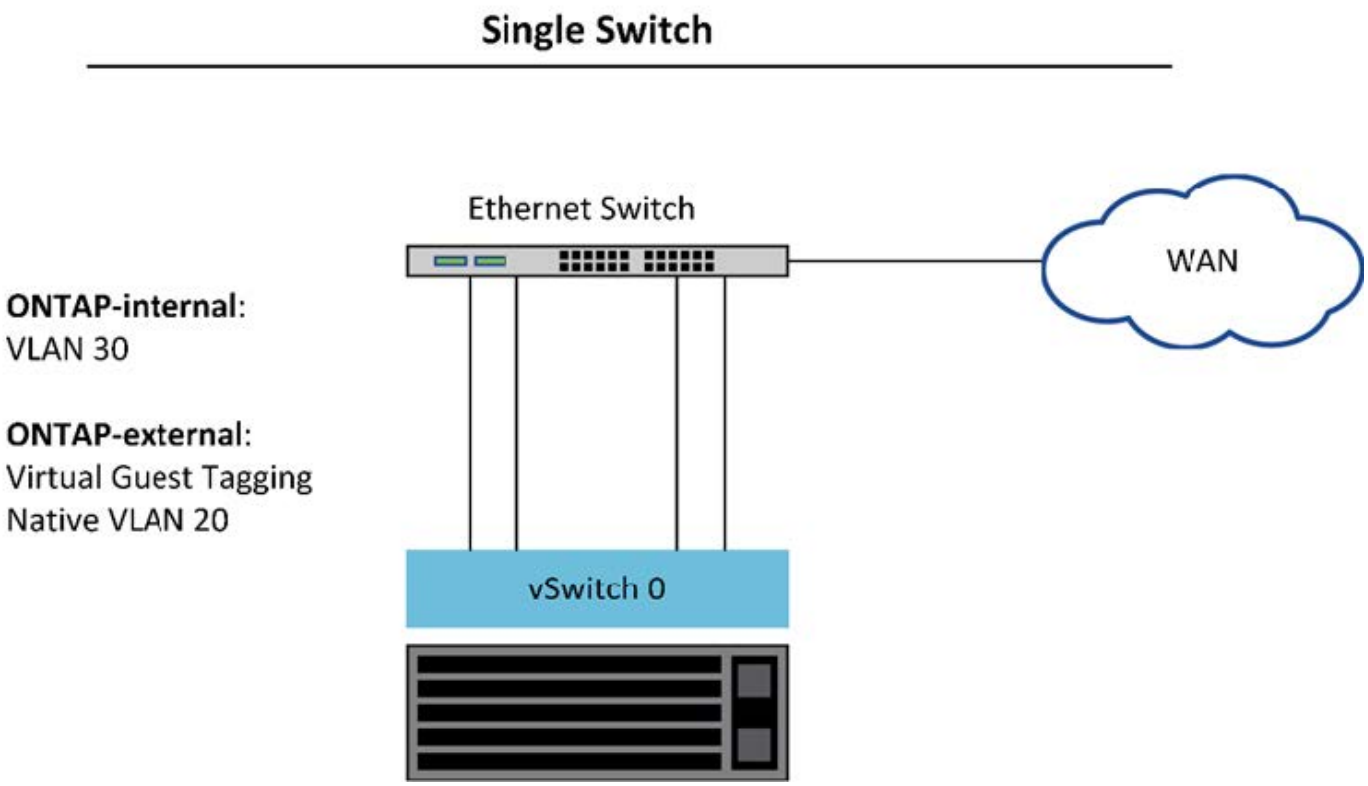
ONTAP Select internal network traffic occurs using virtual interfaces defined with link local IP addresses. Because these IP addresses are non-routable, internal traffic between cluster nodes must flow across a single layer-2 network. Route hops between ONTAP Select cluster nodes are unsupported.

Shared physical switch

The following figure depicts a possible switch configuration used by one node in a multi-node ONTAP Select cluster. In this example, the physical NICs used by the vSwitches hosting both the internal and external network port groups are cabled to the same upstream switch. Switch traffic is kept isolated using broadcast domains contained within separate VLANs.

For the ONTAP Select internal network, tagging is done at the port group level. While the following example uses VGT for the external network, both VGT and VST are supported on that port group.

Network configuration using shared physical switch

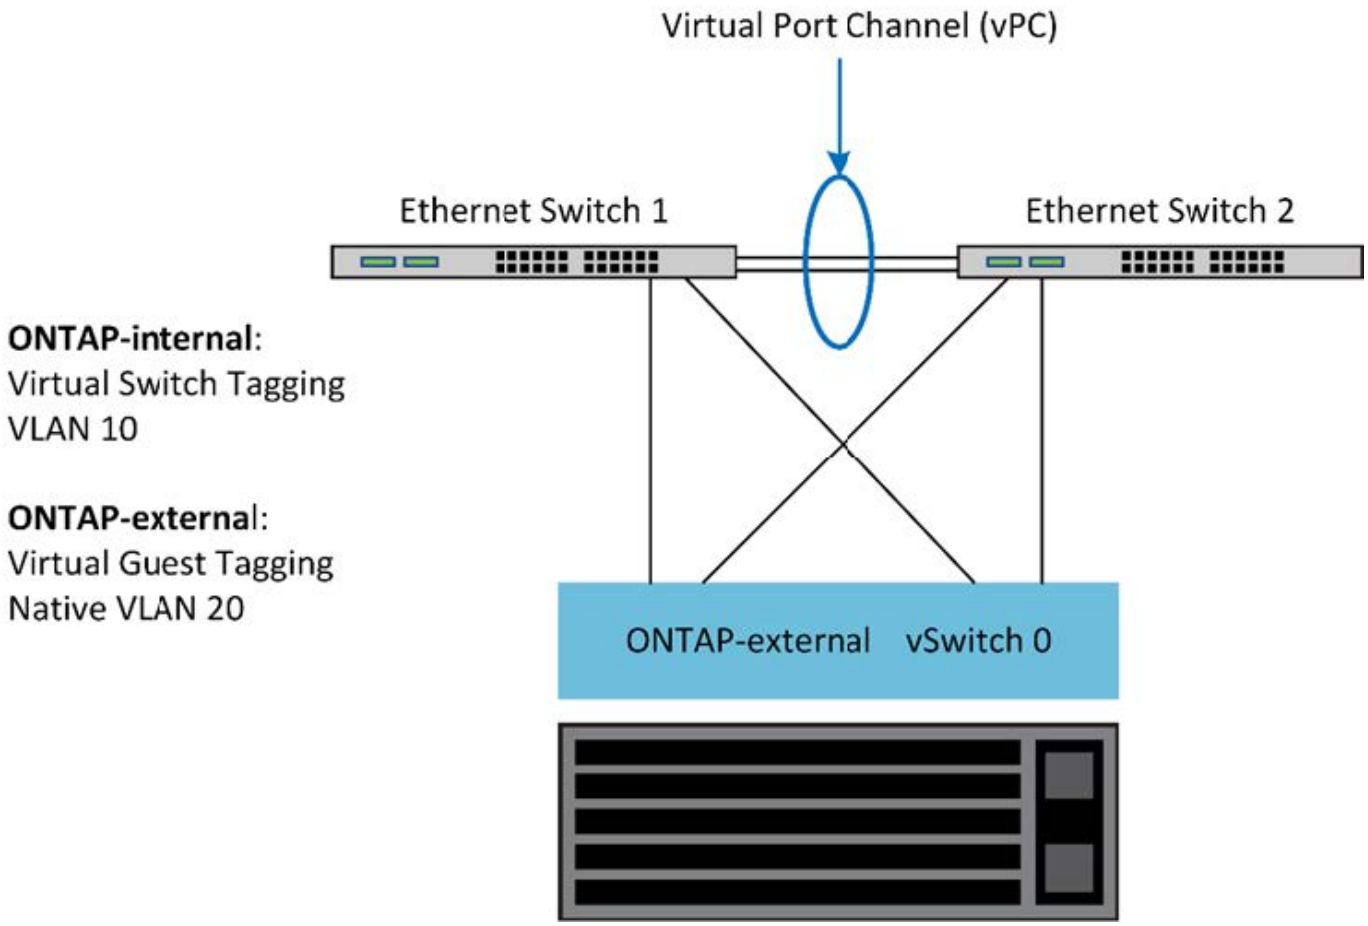


In this configuration, the shared switch becomes a single point of failure. If possible, multiple switches should be used to prevent a physical hardware failure from causing a cluster network outage.

Multiple physical switches

When redundancy is needed, multiple physical network switches should be used. The following figure shows a recommended configuration used by one node in a multi-node ONTAP Select cluster. NICs from both the internal and external port groups are cabled into different physical switches, protecting the user from a single hardware-switch failure. A virtual port channel is configured between switches to prevent spanning tree issues.

Network configuration using multiple physical switches



ONTAP Select data and management traffic separation

Isolate data traffic and management traffic into separate layer-2 networks.

ONTAP Select external network traffic is defined as data (CIFS, NFS, and iSCSI), management, and replication (SnapMirror) traffic. Within an ONTAP cluster, each style of traffic uses a separate logical interface that must be hosted on a virtual network port. On the multi-node configuration of ONTAP Select, these are designated as ports e0a and e0b/e0g. On the single node configuration, these are designated as e0a and e0b/e0c, while the remaining ports are reserved for internal cluster services.

NetApp recommends isolating data traffic and management traffic into separate layer-2 networks. In the ONTAP Select environment, this is done using VLAN tags. This can be achieved by assigning a VLAN-tagged port group to network adapter 1 (port e0a) for management traffic. Then you can assign a separate port group(s) to ports e0b and e0c (single-node clusters) and e0b and e0g (multi-node clusters) for data traffic.

If the VST solution described earlier in this document is not sufficient, collocating both data and management LIFs on the same virtual port might be required. To do so, use a process known as VGT, in which VLAN tagging is performed by the VM.

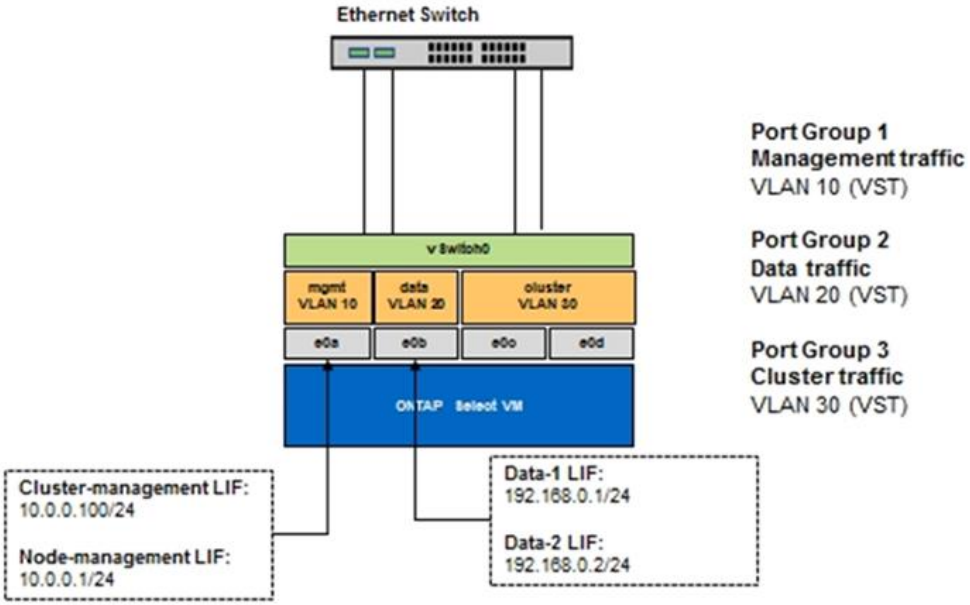
Data and management network separation through VGT is not available when using the ONTAP Deploy utility. This process must be performed after cluster setup is complete.

There is an additional caveat when using VGT and two-node clusters. In two-node cluster configurations, the node management IP address is used to establish connectivity to the mediator before ONTAP is fully available. Therefore, only EST and VST tagging is supported on the port group mapped to the node management LIF

(port e0a). Furthermore, if both the management and the data traffic are using the same port group, only EST/VST are supported for the entire two-node cluster.

Both configuration options, VST and VGT, are supported. The following figure shows the first scenario, VST, in which traffic is tagged at the vSwitch layer through the assigned port group. In this configuration, cluster and node management LIFs are assigned to ONTAP port e0a and tagged with VLAN ID 10 through the assigned port group. Data LIFs are assigned to port e0b and either e0c or e0g and given VLAN ID 20 using a second port group. The cluster ports use a third port group and are on VLAN ID 30.

Data and management separation using VST

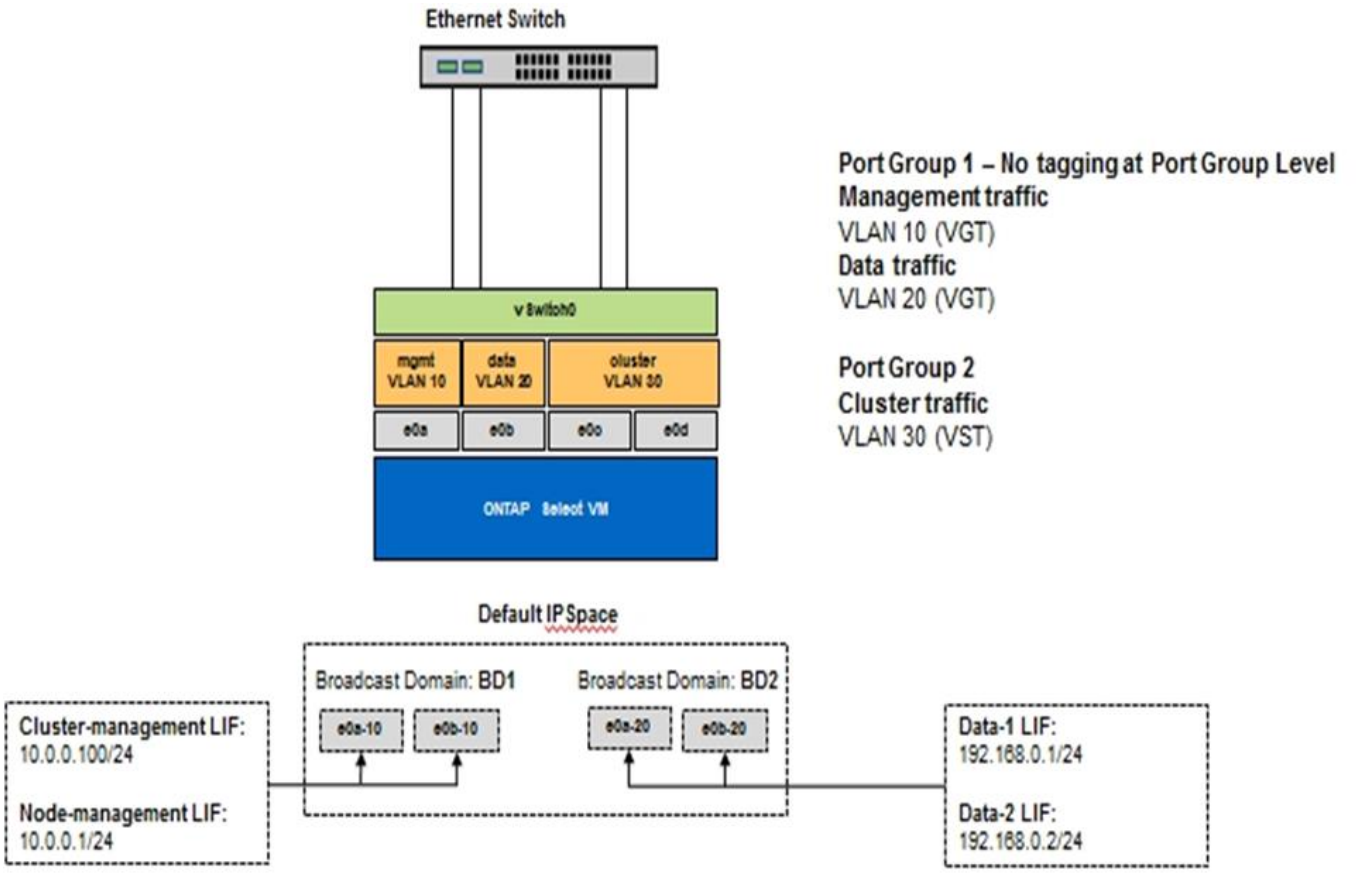


The following figure shows the second scenario, VGT, in which traffic is tagged by the ONTAP VM using VLAN ports that are placed into separate broadcast domains. In this example, virtual ports e0a-10/e0b-10/(e0c or e0g)-10 and e0a-20/e0b-20 are placed on top of VM ports e0a and e0b. This configuration allows network tagging to be performed directly within ONTAP, rather than at the vSwitch layer. Management and data LIFs are placed on these virtual ports, allowing further layer-2 subdivision within a single VM port. The cluster VLAN (VLAN ID 30) is still tagged at the port group.

Notes:

- This style of configuration is especially desirable when using multiple IPspaces. Group VLAN ports into separate custom IPspaces if further logical isolation and multitenancy are desired.
- To support VGT, the ESXi/ESX host network adapters must be connected to trunk ports on the physical switch. The port groups connected to the virtual switch must have their VLAN ID set to 4095 to enable trunking on the port group.

Data and management separation using VGT



High availability architecture

ONTAP Select high availability configurations

Discover high availability options to select the best HA configuration for your environment.

Although customers are starting to move application workloads from enterprise-class storage appliances to software-based solutions running on commodity hardware, the expectations and needs around resiliency and fault tolerance have not changed. An HA solution providing a zero recovery point objective (RPO) protects the customer from data loss due to a failure from any component in the infrastructure stack.

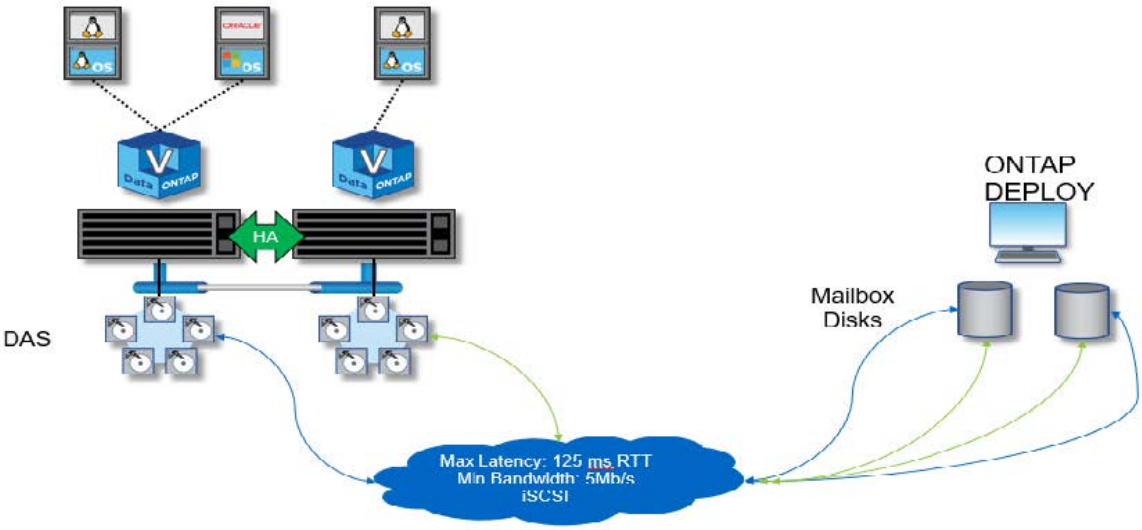
A large portion of the SDS market is built on the notion of shared-nothing storage, with software replication providing data resiliency by storing multiple copies of user data across different storage silos. ONTAP Select builds on this premise by using the synchronous replication features (RAID SyncMirror) provided by ONTAP to store an extra copy of user data within the cluster. This occurs within the context of an HA pair. Every HA pair stores two copies of user data: one on storage provided by the local node, and one on storage provided by the HA partner. Within an ONTAP Select cluster, HA and synchronous replication are tied together, and the functionality of the two cannot be decoupled or used independently. As a result, the synchronous replication functionality is only available in the multi-node offering.

In an ONTAP Select cluster, synchronous replication functionality is a function of HA implementation, not a replacement for the asynchronous SnapMirror or SnapVault replication engines. Synchronous replication cannot be used independently from HA.

There are two ONTAP Select HA deployment models: the multi-node clusters (four, six, eight, ten, or twelve nodes) and the two-node clusters. The salient feature of a two-node ONTAP Select cluster is the use of an external mediator service to resolve split-brain scenarios. The ONTAP Deploy VM serves as the default mediator for all of the two-node HA pairs that it configures.

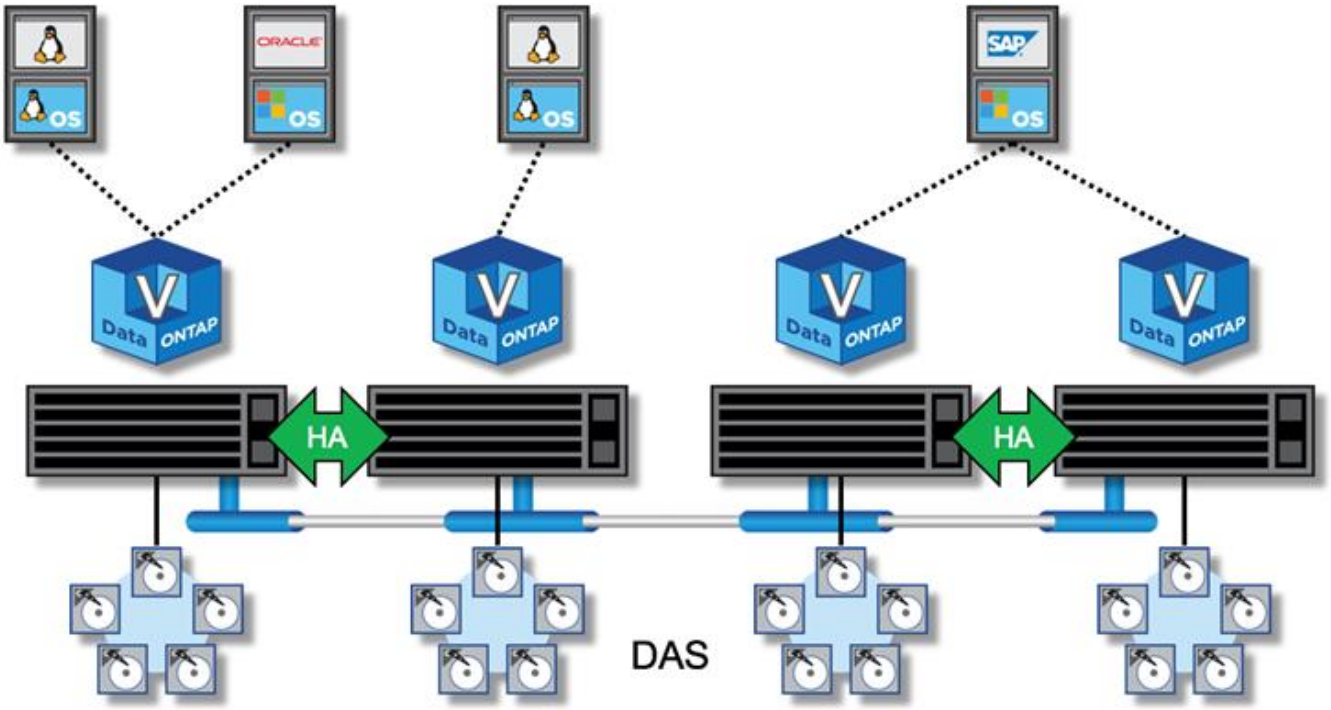
The two architectures are represented in the following figures.

Two-node ONTAP Select cluster with remote mediator and using local-attached storage



The two-node ONTAP Select cluster is composed of one HA pair and a mediator. Within the HA pair, data aggregates on each cluster node are synchronously mirrored, and, in the event of a failover, there is no loss of data.

Four-node ONTAP Select cluster using local-attached storage



- The four-node ONTAP Select cluster is composed of two HA pairs. Six-node, eight-node, ten-node, and twelve-node clusters are composed of three, four, five, and six HA pairs, respectively. Within each HA pair, data aggregates on each cluster node are synchronously mirrored, and, in the event of a failover, there is no loss of data.
- Only one ONTAP Select instance can be present on a physical server when using DAS storage. ONTAP Select requires unshared access to the local RAID controller of the system and is designed to manage the locally attached disks, which would be impossible without physical connectivity to the storage.

Two-node HA versus multi-node HA

Unlike FAS arrays, ONTAP Select nodes in an HA pair communicate exclusively over the IP network. That means that the IP network is a single point of failure (SPOF), and protecting against network partitions and split-brain scenarios becomes an important aspect of the design. The multi-node cluster can sustain single-node failures because the cluster quorum can be established by the three or more surviving nodes. The two-node cluster relies on the mediator service hosted by the ONTAP Deploy VM to achieve the same result.

The heartbeat network traffic between the ONTAP Select nodes and the ONTAP Deploy mediator service is minimal and resilient so that the ONTAP Deploy VM can be hosted in a different data center than the ONTAP Select two-node cluster.

> The ONTAP Deploy VM becomes an integral part of a two-node cluster when serving as the mediator for that cluster. If the mediator service is not available, the two-node cluster continues serving data, but the storage failover capabilities of the ONTAP Select cluster are disabled. Therefore, the ONTAP Deploy mediator service must maintain constant communication with each ONTAP Select node in the HA pair. A minimum bandwidth of 5Mbps and a maximum round-trip time (RTT) latency of 125ms are required to allow proper functioning of the cluster quorum.

If the ONTAP Deploy VM acting as a mediator is temporarily or potentially permanently unavailable, a secondary ONTAP Deploy VM can be used to restore the two-node cluster quorum. This results in a

configuration in which the new ONTAP Deploy VM is unable to manage the ONTAP Select nodes, but it successfully participates in the cluster quorum algorithm. The communication between the ONTAP Select nodes and the ONTAP Deploy VM is done by using the iSCSI protocol over IPv4. The ONTAP Select node management IP address is the initiator, and the ONTAP Deploy VM IP address is the target. Therefore, it is not possible to support IPv6 addresses for the node management IP addresses when creating a two-node cluster. The ONTAP Deploy hosted mailbox disks are automatically created and masked to the proper ONTAP Select node management IP addresses at the time of two-node cluster creation. The entire configuration is automatically performed during setup, and no further administrative action is required. The ONTAP Deploy instance creating the cluster is the default mediator for that cluster.

An administrative action is required if the original mediator location must be changed. It is possible to recover a cluster quorum even if the original ONTAP Deploy VM is lost. However, NetApp recommends that you back up the ONTAP Deploy database after every two-node cluster is instantiated.

Two-node HA versus two-node stretched HA (MetroCluster SDS)

It is possible to stretch a two-node, active/active HA cluster across larger distances and potentially place each node in a different data center. The only distinction between a two-node cluster and a two-node stretched cluster (also referred to as MetroCluster SDS) is the network connectivity distance between nodes.

The two-node cluster is defined as a cluster for which both nodes are located in the same data center within a distance of 300m. In general, both nodes have uplinks to the same network switch or set of interswitch link (ISL) network switches.

Two-node MetroCluster SDS is defined as a cluster with nodes that are physically separated (different rooms, different buildings, and different data centers) by more than 300m. In addition, each node's uplink connections are connected to separate network switches. The MetroCluster SDS does not require dedicated hardware. However, the environment should adhere to requirements for latency (a maximum of 5ms for RTT and 5ms for jitter, for a total of 10ms).

MetroCluster SDS is a premium feature and requires a Premium license or a Premium XL license. The Premium license supports the creation of both small and medium VMs, as well as HDD and SSD media. The Premium XL license also supports the creation of NVMe drives.



MetroCluster SDS is supported with both local attached storage (DAS) and shared storage (vNAS). Note that vNAS configurations usually have a higher innate latency because of the network between the ONTAP Select VM and shared storage. MetroCluster SDS configurations must provide a maximum of 10ms of latency between the nodes, including the shared storage latency. In other words, only measuring the latency between the Select VMs is not adequate because shared storage latency is not negligible for these configurations.

ONTAP Select HA RSM and mirrored aggregates

Prevent data loss using RAID SyncMirror (RSM), mirrored aggregates, and the write path.

Synchronous replication

The ONTAP HA model is built on the concept of HA partners. ONTAP Select extends this architecture into the nonshared commodity server world by using the RAID SyncMirror (RSM) functionality that is present in ONTAP to replicate data blocks between cluster nodes, providing two copies of user data spread across an HA pair.

A two-node cluster with a mediator can span two data centers. For more information, refer to the section Two-node stretched HA (MetroCluster SDS) best practices.

Mirrored aggregates

An ONTAP Select cluster is composed of two to twelve nodes. Each HA pair contains two copies of user data, synchronously mirrored across nodes over an IP network. This mirroring is transparent to the user, and it is a property of the data aggregate, which is automatically configured during the data aggregate creation process.

All aggregates in an ONTAP Select cluster must be mirrored for data availability in the event of a node failover and to avoid an SPOF in case of hardware failure. Aggregates in an ONTAP Select cluster are built from virtual disks provided from each node in the HA pair and use the following disks:

- A local set of disks (contributed by the current ONTAP Select node)
- A mirrored set of disks (contributed by the HA partner of the current node)

> The local and mirror disks used to build a mirrored aggregate must be the same size. These aggregates are referred to as plex 0 and plex 1 (to indicate the local and remote mirror pairs, respectively). The actual plex numbers can be different in your installation.

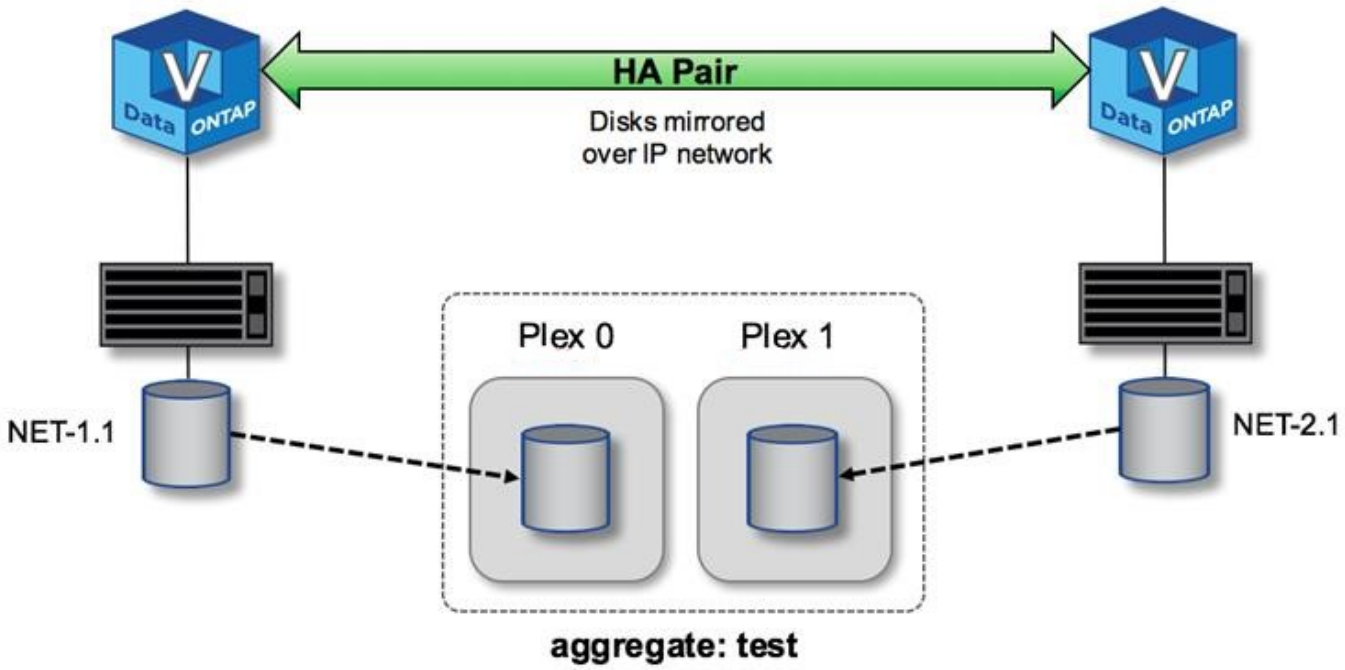
This approach is fundamentally different from the way standard ONTAP clusters work. This applies to all root and data disks within the ONTAP Select cluster. The aggregate contains both local and mirror copies of data. Therefore, an aggregate that contains N virtual disks offers N/2 disks' worth of unique storage, because the second copy of data resides on its own unique disks.

The following figure shows an HA pair within a four-node ONTAP Select cluster. Within this cluster is a single aggregate (test) that uses storage from both HA partners. This data aggregate is composed of two sets of virtual disks: a local set, contributed by the ONTAP Select owning cluster node (Plex 0), and a remote set, contributed by the failover partner (Plex 1).

Plex 0 is the bucket that holds all local disks. Plex 1 is the bucket that holds mirror disks, or disks responsible for storing a second replicated copy of user data. The node that owns the aggregate contributes disks to Plex 0, and the HA partner of that node contributes disks to Plex 1.

In the following figure, there is a mirrored aggregate with two disks. The contents of this aggregate are mirrored across our two cluster nodes, with local disk NET-1.1 placed into the Plex 0 bucket and remote disk NET-2.1 placed into the Plex 1 bucket. In this example, aggregate test is owned by the cluster node to the left and uses local disk NET-1.1 and HA partner mirror disk NET-2.1.

ONTAP Select mirrored aggregate



When an ONTAP Select cluster is deployed, all virtual disks present on the system are automatically assigned to the correct plex, requiring no additional step from the user regarding disk assignment. This prevents the accidental assignment of disks to an incorrect plex and provides optimal mirror disk configuration.

Write Path

Synchronous mirroring of data blocks between cluster nodes and the requirement for no data loss with a system failure have a significant impact on the path an incoming write takes as it propagates through an ONTAP Select cluster. This process consists of two stages:

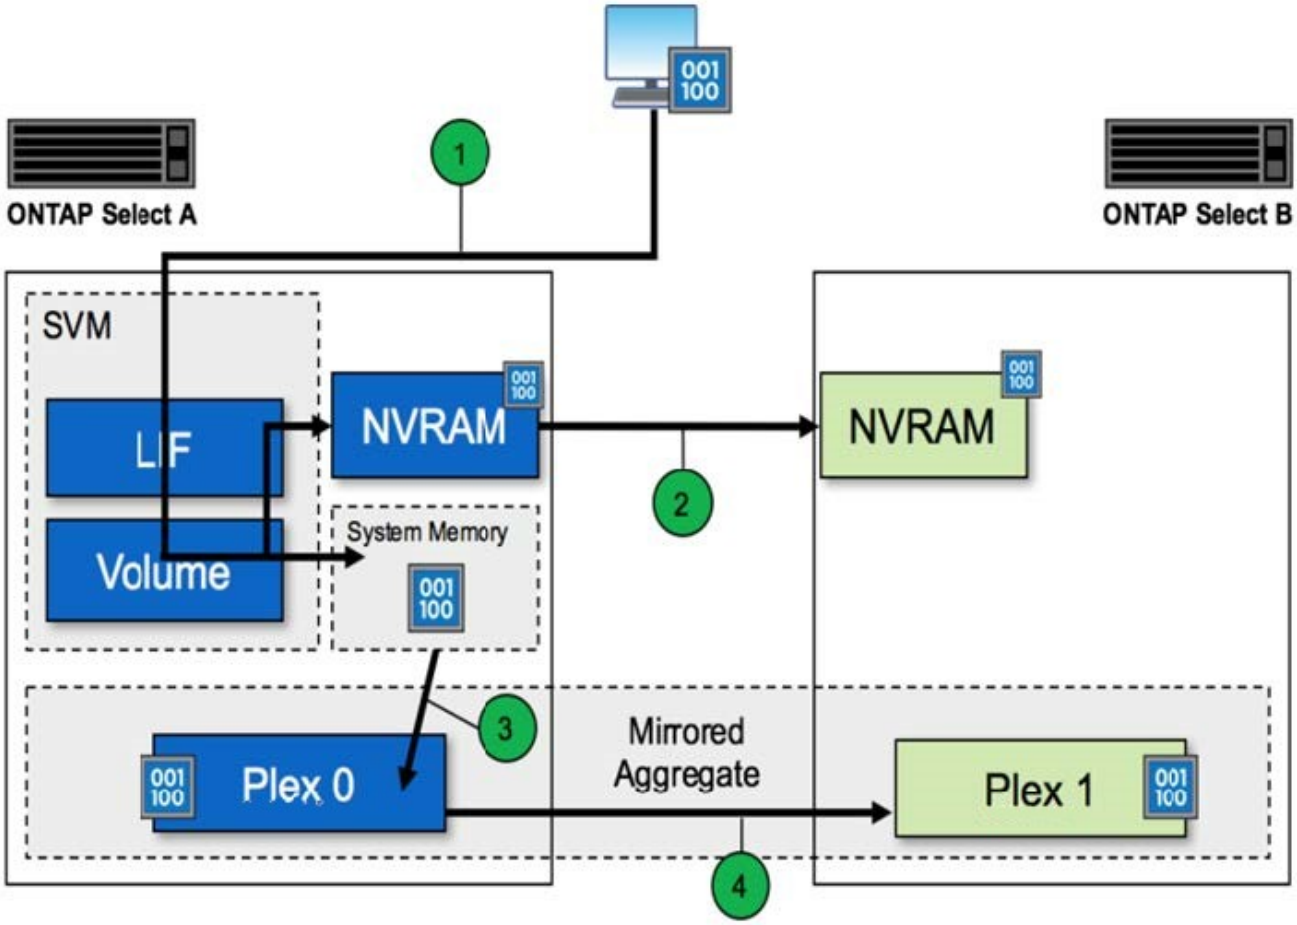
- Acknowledgment
- Destaging

Writes to a target volume occur over a data LIF and are committed to the virtualized NVRAM partition, present on a system disk of the ONTAP Select node, before being acknowledged back to the client. On an HA configuration, an additional step occurs, because these NVRAM writes are immediately mirrored to the HA partner of the target volume's owner before being acknowledged. This process makes sure of the file system consistency on the HA partner node, if there is a hardware failure on the original node.

After the write has been committed to NVRAM, ONTAP periodically moves the contents of this partition to the appropriate virtual disk, a process known as destaging. This process only happens once, on the cluster node owning the target volume, and does not happen on the HA partner.

The following figure shows the write path of an incoming write request to an ONTAP Select node.

ONTAP Select write path workflow



Incoming write acknowledgment includes the following steps:

- Writes enter the system through a logical interface owned by ONTAP Select node A.
- Writes are committed to the NVRAM of node A and mirrored to the HA partner, node B.
- After the I/O request is present on both HA nodes, the request is then acknowledged back to the client.

ONTAP Select destaging from NVRAM to the data aggregate (ONTAP CP) includes the following steps:

- Writes are destaged from virtual NVRAM to virtual data aggregate.
- Mirror engine synchronously replicates blocks to both plexes.

Data protection enhancements with ONTAP Select HA

High availability (HA) disk heartbeating, HA mailbox, HA heartbeating, HA Failover, and Giveback work to enhance data protection.

Disk heartbeating

Although the ONTAP Select HA architecture leverages many of the code paths used by the traditional FAS arrays, some exceptions exist. One of these exceptions is in the implementation of disk-based heartbeating, a nonnetwork-based method of communication used by cluster nodes to prevent network isolation from causing split-brain behavior. A split-brain scenario is the result of cluster partitioning, typically caused by network failures, whereby each side believes the other is down and attempts to take over cluster resources.

Enterprise-class HA implementations must gracefully handle this type of scenario. ONTAP does this through a customized, disk-based method of heartbeating. This is the job of the HA mailbox, a location on physical storage that is used by cluster nodes to pass heartbeat messages. This helps the cluster determine connectivity and therefore define quorum in the event of a failover.

On FAS arrays, which use a shared storage HA architecture, ONTAP resolves split-brain issues in the following ways:

- SCSI persistent reservations
- Persistent HA metadata
- HA state sent over HA interconnect

However, within the shared-nothing architecture of an ONTAP Select cluster, a node is only able to see its own local storage and not that of the HA partner. Therefore, when network partitioning isolates each side of an HA pair, the preceding methods of determining cluster quorum and failover behavior are unavailable.

Although the existing method of split-brain detection and avoidance cannot be used, a method of mediation is still required, one that fits within the constraints of a shared-nothing environment. ONTAP Select extends the existing mailbox infrastructure further, allowing it to act as a method of mediation in the event of network partitioning. Because shared storage is unavailable, mediation is accomplished through access to the mailbox disks over NAS. These disks are spread throughout the cluster, including the mediator in a two-node cluster, using the iSCSI protocol. Therefore, intelligent failover decisions can be made by a cluster node based on access to these disks. If a node can access the mailbox disks of other nodes outside of its HA partner, it is likely up and healthy.



The mailbox architecture and disk-based heartbeating method of resolving cluster quorum and split-brain issues are the reasons the multi-node variant of ONTAP Select requires either four separate nodes or a mediator for a two-node cluster.

HA mailbox posting

The HA mailbox architecture uses a message post model. At repeated intervals, cluster nodes post messages to all other mailbox disks across the cluster, including the mediator, stating that the node is up and running. Within a healthy cluster at any point in time, a single mailbox disk on a cluster node has messages posted from all other cluster nodes.

Attached to each Select cluster node is a virtual disk that is used specifically for shared mailbox access. This disk is referred to as the mediator mailbox disk, because its main function is to act as a method of cluster mediation in the event of node failures or network partitioning. This mailbox disk contains partitions for each cluster node and is mounted over an iSCSI network by other Select cluster nodes. Periodically, these nodes post health statuses to the appropriate partition of the mailbox disk. Using network-accessible mailbox disks spread throughout the cluster allows you to infer node health through a reachability matrix. For example, cluster nodes A and B can post to the mailbox of cluster node D, but not to the mailbox of node C. In addition, cluster node D cannot post to the mailbox of node C, so it is likely that node C is either down or network isolated and should be taken over.

HA heartbeating

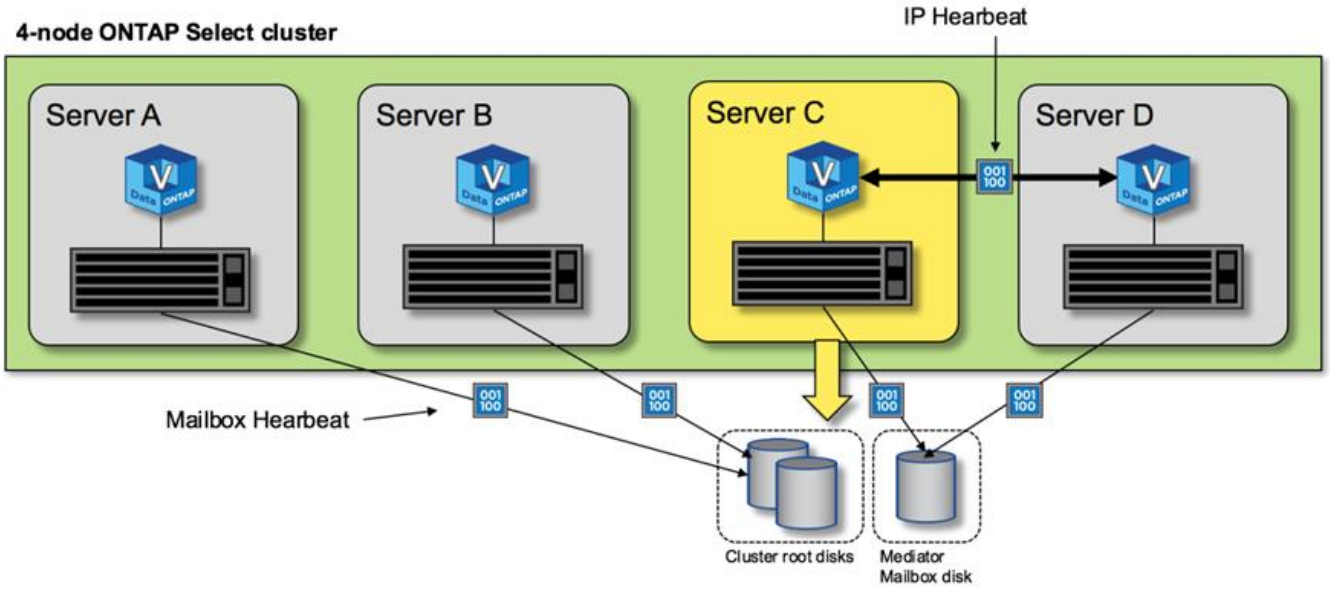
Like with NetApp FAS platforms, ONTAP Select periodically sends HA heartbeat messages over the HA interconnect. Within the ONTAP Select cluster, this is performed over a TCP/IP network connection that exists between HA partners. Additionally, disk-based heartbeat messages are passed to all HA mailbox disks, including mediator mailbox disks. These messages are passed every few seconds and read back periodically. The frequency with which these are sent and received allows the ONTAP Select cluster to detect HA failure

events within approximately 15 seconds, the same window available on FAS platforms. When heartbeat messages are no longer being read, a failover event is triggered.

The following figure shows the process of sending and receiving heartbeat messages over the HA interconnect and mediator disks from the perspective of a single ONTAP Select cluster node, node C.

Network heartbeats are sent over the HA interconnect to the HA partner, node D, while disk heartbeats use mailbox disks across all cluster nodes, A, B, C, and D.

HA heartbeating in a four-node cluster: steady state



HA failover and giveback

During a failover operation, the surviving node assumes the data serving responsibilities for its peer node using the local copy of its HA partner's data. Client I/O can continue uninterrupted, but changes to this data must be replicated back before giveback can occur. Note that ONTAP Select does not support a forced giveback because this causes changes stored on the surviving node to be lost.

The sync back operation is automatically triggered when the rebooted node rejoins the cluster. The time required for the sync back depends on several factors. These factors include the number of changes that must be replicated, the network latency between the nodes, and the speed of the disk subsystems on each node. It is possible that the time required for sync back will exceed the auto give back window of 10 minutes. In this case, a manual giveback after the sync back is required. The progress of the sync back can be monitored using the following command:

```
storage aggregate status -r -aggregate <aggregate name>
```

Learn about ONTAP Select performance

The performance of an ONTAP Select cluster can vary considerably due to the characteristics of the underlying hardware and configuration. The specific hardware configuration is the biggest factor in the performance of a particular ONTAP Select

instance. Here are some of the factors that affect the performance of a specific ONTAP Select instance:

- **Core frequency**. In general, a higher frequency is preferable.
- **Single socket versus multisocket**. ONTAP Select does not use multisocket features, but the hypervisor overhead for supporting multisocket configurations accounts for some amount of deviation in total performance.
- **RAID card configuration and associated hypervisor driver**. The default driver provided by the hypervisor might need to be replaced by the hardware vendor driver.
- **Drive type and number of drives in the RAID group(s)**.
- **Hypervisor version and patch level**.

Automate with REST

Concepts

Learn about REST web services for deploying and managing ONTAP Select clusters

Representational State Transfer (REST) is a style for creating distributed web applications. When applied to the design of a web services API, it establishes a set of technologies and best practices for exposing server-based resources and managing their states. It uses mainstream protocols and standards to provide a flexible foundation for deploying and managing ONTAP Select clusters.

Architecture and classic constraints

REST was formally articulated by Roy Fielding in his PhD dissertation at UC Irvine in 2000. It defines an architectural style through a set of constraints, which collectively have improves web-based applications and the underlying protocols. The constraints establish a RESTful web services application based on a client/server architecture using a stateless communication protocol.

Resources and state representation

Resources are the basic components of a web-based system. When creating a REST web services application, early design tasks include:

- Identification of system or server-based resources
  Every system uses and maintains resources. A resource can be a file, business transaction, process, or administrative entity. One of the first tasks in designing an application based on REST web services is to identify the resources.
- Definition of resource states and associated state operations
  Resources are always in one of a finite number of states. The states, as well as the associated operations used to affect the state changes, must be clearly defined.

Messages are exchanged between the client and server to access and change the state of the resources according to the generic CRUD (Create, Read, Update, and Delete) model.

URI endpoints

Every REST resource must be defined and made available using a well-defined addressing scheme. The endpoints where the resources are located and identified use a Uniform Resource Identifier (URI). The URI provides a general framework for creating a unique name for each resource in the network. The Uniform Resource Locator (URL) is a type of URI used with web services to identify and access resources. Resources are typically exposed in a hierarchical structure similar to a file directory.

HTTP messages

Hypertext Transfer Protocol (HTTP) is the protocol used by the web services client and server to exchange request and response messages about the resources. As part of designing a web services application, HTTP verbs (such as GET and POST) are mapped to the resources and corresponding state management actions.

HTTP is stateless. Therefore, to associate a set of related requests and responses under one transaction,

additional information must be included in the HTTP headers carried with the request/response data flows.

JSON formatting

While information can be structured and transferred between a client and server in several ways, the most popular option (and the one used with the Deploy REST API) is JavaScript Object Notation (JSON). JSON is an industry standard for representing simple data structures in plain text and is used to transfer state information describing the resources.

Learn about accessing the ONTAP Select Deploy API

Because of the inherent flexibility of REST web services, the ONTAP Select Deploy API can be accessed in several different ways.

> The REST API included with ONTAP Select Deploy is assigned a version number. The API version number is independent of the Deploy release number. The ONTAP Select 9.19.1 Deploy administration utility includes version 3 of the REST API.

Deploy utility native user interface

The primary way you access the API is through the ONTAP Select Deploy web user interface. The browser makes calls to the API and reformats the data according to the design of the user interface. You also access the API through the Deploy utility command line interface.

ONTAP Select Deploy online documentation page

The ONTAP Select Deploy online documentation page provides an alternative access point when using a browser. In addition to providing a way to execute individual API calls directly, the page also includes a detailed description of the API, including input parameters and other options for each call. The API calls are organized into several different functional areas or categories.

Custom program

You can access the Deploy API using any of several different programming languages and tools. Popular choices include Python, Java, and cURL. A program, script, or tool that uses the API acts as a REST web services client. Using a programming language allows you to better understand the API and provides an opportunity to automate the ONTAP Select deployments.

ONTAP Select Deploy API basic operational characteristics

While REST establishes a common set of technologies and best practices, the details of each API can vary based on the design choices. You should be aware of the details and operational characteristics of the ONTAP Select Deploy API before using the API.

Hypervisor host versus ONTAP Select node

A *hypervisor host* is the core hardware platform that hosts an ONTAP Select virtual machine. When an ONTAP Select virtual machine is deployed and active on a hypervisor host, the virtual machine is considered to be an *ONTAP Select node*. With version 3 of the Deploy REST API, the host and node objects are separate and distinct. This allows a one-to-many relationship, where one or more ONTAP Select nodes can run on the same hypervisor host.

Object identifiers

Each resource instance or object is assigned a unique identifier when it is created. These identifiers are globally unique within a specific instance of ONTAP Select Deploy. After issuing an API call that creates a new object instance, the associated id value is returned to the caller in the `location` header of the HTTP response. You can extract the identifier and use it on subsequent calls when referring to the resource instance.



The content and internal structure of the object identifiers can change at any time. You should only use the identifiers on the applicable API calls as needed when referring to the associated objects.

Request identifiers

Every successful API request is assigned a unique identifier. The identifier is returned in the `request-id` header of the associated HTTP response. You can use a request identifier to collectively refer to the activities of a single specific API request-response transaction. For example, you can retrieve all the event messages for a transaction based on the request id.

Synchronous and asynchronous calls

There are two primary ways that a server performs an HTTP request received from a client:

- Synchronous
  The server performs the request immediately and responds with a status code of 200, 201, or 204.
- Asynchronous
  The server accepts the request and responds with a status code of 202. This indicates the server has accepted the client request and started a background task to complete the request. Final success or failure is not immediately available and must be determined through additional API calls.

Confirm the completion of a long-running job

Generally, any operation that can take a long time to complete is processed asynchronously using a background task at the server. With the Deploy REST API, every background task is anchored by a Job object which tracks the task and provides information, such as the current state. A Job object, including its unique identifier, is returned in the HTTP response after a background task is created.

You can query the Job object directly to determine the success or failure of the associated API call. Refer to *asynchronous processing using the Job object* for additional information.

In addition to using the Job object, there are other ways you can determine the success or failure of a request, including:

- Event messages
  You can retrieve all the event messages associated with a specific API call using the request id returned with the original response. The event messages typically contain an indication of success or failure, and can also be useful when debugging an error condition.
- Resource state or status
  Several of the resources maintain a state or status value which you can query to indirectly determine the success or failure of a request.

Security

The Deploy API uses the following security technologies:

- Transport Layer Security
  All traffic sent over the network between the Deploy server and client is encrypted through TLS. Using the HTTP protocol over an unencrypted channel is not supported. TLS version 1.2 is supported.
- HTTP authentication
  Basic authentication is used for every API transaction. An HTTP header, which includes the user name and password in a base64 string, is added to every request.

Request and response API transaction for ONTAP Select

Every Deploy API call is performed as an HTTP request to the Deploy virtual machine which generates an associated response to the client. This request/response pair is considered an API transaction. Before using the Deploy API, you should be familiar with the input variables available to control a request and the contents of the response output.

Input variables controlling an API request

You can control how an API call is processed through parameters set in the HTTP request.

Request headers

You must include several headers in the HTTP request, including:

- content-type
  If the request body includes JSON, this header must be set to application/json.
- accept
  If the response body will include JSON, this header must be set to application/json.
- authorization
  Basic authentication must be set with the user name and password encoded in a base64 string.

Request body

The content of the request body varies depending on the specific call. The HTTP request body consists of one of the following:

- JSON object with input variables (such as, the name of a new cluster)
- Empty

Filter objects

When issuing an API call that uses GET, you can limit or filter the returned objects based on any attribute. For example, you can specify an exact value to match:

```
<field>=<query value>
```

In addition to an exact match, there are other operators available to return a set of objects over a range of values. ONTAP Select supports the filtering operators shown below.

Operator	Description
=	Equal to
<	Less than
>	Greater than
⇐	Less than or equal to
>=	Greater than or equal to
	Or
!	Not equal to
*	Greedy wildcard

You can also return a set of objects based on whether a specific field is set or not set by using the null keyword or its negation (!null) as part of the query.

Selecting object fields

By default, issuing an API call using GET returns only the attributes that uniquely identify the object or objects. This minimum set of fields acts as a key for each object and varies based on the object type. You can select additional object properties using the fields query parameter in the following ways:

- Inexpensive fields
  Specify `fields=*` to retrieve the object fields that are maintained in local server memory or require little processing to access.
- Expensive fields
  Specify `fields=**` to retrieve all the object fields, including those requiring additional server processing to access.
- Custom field selection
  Use `fields=FIELDNAME` to specify the exact field you want. When requesting multiple fields, the values must be separated using commas without spaces.



As a best practice, you should always identify the specific fields you want. You should only retrieve the set of inexpensive or expensive fields when needed. The inexpensive and expensive classification is determined by NetApp based on internal performance analysis. The classification for a given field can change at any time.

Sort objects in the output set

The records in a resource collection are returned in the default order defined by the object. You can change the order using the order_by query parameter with the field name and sort direction as follows:
`order_by=<field name> asc|desc`

For example, you can sort the type field in descending order followed by id in ascending order:
`order_by=type desc, id asc`

When including multiple parameters, you must separate the fields with a comma.

Pagination

When issuing an API call using GET to access a collection of objects of the same type, all matching objects are returned by default. If needed, you can limit the number of records returned using the max_records query parameter with the request. For example:

```
max_records=20
```

If needed, you can combine this parameter with other query parameters to narrow the result set. For example, the following returns up to 10 system events generated after the specified time:

```
time⇒ 2019-04-04T15:41:29.140265Z&max_records=10
```

You can issue multiple requests to page through the events (or any object type). Each subsequent API call should use a new time value based on the latest event in the last result set.

Interpret an API response

Each API request generates a response back to the client. You can examine the response to determine whether it was successful and retrieve additional data as needed.

HTTP status code

The HTTP status codes used by the Deploy REST API are described below.

Code	Meaning	Description
200	OK	Indicates success for calls that do not create a new object.
201	Created	An object is successfully created; the location response header includes the unique identifier for the object.
202	Accepted	A long-running background job has been started to perform the request, but the operation has not completed yet.
400	Bad request	The request input is not recognized or is inappropriate.
403	Forbidden	Access is denied due to an authorization error.
404	Not found	The resource referred to in the request does not exist.
405	Method not allowed	The HTTP verb in the request is not supported for the resource.
409	Conflict	An attempt to create an object failed because the object already exists.
500	Internal error	A general internal error occurred at the server.
501	Not implemented	The URI is known but is not capable of performing the request.

Response headers

Several headers are included in the HTTP response generated by the Deploy server, including:

- request-id
  Every successful API request is assigned a unique request identifier.
- location
  When an object is created, the location header includes the complete URL to the new object including the unique object identifier.

Response body

The content of the response associated with an API request differs based on the object, processing type, and the success or failure of the request. The response body is rendered in JSON.

- Single object
  A single object can be returned with a set of fields based on the request. For example, you can use GET to retrieve selected properties of a cluster using the unique identifier.
- Multiple objects
  Multiple objects from a resource collection can be returned. In all cases, there is a consistent format used, with `num_records` indicating the number of records and records containing an array of the object instances. For example, you can retrieve all the nodes defined in a specific cluster.
- Job object
  If an API call is processed asynchronously, a Job object is returned which anchors the background task. For example, the POST request used to deploy a cluster is processed asynchronously and returns a Job object.
- Error object
  If an error occurs, an Error object is always returned. For example, you will receive an error when attempting to create a cluster with a name that already exists.
- Empty
  In certain cases, no data is returned and the response body is empty. For example, the response body is empty after using DELETE to delete an existing host.

Asynchronous processing using the Job object for ONTAP Select

Some of the Deploy API calls, particularly those that create or modify a resource, can take longer to complete than other calls. ONTAP Select Deploy processes these long-running requests asynchronously.

Asynchronous requests described using Job object

After making an API call that runs asynchronously, the HTTP response code 202 indicates the request has been successfully validated and accepted, but not yet completed. The request is processed as a background task which continues to run after the initial HTTP response to the client. The response includes the Job object anchoring the request, including its unique identifier.

> You should refer to the ONTAP Select Deploy online documentation page to determine which API calls operate asynchronously.

Query the Job object associated with an API request

The Job object returned in the HTTP response contains several properties. You can query the state property to determine if the request completed successfully. A Job object can be in one of the following states:

- Queued
- Running
- Success
- Failure

There are two techniques you can use when polling a Job object to detect a terminal state for the task, either

success or failure:

- Standard polling request
  Current Job state is returned immediately
- Long polling request
  Job state is returned only when one of the following occurs:
  - State has changed more recently than the date-time value provided on the poll request
  - Timeout value has expired (1 to 120 seconds)

Standard polling and long polling use the same API call to query a Job object. However, a long polling request includes two query parameters: `poll_timeout` and `last_modified`.



You should always use long polling to reduce the workload on the Deploy virtual machine.

General procedure for issuing an asynchronous request

You can use the following high-level procedure to complete an asynchronous API call:

1. Issue the asynchronous API call.
2. Receive an HTTP response 202 indicating successful acceptance of the request.
3. Extract the identifier for the Job object from the response body.
4. Within a loop, perform the following in each cycle:
   a. Get the current state of the Job with a long-poll request
   b. If the Job is in a non-terminal state (queued, running), perform loop again.
5. Stop when the Job reaches a terminal state (success, failure).

Access with a browser

Learn about accessing the ONTAP Select Deploy API with a browser

There are several things you should be aware of before using the Deploy online documentation page.

Deployment plan

If you intend to issue API calls as part of performing specific deployment or administrative tasks, you should consider creating a deployment plan. These plans can be formal or informal, and generally contain your goals and the API calls to be used. Refer to Workflow processes using the Deploy REST API for more information.

JSON examples and parameter definitions

Each API call is described on the documentation page using a consistent format. The content includes implementation notes, query parameters, and HTTP status codes. In addition, you can display details about the JSON used with the API requests and responses as follows:

- Example Value
  If you select *Example Value* on an API call, a typical JSON structure for the call is displayed. You can modify the example as needed and use it as input for your request.

- Model
  If you select *Model*, a complete list of the JSON parameters is displayed, with a description for each parameter.

Caution when issuing API calls

All API operations you perform using the Deploy documentation page are live operations. You should be careful not to mistakenly create, update, or delete configuration or other data.

Access the ONTAP Select Deploy documentation page

You must access the ONTAP Select Deploy online documentation page to display the API documentation, as well as to manually issue an API call.

Before you begin

You must have the following:

- IP address or domain name of the ONTAP Select Deploy virtual machine
- User name and password for the administrator

Steps

1. Type the URL in your browser and select **Enter**:

   ```
   https://<ip_address>/api/ui
   ```

2. Sign in using the administrator user name and password.

Result

The Deploy documentation web page is displayed with the calls organized by category at the bottom of the page.

Understand and execute an ONTAP Select Deploy API call

The details of all the API calls are documented and displayed using a common format on the ONTAP Select Deploy online documentation web page. By understanding a single API call, you can access and interpret the details of all the API calls.

Before you begin

You must be signed in to the ONTAP Select Deploy online documentation web page. You must have the unique identifier assigned to your ONTAP Select cluster when the cluster was created.

About this task

You can retrieve the configuration information describing an ONTAP Select cluster using its unique identifier. In this example, all fields classified as inexpensive are returned. However, as a best practice you should request only the specific fields that are needed.

Steps

1. On the main page, scroll to the bottom and select **Cluster**.
2. Select **GET /clusters/{cluster_id}** to display the details of the API call used to return information about an ONTAP Select cluster.

Workflow processes

Learn about using the ONTAP Select Deploy API workflows

You should prepare to review and use the workflow processes.

Understand the API calls used in the workflows

The ONTAP Select online documentation page includes the details of every REST API call. Rather than repeat those details here, each API call used in the workflow samples includes only the information you need to locate the call on the documentation page. After locating a specific API call, you can review the complete details of the call, including the input parameters, output formats, HTTP status codes, and request processing type.

The following information is included for each API call within a workflow to help locate the call on the documentation page:

- Category
  The API calls are organized on the documentation page into functionally related areas or categories. To locate a specific API call, scroll to the bottom of the page and select the applicable API category.
- HTTP verb
  The HTTP verb identifies the action performed on a resource. Each API call is executed through a single HTTP verb.
- Path
  The path determines the specific resource which the action applies to as part of performing a call. The path string is appended to the core URL to form the complete URL identifying the resource.

Construct a URL to directly access the REST API

In addition to the ONTAP Select documentation page, you can also access the Deploy REST API directly through a programming language such as Python. In this case, the core URL is slightly different than the URL used when accessing the online documentation page. When accessing the API directly, you must append /api to the domain and port string. For example:
`http://deploy.mycompany.com/api`

Create an ONTAP Select single-node evaluation cluster on ESXi

You can deploy a single-node ONTAP Select cluster on a VMware ESXi host managed by vCenter. The cluster is created with an evaluation license.

The cluster creation workflow differs in the following situations:

- The ESXi host is not managed by vCenter (standalone host).
- Multiple nodes or hosts are used within the cluster.
- The cluster is deployed in a production environment with a purchased license.
- The KVM hypervisor is used instead of VMware ESXi.

Step 1: Register vCenter server credential

When deploying to an ESXi host managed by a vCenter server, you need to add a credential before registering the host. The Deploy administration utility can then use the credential to authenticate to vCenter.

Category	HTTP verb	Path
Deploy	POST	/security/credentials

Curl

```
curl -iX POST -H 'Content-Type: application/json' -u admin:<password> -k
-d @step01 'https://10.21.191.150/api/security/credentials'
```

JSON input (step01)

```
{
  "hostname": "vcenter.company-demo.com",
  "type": "vcenter",
  "username": "misteradmin@vsphere.local",
  "password": "mypassword"
}
```

Processing type

Asynchronous

Output

- Credential ID in the location response header
- Job object

Step 2: Register a hypervisor host

Add a hypervisor host where the virtual machine containing the ONTAP Select node will run.

Category	HTTP verb	Path
Cluster	POST	/hosts

Curl

```
curl -iX POST -H 'Content-Type: application/json' -u admin:<password> -k
-d @step02 'https://10.21.191.150/api/hosts'
```

JSON input (step02)

```
{
  "hosts": [
    {
      "hypervisor_type": "ESX",
      "management_server": "vcenter.company-demo.com",
```

```
        "name": "esx1.company-demo.com"
      }
    ]
}
```

Processing type

Asynchronous

Output

- Host ID in the location response header
- Job object

Step 3: Create a cluster

When you create an ONTAP Select cluster, the basic cluster configuration is registered and Deploy generates the node names automatically.

Category	HTTP verb	Path
Cluster	POST	/clusters

Curl

You should set the query parameter node_count to 1 for a single-node cluster.

```
curl -iX POST -H 'Content-Type: application/json' -u admin:<password> -k
-d @step03 'https://10.21.191.150/api/clusters? node_count=1'
```

JSON input (step03)

```
{
  "name": "my_cluster"
}
```

Processing type

Synchronous

Output

- Cluster ID in the location response header

Step 4: Configure the cluster

You must provide the following attributes as part of configuring the cluster.

Category	HTTP verb	Path
Cluster	PATCH	/clusters/{cluster_id}

Curl

Provide the cluster ID.

```
curl -iX PATCH -H 'Content-Type: application/json' -u admin:<password> -k
-d @step04 'https://10.21.191.150/api/clusters/CLUSTERID'
```

JSON input (step04)

```
{
  "dns_info": {
    "domains": ["lab1.company-demo.com"],
    "dns_ips": ["10.206.80.135", "10.206.80.136"]
    },
    "ontap_image_version": "9.5",
    "gateway": "10.206.80.1",
    "ip": "10.206.80.115",
    "netmask": "255.255.255.192",
    "ntp_servers": {"10.206.80.183"}
}
```

Processing type

Synchronous

Output

None

Step 5: Retrieve the node name

The Deploy administration utility automatically generates the node identifiers and names when a cluster is created. Before you can configure a node, you need to retrieve the assigned ID.

Category	HTTP verb	Path
Cluster	GET	/clusters/{cluster_id}/nodes

Curl

Provide the cluster ID.

```
curl -iX GET -u admin:<password> -k
'https://10.21.191.150/api/clusters/CLUSTERID/nodes?fields=id,name'
```

Processing type

Synchronous

Output

- Array records each describing a single node with the unique ID and name

Step 6: Configure the nodes

Provide the basic configuration for the node, which is the first of three API calls used to configure a node.

Category	HTTP verb	Path
Cluster	PATH	/clusters/{cluster_id}/nodes/{node_id}

Curl

Provide the cluster ID and node ID.

```
curl -iX PATCH -H 'Content-Type: application/json' -u admin:<password> -k
-d @step06 'https://10.21.191.150/api/clusters/CLUSTERID/nodes/NODEID'
```

JSON input (step06)

Provide the host ID where the ONTAP Select node will run.

```
{
  "host": {
    "id": "HOSTID"
    },
  "instance_type": "small",
  "ip": "10.206.80.101",
  "passthrough_disks": false
}
```

Processing type

Synchronous

Output

None

Step 7: Retrieve the node networks

Identify the data and management networks used by the node in the single-node cluster. The internal network is not used with a single-node cluster.

Category	HTTP verb	Path
Cluster	GET	/clusters/{cluster_id}/nodes/{node_id}/networks

Curl

Provide the cluster ID and node ID.

```
curl -iX GET -u admin:<password> -k 'https://10.21.191.150/api/
```

```
clusters/CLUSTERID/nodes/NODEID/networks?fields=id,purpose'
```

Processing type

Synchronous

Output

- Array of two records each describing a single network for the node, including the unique ID and purpose

Step 8: Configure the node networking

Configure the data and management networks. The internal network is not used with a single-node cluster.

Issue the following API call two times, once for each network.

Category	HTTP verb	Path
Cluster	PATCH	/clusters/{cluster_id}/nodes/{node_id}/networks/{network_id}

Curl

Provide the cluster ID, node ID, and network ID.

```
curl -iX PATCH -H 'Content-Type: application/json' -u admin:<password> -k
-d @step08 'https://10.21.191.150/api/clusters/
CLUSTERID/nodes/NODEID/networks/NETWORKID'
```

JSON input (step08)

Provide the name of the network.

```
{
  "name": "sDOT_Network"
}
```

Processing type

Synchronous

Output

None

Step 9: Configure the node storage pool

The final step in configuring a node is to attach a storage pool. You can determine the available storage pools through the vSphere web client, or optionally through the Deploy REST API.

Category	HTTP verb	Path
Cluster	PATCH	/clusters/{cluster_id}/nodes/{node_id}/networks/{network_id}

Curl

Provide the cluster ID, node ID, and network ID.

```
curl -iX PATCH -H 'Content-Type: application/json' -u admin:<password> -k
-d @step09 'https://10.21.191.150/api/clusters/ CLUSTERID/nodes/NODEID'
```

JSON input (step09)

The pool capacity is 2 TB.

```
{
  "pool_array": [
    {
      "name": "sDOT-01",
      "capacity": 2147483648000
    }
  ]
}
```

Processing type

Synchronous

Output

None

Step 10: Deploy the cluster

After the cluster and node have been configured, you can deploy the cluster.

Category	HTTP verb	Path
Cluster	POST	/clusters/{cluster_id}/deploy

Curl

Provide the cluster ID.

```
curl -iX POST -H 'Content-Type: application/json' -u admin:<password> -k
-d @step10 'https://10.21.191.150/api/clusters/CLUSTERID/deploy'
```

JSON input (step10)

Provide the password for the ONTAP administrator account.

```
{
  "ontap_credentials": {
    "password": "mypassword"
  }
```

```
}
```

Processing type

Asynchronous

Output

- Job object

Related information

Deploy a 90-day evaluation instance of an ONTAP Select cluster

Access with Python

Learn about accessing the ONTAP Select Deploy API using Python

You must prepare the environment before running the sample Python scripts.

Before you run the Python scripts, you must make sure the environment is configured correctly:

- The latest applicable version of Python2 must be installed.
  The sample codes have been tested using Python2. They should also be portable to Python3, but have not been tested for compatibility.
- The Requests and urllib3 libraries must be installed.
  You can use pip or another Python management tool as appropriate for your environment.
- The client workstation where the scripts run must have network access to the ONTAP Select Deploy virtual machine.

In addition, you must have the following information:

- IP address of the Deploy virtual machine
- User name and password of a Deploy administrator account

Learn about the Python scripts for ONTAP Select Deploy

The sample Python scripts allow you to perform several different tasks. You should understand the scripts before using them at a live Deploy instance.

Common design characteristics

The scripts have been designed with the following common characteristics:

- Execute from command line interface at a client machine
  You can run the Python scripts from any properly configured client machine. See *Before you begin* for more information.
- Accept CLI input parameters
  Each script is controlled at the CLI through input parameters.
- Read input file
  Each script reads an input file based on its purpose. When creating or deleting a cluster, you must provide

a JSON configuration file. When adding a node license, you must provide a valid license file.

- Use a common support module
  The common support module *deploy_requests.py* contains a single class. It is imported and used by each of the scripts.

Create a cluster

You can create an ONTAP Select cluster using the script cluster.py. Based on the CLI parameters and contents of the JSON input file, you can modify the script to your deployment environment as follows:

- Hypervisor
  You can deploy to ESXI or KVM (depending on the Deploy release). When deploying to ESXi, the hypervisor can be managed by vCenter or can be a standalone host.
- Cluster size
  You can deploy a single-node or multiple-node cluster.
- Evaluation or production license
  You can deploy a cluster with an evaluation or purchased license for production.

The CLI input parameters for the script include:

- Host name or IP address of the Deploy server
- Password for the admin user account
- Name of the JSON configuration file
- Verbose flag for message output

Add a node license

If you choose to deploy a production cluster, you must add a license for each node using the script *add_license.py*. You can add the license before or after you deploy the cluster.

The CLI input parameters for the script include:

- Host name or IP address of the Deploy server
- Password for the admin user account
- Name of the license file
- ONTAP user name with privileges to add the license
- Password for the ONTAP user

Delete a cluster

You can delete an existing ONTAP Select cluster using the script *delete_cluster.py*.

The CLI input parameters for the script include:

- Host name or IP address of the Deploy server
- Password for the admin user account
- Name of the JSON configuration file

Python code samples

Script to create an ONTAP Select cluster

You can use the following script to create a cluster based on parameters defined within the script and a JSON input file.

```
#!/usr/bin/env python
##-----------------------------------------------------------------------
#
# File: cluster.py
#
# (C) Copyright 2019 NetApp, Inc.
#
# This sample code is provided AS IS, with no support or warranties of
# any kind, including but not limited for warranties of
merchantability
# or fitness of any kind, expressed or implied. Permission to use,
# reproduce, modify and create derivatives of the sample code is
granted
# solely for the purpose of researching, designing, developing and
# testing a software application product for use with NetApp products,
# provided that the above copyright notice appears in all copies and
# that the software application product is distributed pursuant to
terms
# no less restrictive than those set forth herein.
#
##-----------------------------------------------------------------------

import traceback
import argparse
import json
import logging

from deploy_requests import DeployRequests


def add_vcenter_credentials(deploy, config):
    """ Add credentials for the vcenter if present in the config """
    log_debug_trace()

    vcenter = config.get('vcenter', None)
    if vcenter and not deploy.resource_exists('/security/credentials',
                                              'hostname', vcenter[
'hostname']):
        log_info("Registering vcenter {} credentials".format(vcenter[
```

```
    'hostname']))
36          data = {k: vcenter[k] for k in ['hostname', 'username',
    'password']}
37          data['type'] = "vcenter"
38          deploy.post('/security/credentials', data)
39
40
41  def add_standalone_host_credentials(deploy, config):
42      """ Add credentials for standalone hosts if present in the config.
43          Does nothing if the host credential already exists on the
    Deploy.
44      """
45      log_debug_trace()
46
47      hosts = config.get('hosts', [])
48      for host in hosts:
49          # The presense of the 'password' will be used only for
    standalone hosts.
50          # If this host is managed by a vcenter, it should not have a
    host 'password' in the json.
51          if 'password' in host and not deploy.resource_exists(
    '/security/credentials',
52                                                               '
    hostname', host['name']):
53              log_info("Registering host {} credentials".format(host[
    'name']))
54              data = {'hostname': host['name'], 'type': 'host',
55                      'username': host['username'], 'password': host[
    'password']}
56              deploy.post('/security/credentials', data)
57
58
59  def register_unkown_hosts(deploy, config):
60      ''' Registers all hosts with the deploy server.
61          The host details are read from the cluster config json file.
62
63          This method will skip any hosts that are already registered.
64          This method will exit the script if no hosts are found in the
    config.
65      '''
66      log_debug_trace()
67
68      data = {"hosts": []}
69      if 'hosts' not in config or not config['hosts']:
70          log_and_exit("The cluster config requires at least 1 entry in
    the 'hosts' list got {}".format(config))
```

```
71
72      missing_host_cnt = 0
73      for host in config['hosts']:
74          if not deploy.resource_exists('/hosts', 'name', host['name']):
75              missing_host_cnt += 1
76              host_config = {"name": host['name'], "hypervisor_type":
    host['type']}
77              if 'mgmt_server' in host:
78                  host_config["management_server"] = host['mgmt_server']
79                  log_info(
80                     "Registering from vcenter {mgmt_server}".format(
    **host))
81
82              if 'password' in host and 'user' in host:
83                  host_config['credential'] = {
84                      "password": host['password'], "username": host[
    'user']}
85
86              log_info("Registering {type} host {name}".format(**host))
87              data["hosts"].append(host_config)
88
89      # only post /hosts if some missing hosts were found
90      if missing_host_cnt:
91          deploy.post('/hosts', data, wait_for_job=True)
92
93
94  def add_cluster_attributes(deploy, config):
95      ''' POST a new cluster with all needed attribute values.
96          Returns the cluster_id of the new config
97      '''
98      log_debug_trace()
99
100     cluster_config = config['cluster']
101     cluster_id = deploy.find_resource('/clusters', 'name',
    cluster_config['name'])
102
103     if not cluster_id:
104         log_info("Creating cluster config named {name}".format(
    **cluster_config))
105
106         # Filter to only the valid attributes, ignores anything else
    in the json
107         data = {k: cluster_config[k] for k in [
108             'name', 'ip', 'gateway', 'netmask', 'ontap_image_version',
    'dns_info', 'ntp_servers']}
109
```

```
110         num_nodes = len(config['nodes'])
111
112         log_info("Cluster properties: {}".format(data))
113
114         resp = deploy.post('/v3/clusters?node_count={}'.format
    (num_nodes), data)
115         cluster_id = resp.headers.get('Location').split('/')[-1]
116
117     return cluster_id
118
119
120 def get_node_ids(deploy, cluster_id):
121     ''' Get the the ids of the nodes in a cluster. Returns a list of
    node_ids.'''
122     log_debug_trace()
123
124     response = deploy.get('/clusters/{}/nodes'.format(cluster_id))
125     node_ids = [node['id'] for node in response.json().get('records')]
126     return node_ids
127
128
129 def add_node_attributes(deploy, cluster_id, node_id, node):
130     ''' Set all the needed properties on a node '''
131     log_debug_trace()
132
133     log_info("Adding node '{}' properties".format(node_id))
134
135     data = {k: node[k] for k in ['ip', 'serial_number', '
    instance_type',
136                                  'is_storage_efficiency_enabled'] if k
    in node}
137     # Optional: Set a serial_number
138     if 'license' in node:
139         data['license'] = {'id': node['license']}
140
141     # Assign the host
142     host_id = deploy.find_resource('/hosts', 'name', node['host_name
    '])
143     if not host_id:
144         log_and_exit("Host names must match in the 'hosts' array, and
    the nodes.host_name property")
145
146     data['host'] = {'id': host_id}
147
148     # Set the correct raid_type
149     is_hw_raid = not node['storage'].get('disks')  # The presence of a
```

```
    list of disks indicates sw_raid
150     data['passthrough_disks'] = not is_hw_raid
151
152     # Optionally set a custom node name
153     if 'name' in node:
154         data['name'] = node['name']
155
156     log_info("Node properties: {}".format(data))
157     deploy.patch('/clusters/{}/nodes/{}'.format(cluster_id, node_id),
    data)
158
159
160 def add_node_networks(deploy, cluster_id, node_id, node):
161     ''' Set the network information for a node '''
162     log_debug_trace()
163
164     log_info("Adding node '{}' network properties".format(node_id))
165
166     num_nodes = deploy.get_num_records('/clusters/{}/nodes'.format
    (cluster_id))
167
168     for network in node['networks']:
169
170         # single node clusters do not use the 'internal' network
171         if num_nodes == 1 and network['purpose'] == 'internal':
172             continue
173
174         # Deduce the network id given the purpose for each entry
175         network_id = deploy.find_resource(
    '/clusters/{}/nodes/{}/networks'.format(cluster_id, node_id),
176                                           'purpose', network['purpose
    '])
177         data = {"name": network['name']}
178         if 'vlan' in network and network['vlan']:
179             data['vlan_id'] = network['vlan']
180
181         deploy.patch('/clusters/{}/nodes/{}/networks/{}'.format
    (cluster_id, node_id, network_id), data)
182
183
184 def add_node_storage(deploy, cluster_id, node_id, node):
185     ''' Set all the storage information on a node '''
186     log_debug_trace()
187
188     log_info("Adding node '{}' storage properties".format(node_id))
189     log_info("Node storage: {}".format(node['storage']['pools']))
```

```
190
191     data = {'pool_array': node['storage']['pools']}  # use all the
    json properties
192     deploy.post(
193         '/clusters/{}/nodes/{}/storage/pools'.format(cluster_id,
    node_id), data)
194
195     if 'disks' in node['storage'] and node['storage']['disks']:
196         data = {'disks': node['storage']['disks']}
197         deploy.post(
198             '/clusters/{}/nodes/{}/storage/disks'.format(cluster_id,
    node_id), data)
199
200
201 def create_cluster_config(deploy, config):
202     ''' Construct a cluster config in the deploy server using the
    input json data '''
203     log_debug_trace()
204
205     cluster_id = add_cluster_attributes(deploy, config)
206
207     node_ids = get_node_ids(deploy, cluster_id)
208     node_configs = config['nodes']
209
210     for node_id, node_config in zip(node_ids, node_configs):
211         add_node_attributes(deploy, cluster_id, node_id, node_config)
212         add_node_networks(deploy, cluster_id, node_id, node_config)
213         add_node_storage(deploy, cluster_id, node_id, node_config)
214
215     return cluster_id
216
217
218 def deploy_cluster(deploy, cluster_id, config):
219     ''' Deploy the cluster config to create the ONTAP Select VMs. '''
220     log_debug_trace()
221     log_info("Deploying cluster: {}".format(cluster_id))
222
223     data = {'ontap_credential': {'password': config['cluster'][
    'ontap_admin_password']}}
224     deploy.post('/clusters/{}/deploy?inhibit_rollback=true'.format
    (cluster_id),
225                 data, wait_for_job=True)
226
227
228 def log_debug_trace():
229     stack = traceback.extract_stack()
```

```
    parent_function = stack[-2][2]
    logging.getLogger('deploy').debug('Calling %s()' %
parent_function)


def log_info(msg):
    logging.getLogger('deploy').info(msg)


def log_and_exit(msg):
    logging.getLogger('deploy').error(msg)
    exit(1)


def configure_logging(verbose):
    FORMAT = '%(asctime)-15s:%(levelname)s:%(name)s: %(message)s'
    if verbose:
        logging.basicConfig(level=logging.DEBUG, format=FORMAT)
    else:
        logging.basicConfig(level=logging.INFO, format=FORMAT)
        logging.getLogger('requests.packages.urllib3.connectionpool'
).setLevel(
            logging.WARNING)


def main(args):
    configure_logging(args.verbose)
    deploy = DeployRequests(args.deploy, args.password)

    with open(args.config_file) as json_data:
        config = json.load(json_data)

        add_vcenter_credentials(deploy, config)

        add_standalone_host_credentials(deploy, config)

        register_unkown_hosts(deploy, config)

        cluster_id = create_cluster_config(deploy, config)

        deploy_cluster(deploy, cluster_id, config)


def parseArgs():
    parser = argparse.ArgumentParser(description='Uses the ONTAP
Select Deploy API to construct and deploy a cluster.')
```

```
    parser.add_argument('-d', '--deploy', help='Hostname or IP address
of Deploy server')
    parser.add_argument('-p', '--password', help='Admin password of
Deploy server')
    parser.add_argument('-c', '--config_file', help='Filename of the
cluster config')
    parser.add_argument('-v', '--verbose', help='Display extra
debugging messages for seeing exact API calls and responses',
                        action='store_true', default=False)
    return parser.parse_args()

if __name__ == '__main__':
    args = parseArgs()
    main(args)
```

JSON for script to create an ONTAP Select cluster

When creating or deleting an ONTAP Select cluster using the Python code samples, you must provide a JSON file as input to the script. You can copy and modify the appropriate JSON sample based on your deployment plans.

Single-node cluster on ESXi

```
{
  "hosts": [
    {
      "password": "mypassword1",
      "name": "host-1234",
      "type": "ESX",
      "username": "admin"
    }
  ],

  "cluster": {
    "dns_info": {
      "domains": ["lab1.company-demo.com", "lab2.company-demo.com",
        "lab3.company-demo.com", "lab4.company-demo.com"
        ],

      "dns_ips": ["10.206.80.135", "10.206.80.136"]
      },
      "ontap_image_version": "9.7",
      "gateway": "10.206.80.1",
      "ip": "10.206.80.115",
      "name": "mycluster",
```

```
      "ntp_servers": ["10.206.80.183", "10.206.80.142"],
      "ontap_admin_password": "mypassword2",
      "netmask": "255.255.254.0"
  },

  "nodes": [
    {
      "serial_number": "3200000nn",
      "ip": "10.206.80.114",
      "name": "node-1",
      "networks": [
        {
          "name": "ontap-external",
          "purpose": "mgmt",
          "vlan": 1234
        },
        {
          "name": "ontap-external",
          "purpose": "data",
          "vlan": null
        },
        {
          "name": "ontap-internal",
          "purpose": "internal",
          "vlan": null
        }
      ],
      "host_name": "host-1234",
      "is_storage_efficiency_enabled": false,
      "instance_type": "small",
      "storage": {
        "disk": [],
        "pools": [
          {
            "name": "storage-pool-1",
            "capacity": 4802666790125
          }
        ]
      }
    }
  ]
}
```

Single-node cluster on ESXi using vCenter

```
{
  "hosts": [
    {
      "name":"host-1234",
      "type":"ESX",
      "mgmt_server":"vcenter-1234"
    }
  ],

  "cluster": {
    "dns_info": {"domains": ["lab1.company-demo.com", "lab2.company-
demo.com",
      "lab3.company-demo.com", "lab4.company-demo.com"
      ],
      "dns_ips": ["10.206.80.135","10.206.80.136"]
  },

  "ontap_image_version":"9.7",
  "gateway":"10.206.80.1",
  "ip":"10.206.80.115",
  "name":"mycluster",
  "ntp_servers": ["10.206.80.183","10.206.80.142"],
  "ontap_admin_password":"mypassword2",
  "netmask":"255.255.254.0"
  },

  "vcenter": {
    "password":"mypassword2",
    "hostname":"vcenter-1234",
    "username":"selectadmin"
  },

  "nodes": [
    {
      "serial_number": "3200000nn",
      "ip":"10.206.80.114",
      "name":"node-1",
      "networks": [
        {
          "name":"ONTAP-Management",
          "purpose":"mgmt",
          "vlan":null
        },
        {
          "name": "ONTAP-External",
          "purpose":"data",
```

```
          "vlan":null
        },
        {
          "name": "ONTAP-Internal",
          "purpose":"internal",
          "vlan":null
        }
      ],

      "host_name": "host-1234",
      "is_storage_efficiency_enabled": false,
      "instance_type": "small",
      "storage": {
        "disk":[],
        "pools": [
          {
            "name": "storage-pool-1",
            "capacity":5685190380748
          }
        ]
      }
    }
  ]
}
```

Single-node cluster on KVM

```
{
  "hosts": [
    {
      "password": "mypassword1",
      "name":"host-1234",
      "type":"KVM",
      "username":"root"
    }
  ],

  "cluster": {
    "dns_info": {
      "domains": ["lab1.company-demo.com", "lab2.company-demo.com",
        "lab3.company-demo.com", "lab4.company-demo.com"
      ],

      "dns_ips": ["10.206.80.135", "10.206.80.136"]
    },
```

```
    "ontap_image_version": "9.7",
    "gateway":"10.206.80.1",
    "ip":"10.206.80.115",
    "name":"CBF4ED97",
    "ntp_servers": ["10.206.80.183", "10.206.80.142"],
    "ontap_admin_password": "mypassword2",
    "netmask":"255.255.254.0"
  },
  "nodes": [
    {
      "serial_number":"3200000nn",
      "ip":"10.206.80.115",
      "name": "node-1",
      "networks": [
        {
          "name": "ontap-external",
          "purpose": "mgmt",
          "vlan":1234
        },
        {
          "name": "ontap-external",
          "purpose": "data",
          "vlan": null
        },
        {
          "name": "ontap-internal",
          "purpose": "internal",
          "vlan": null
        }
      ],

      "host_name": "host-1234",
      "is_storage_efficiency_enabled": false,
      "instance_type": "small",
      "storage": {
        "disk": [],
        "pools": [
          {
            "name": "storage-pool-1",
            "capacity": 4802666790125
          }
        ]
      }
    }
  ]
```

```
}
```

Script to add an ONTAP Select node license

You can use the following script to add a license for an ONTAP Select node.

```
#!/usr/bin/env python
##-----------------------------------------------------------------------
#
# File: add_license.py
#
# (C) Copyright 2019 NetApp, Inc.
#
# This sample code is provided AS IS, with no support or warranties of
# any kind, including but not limited for warranties of
merchantability
# or fitness of any kind, expressed or implied. Permission to use,
# reproduce, modify and create derivatives of the sample code is
granted
# solely for the purpose of researching, designing, developing and
# testing a software application product for use with NetApp products,
# provided that the above copyright notice appears in all copies and
# that the software application product is distributed pursuant to
terms
# no less restrictive than those set forth herein.
#
##-----------------------------------------------------------------------

import argparse
import logging
import json

from deploy_requests import DeployRequests


def post_new_license(deploy, license_filename):
    log_info('Posting a new license: {}'.format(license_filename))

    # Stream the file as multipart/form-data
    deploy.post('/licensing/licenses', data={},
                files={'license_file': open(license_filename, 'rb')})

    # Alternative if the NLF license data is converted to a string.
    # with open(license_filename, 'rb') as f:
    #    nlf_data = f.read()
```

```
37      #     r = deploy.post('/licensing/licenses', data={},
38      #                     files={'license_file': (license_filename,
    nlf_data)})
39
40
41  def put_license(deploy, serial_number, data, files):
42      log_info('Adding license for serial number: {}'.format
    (serial_number))
43
44      deploy.put('/licensing/licenses/{}'.format(serial_number), data
    =data, files=files)
45
46
47  def put_used_license(deploy, serial_number, license_filename,
    ontap_username, ontap_password):
48      ''' If the license is used by an 'online' cluster, a
    username/password must be given. '''
49
50      data = {'ontap_username': ontap_username, 'ontap_password':
    ontap_password}
51      files = {'license_file': open(license_filename, 'rb')}
52
53      put_license(deploy, serial_number, data, files)
54
55
56  def put_free_license(deploy, serial_number, license_filename):
57      data = {}
58      files = {'license_file': open(license_filename, 'rb')}
59
60      put_license(deploy, serial_number, data, files)
61
62
63  def get_serial_number_from_license(license_filename):
64      ''' Read the NLF file to extract the serial number '''
65      with open(license_filename) as f:
66          data = json.load(f)
67
68          statusResp = data.get('statusResp', {})
69          serialNumber = statusResp.get('serialNumber')
70          if not serialNumber:
71              log_and_exit("The license file seems to be missing the
    serialNumber")
72
73          return serialNumber
74
75
```

```
76 def log_info(msg):
77     logging.getLogger('deploy').info(msg)
78
79
80 def log_and_exit(msg):
81     logging.getLogger('deploy').error(msg)
82     exit(1)
83
84
85 def configure_logging():
86     FORMAT = '%(asctime)-15s:%(levelname)s:%(name)s: %(message)s'
87     logging.basicConfig(level=logging.INFO, format=FORMAT)
88     logging.getLogger('requests.packages.urllib3.connectionpool'
   ).setLevel(logging.WARNING)
89
90
91 def main(args):
92     configure_logging()
93     serial_number = get_serial_number_from_license(args.license)
94
95     deploy = DeployRequests(args.deploy, args.password)
96
97     # First check if there is already a license resource for this
   serial-number
98     if deploy.find_resource('/licensing/licenses', 'id',
   serial_number):
99
100         # If the license already exists in the Deploy server,
   determine if its used
101         if deploy.find_resource('/clusters', 'nodes.serial_number',
   serial_number):
102
103             # In this case, requires ONTAP creds to push the license
   to the node
104             if args.ontap_username and args.ontap_password:
105                 put_used_license(deploy, serial_number, args.license,
106                                  args.ontap_username, args
   .ontap_password)
107             else:
108                 print("ERROR: The serial number for this license is in
   use. Please provide ONTAP credentials.")
109         else:
110             # License exists, but its not used
111             put_free_license(deploy, serial_number, args.license)
112     else:
113         # No license exists, so register a new one as an available
```

```
    license for later use
114         post_new_license(deploy, args.license)
115
116
117 def parseArgs():
118     parser = argparse.ArgumentParser(description='Uses the ONTAP
    Select Deploy API to add or update a new or used NLF license file.')
119     parser.add_argument('-d', '--deploy', required=True, type=str,
    help='Hostname or IP address of ONTAP Select Deploy')
120     parser.add_argument('-p', '--password', required=True, type=str,
    help='Admin password of Deploy server')
121     parser.add_argument('-l', '--license', required=True, type=str,
    help='Filename of the NLF license data')
122     parser.add_argument('-u', '--ontap_username', type=str,
123                         help='ONTAP Select username with privelege to
    add the license. Only provide if the license is used by a Node.')
124     parser.add_argument('-o', '--ontap_password', type=str,
125                         help='ONTAP Select password for the
    ontap_username. Required only if ontap_username is given.')
126     return parser.parse_args()
127
128 if __name__ == '__main__':
129     args = parseArgs()
130     main(args)
```

Script to delete an ONTAP Select cluster

You can use the following CLI script to delete an existing cluster.

```
 1 #!/usr/bin/env python
 2 ##-----------------------------------------------------------------------
 3 #
 4 # File: delete_cluster.py
 5 #
 6 # (C) Copyright 2019 NetApp, Inc.
 7 #
 8 # This sample code is provided AS IS, with no support or warranties of
 9 # any kind, including but not limited for warranties of merchantability
10 # or fitness of any kind, expressed or implied. Permission to use,
11 # reproduce, modify and create derivatives of the sample code is
   granted
12 # solely for the purpose of researching, designing, developing and
13 # testing a software application product for use with NetApp products,
14 # provided that the above copyright notice appears in all copies and
15 # that the software application product is distributed pursuant to
```

```
terms
# no less restrictive than those set forth herein.
#
##----------------------------------------------------------------------

import argparse
import json
import logging

from deploy_requests import DeployRequests

def find_cluster(deploy, cluster_name):
    return deploy.find_resource('/clusters', 'name', cluster_name)


def offline_cluster(deploy, cluster_id):
    # Test that the cluster is online, otherwise do nothing
    response = deploy.get('/clusters/{}?fields=state'.format(
cluster_id))
    cluster_data = response.json()['record']
    if cluster_data['state'] == 'powered_on':
        log_info("Found the cluster to be online, modifying it to be
powered_off.")
        deploy.patch('/clusters/{}'.format(cluster_id), {'availability
': 'powered_off'}, True)


def delete_cluster(deploy, cluster_id):
    log_info("Deleting the cluster({}).".format(cluster_id))
    deploy.delete('/clusters/{}'.format(cluster_id), True)
    pass


def log_info(msg):
    logging.getLogger('deploy').info(msg)


def configure_logging():
    FORMAT = '%(asctime)-15s:%(levelname)s:%(name)s: %(message)s'
    logging.basicConfig(level=logging.INFO, format=FORMAT)
    logging.getLogger('requests.packages.urllib3.connectionpool'
).setLevel(logging.WARNING)


def main(args):
    configure_logging()
```

```
    deploy = DeployRequests(args.deploy, args.password)

    with open(args.config_file) as json_data:
        config = json.load(json_data)

        cluster_id = find_cluster(deploy, config['cluster']['name'])

        log_info("Found the cluster {} with id: {}.".format(config[
'cluster']['name'], cluster_id))

        offline_cluster(deploy, cluster_id)

        delete_cluster(deploy, cluster_id)


def parseArgs():
    parser = argparse.ArgumentParser(description='Uses the ONTAP Select
Deploy API to delete a cluster')
    parser.add_argument('-d', '--deploy', required=True, type=str,
help='Hostname or IP address of Deploy server')
    parser.add_argument('-p', '--password', required=True, type=str,
help='Admin password of Deploy server')
    parser.add_argument('-c', '--config_file', required=True, type=str,
help='Filename of the cluster json config')
    return parser.parse_args()

if __name__ == '__main__':
    args = parseArgs()
    main(args)
```

Common support Python module for ONTAP Select

All of the Python scripts use a common Python class in a single module.

```
#!/usr/bin/env python
##-----------------------------------------------------------------------
#
# File: deploy_requests.py
#
# (C) Copyright 2019 NetApp, Inc.
#
# This sample code is provided AS IS, with no support or warranties of
# any kind, including but not limited for warranties of merchantability
# or fitness of any kind, expressed or implied. Permission to use,
# reproduce, modify and create derivatives of the sample code is granted
```

```
# solely for the purpose of researching, designing, developing and
# testing a software application product for use with NetApp products,
# provided that the above copyright notice appears in all copies and
# that the software application product is distributed pursuant to terms
# no less restrictive than those set forth herein.
#
##---------------------------------------------------------------------

import json
import logging
import requests

requests.packages.urllib3.disable_warnings()

class DeployRequests(object):
    '''
    Wrapper class for requests that simplifies the ONTAP Select Deploy
    path creation and header manipulations for simpler code.
    '''

    def __init__(self, ip, admin_password):
        self.base_url = 'https://{}/api'.format(ip)
        self.auth = ('admin', admin_password)
        self.headers = {'Accept': 'application/json'}
        self.logger = logging.getLogger('deploy')

    def post(self, path, data, files=None, wait_for_job=False):
        if files:
            self.logger.debug('POST FILES:')
            response = requests.post(self.base_url + path,
                                     auth=self.auth, verify=False,
                                     files=files)
        else:
            self.logger.debug('POST DATA: %s', data)
            response = requests.post(self.base_url + path,
                                     auth=self.auth, verify=False,
                                     json=data,
                                     headers=self.headers)

        self.logger.debug('HEADERS: %s\nBODY: %s', self.filter_headers
(response), response.text)
        self.exit_on_errors(response)

        if wait_for_job and response.status_code == 202:
            self.wait_for_job(response.json())
        return response
```

```
    def patch(self, path, data, wait_for_job=False):
        self.logger.debug('PATCH DATA: %s', data)
        response = requests.patch(self.base_url + path,
                                  auth=self.auth, verify=False,
                                  json=data,
                                  headers=self.headers)
        self.logger.debug('HEADERS: %s\nBODY: %s', self.filter_headers
(response), response.text)
        self.exit_on_errors(response)

        if wait_for_job and response.status_code == 202:
            self.wait_for_job(response.json())
        return response

    def put(self, path, data, files=None, wait_for_job=False):
        if files:
            print('PUT FILES: {}'.format(data))
            response = requests.put(self.base_url + path,
                                    auth=self.auth, verify=False,
                                    data=data,
                                    files=files)
        else:
            self.logger.debug('PUT DATA:')
            response = requests.put(self.base_url + path,
                                    auth=self.auth, verify=False,
                                    json=data,
                                    headers=self.headers)

        self.logger.debug('HEADERS: %s\nBODY: %s', self.filter_headers
(response), response.text)
        self.exit_on_errors(response)

        if wait_for_job and response.status_code == 202:
            self.wait_for_job(response.json())
        return response

    def get(self, path):
        """ Get a resource object from the specified path """
        response = requests.get(self.base_url + path, auth=self.auth,
verify=False)
        self.logger.debug('HEADERS: %s\nBODY: %s', self.filter_headers
(response), response.text)
        self.exit_on_errors(response)
        return response
```

```
    def delete(self, path, wait_for_job=False):
        """ Delete's a resource from the specified path """
        response = requests.delete(self.base_url + path, auth=self.auth,
verify=False)
        self.logger.debug('HEADERS: %s\nBODY: %s', self.filter_headers
(response), response.text)
        self.exit_on_errors(response)

        if wait_for_job and response.status_code == 202:
            self.wait_for_job(response.json())
        return response

    def find_resource(self, path, name, value):
        ''' Returns the 'id' of the resource if it exists, otherwise None
'''
        resource = None
        response = self.get('{path}?{field}={value}'.format(
                            path=path, field=name, value=value))
        if response.status_code == 200 and response.json().get(
'num_records') >= 1:
            resource = response.json().get('records')[0].get('id')
        return resource

    def get_num_records(self, path, query=None):
        ''' Returns the number of records found in a container, or None on
error '''
        resource = None
        query_opt = '?{}'.format(query) if query else ''
        response = self.get('{path}{query}'.format(path=path, query
=query_opt))
        if response.status_code == 200 :
            return response.json().get('num_records')
        return None

    def resource_exists(self, path, name, value):
        return self.find_resource(path, name, value) is not None

    def wait_for_job(self, response, poll_timeout=120):
        last_modified = response['job']['last_modified']
        job_id = response['job']['id']

        self.logger.info('Event: ' + response['job']['message'])

        while True:
            response = self.get('/jobs/{}?fields=state,message&'
                                'poll_timeout={}&last_modified=>={}'
```

```
.format(
                                        job_id, poll_timeout, last_modified))

            job_body = response.json().get('record', {})

            # Show interesting message updates
            message = job_body.get('message', '')
            self.logger.info('Event: ' + message)

            # Refresh the last modified time for the poll loop
            last_modified = job_body.get('last_modified')

            # Look for the final states
            state = job_body.get('state', 'unknown')
            if state in ['success', 'failure']:
                if state == 'failure':
                    self.logger.error('FAILED background job.\nJOB: %s',
job_body)
                    exit(1)    # End the script if a failure occurs
                break

    def exit_on_errors(self, response):
        if response.status_code >= 400:
            self.logger.error('FAILED request to URL: %s\nHEADERS: %s
\nRESPONSE BODY: %s',
                              response.request.url,
                              self.filter_headers(response),
                              response.text)
        response.raise_for_status()    # Displays the response error, and
exits the script

    @staticmethod
    def filter_headers(response):
        ''' Returns a filtered set of the response headers '''
        return {key: response.headers[key] for key in ['Location',
'request-id'] if key in response.headers}
```

Script to resize ONTAP Select cluster nodes

You can use the following script to resize the nodes in an ONTAP Select cluster.

```
1 #!/usr/bin/env python
2 ##-----------------------------------------------------------------
3 #
4 # File: resize_nodes.py
```

```
#
# (C) Copyright 2019 NetApp, Inc.
#
# This sample code is provided AS IS, with no support or warranties of
# any kind, including but not limited for warranties of
merchantability
# or fitness of any kind, expressed or implied. Permission to use,
# reproduce, modify and create derivatives of the sample code is
granted
# solely for the purpose of researching, designing, developing and
# testing a software application product for use with NetApp products,
# provided that the above copyright notice appears in all copies and
# that the software application product is distributed pursuant to
terms
# no less restrictive than those set forth herein.
#
##----------------------------------------------------------------------

import argparse
import logging
import sys

from deploy_requests import DeployRequests


def _parse_args():
    """ Parses the arguments provided on the command line when
executing this
        script and returns the resulting namespace. If all required
arguments
        are not provided, an error message indicating the mismatch is
printed and
        the script will exit.
    """

    parser = argparse.ArgumentParser(description=(
        'Uses the ONTAP Select Deploy API to resize the nodes in the
cluster.'
        ' For example, you might have a small (4 CPU, 16GB RAM per
node) 2 node'
        ' cluster and wish to resize the cluster to medium (8 CPU,
64GB RAM per'
        ' node). This script will take in the cluster details and then
perform'
        ' the operation and wait for it to complete.'
    ))
```

```
    parser.add_argument('--deploy', required=True, help=(
        'Hostname or IP of the ONTAP Select Deploy VM.'
    ))
    parser.add_argument('--deploy-password', required=True, help=(
        'The password for the ONTAP Select Deploy admin user.'
    ))
    parser.add_argument('--cluster', required=True, help=(
        'Hostname or IP of the cluster management interface.'
    ))
    parser.add_argument('--instance-type', required=True, help=(
        'The desired instance size of the nodes after the operation is
complete.'
    ))
    parser.add_argument('--ontap-password', required=True, help=(
        'The password for the ONTAP administrative user account.'
    ))
    parser.add_argument('--ontap-username', default='admin', help=(
        'The username for the ONTAP administrative user account.
Default: admin.'
    ))
    parser.add_argument('--nodes', nargs='+', metavar='NODE_NAME',
help=(
        'A space separated list of node names for which the resize
operation'
        ' should be performed. The default is to apply the resize to
all nodes in'
        ' the cluster. If a list of nodes is provided, it must be
provided in HA'
        ' pairs. That is, in a 4 node cluster, nodes 1 and 2
(partners) must be'
        ' resized in the same operation.'
    ))
    return parser.parse_args()


def _get_cluster(deploy, parsed_args):
    """ Locate the cluster using the arguments provided """

    cluster_id = deploy.find_resource('/clusters', 'ip', parsed_args
.cluster)
    if not cluster_id:
        return None
    return deploy.get('/clusters/%s?fields=nodes' % cluster_id).
json()['record']


```

```
78 def _get_request_body(parsed_args, cluster):
79     """ Build the request body """
80
81     changes = {'admin_password': parsed_args.ontap_password}
82
83     # if provided, use the list of nodes given, else use all the nodes
   in the cluster
84     nodes = [node for node in cluster['nodes']]
85     if parsed_args.nodes:
86         nodes = [node for node in nodes if node['name'] in
   parsed_args.nodes]
87
88     changes['nodes'] = [
89         {'instance_type': parsed_args.instance_type, 'id': node['id']}
   for node in nodes]
90
91     return changes
92
93
94 def main():
95     """ Set up the resize operation by gathering the necessary data
   and then send
96         the request to the ONTAP Select Deploy server.
97     """
98
99     logging.basicConfig(
100         format='[%(asctime)s] [%(levelname)5s] %(message)s', level
   =logging.INFO,)
101
102     logging.getLogger('requests.packages.urllib3').setLevel(logging
   .WARNING)
103
104     parsed_args = _parse_args()
105     deploy = DeployRequests(parsed_args.deploy, parsed_args
   .deploy_password)
106
107     cluster = _get_cluster(deploy, parsed_args)
108     if not cluster:
109         deploy.logger.error(
110             'Unable to find a cluster with a management IP of %s' %
   parsed_args.cluster)
111         return 1
112
113     changes = _get_request_body(parsed_args, cluster)
114     deploy.patch('/clusters/%s' % cluster['id'], changes,
   wait_for_job=True)
```

```
115
116 if __name__ == '__main__':
117     sys.exit(main())
```

Use the CLI

Sign in to ONTAP Select Deploy using SSH

You need to sign in to the Deploy management shell using SSH. After signing in, you can issue CLI commands to create an ONTAP Select cluster and perform related administrative procedures.

Before you begin

You must have the current password for the Deploy administrator (admin) account. If you are signing in for the first time and used vCenter to install the Deploy virtual machine, you should use the password set during installation.

Steps

1. Sign in using the administrator account and management IP address of the Deploy virtual machine; for example:

   ```
   ssh admin@<10.235.82.22>
   ```

2. If this is the first time signing in and you did not install Deploy using the wizard available with vCenter, provide the following configuration information when prompted:
   - New password for the administrator account (required)
   - Company name (required)
   - Proxy URL (optional)
3. Type **?** and select **Enter** to display a list of the available management shell commands.

Deploy an ONTAP Select cluster using the CLI

You can use the command line interface provided with the ONTAP Select Deploy administration utility to create a single-node or multi-node ONTAP Select cluster.

Step 1: Prepare for the deployment

Before creating an ONTAP Select cluster on a hypervisor, you should understand the required preparation.

Steps

1. Prepare to attach storage to the ONTAP Select node

   Hardware RAID

   If you use a local hardware RAID controller, you must create at least one datastore (ESXi) or one storage pool (KVM) at each node for the system data as well as the root and data aggregates. You must attach the storage pool as part of configuring the ONTAP Select node.

   Software RAID

   If you use software RAID, you must create at least one datastore (ESXi) or one storage pool (KVM) for the system data and make sure the SSD drives are available for the root and data aggregates. You must attach the storage pool and disks as part of configuring the ONTAP Select node.

2. Available ONTAP Select versions

   The Deploy administration utility contains a single version of ONTAP Select. If you want to deploy clusters using an earlier version of ONTAP Select, you must first add the ONTAP Select image to your Deploy instance.

3. License ONTAP Select for a production deployment

   Before deploying an ONTAP Select cluster in a production environment, you must purchase a storage capacity license and download the associated license file. You can license the storage at each node using the *Capacity Tiers* model or license a shared pool using the *Capacity Pools* model.

Step 2: Upload and register a license file

After acquiring a license file with storage capacity, you must upload the file containing the license to the Deploy virtual machine and register it.

If you are deploying a cluster for evaluation only, you can skip this step.

Before you begin

You must have the password for the admin user account.

Steps

1. In a command shell on your local workstation, use the sftp utility to upload the license file to the Deploy virtual machine.

   Example output:

   ```
   sftp admin@10.234.81.101 (provide password when prompted)
   put NLF-320000nnn.txt
   exit
   ```

2. Sign in to the Deploy utility CLI with the administrator account using SSH.
3. Register the license:

   ```
   license add -file-name <file_name>
   ```

   Provide the administrator account password when prompted.

4. Display the licenses in the system to confirm the license was added properly:

   ```
   license show
   ```

Step 3: Add hypervisor hosts

You must register each hypervisor host where an ONTAP Select node will run.

KVM

You must register a hypervisor host where the ONTAP Select node will run. As part of this, the Deploy administration utility authenticates to the KVM host.

About this task

If more than one hypervisor host is needed, use this procedure to add each host.

Steps

1. Sign in to the Deploy utility CLI using SSH with the administrator account.
2. Register the host:

   ```
   host register -name <FQDN|IP> -hypervisor-type KVM -username
   <KVM_username>
   ```

   Example output:

   ```
   host register -name 10.234.81.14 -hypervisor-type KVM -username root
   ```

   Provide the password for the host account when prompted.

3. Display the state of the host and confirm it is authenticated:

   ```
   host show -name <FQDN|IP> -detailed
   ```

   Example output:

   ```
   host show -name 10.234.81.14 -detailed
   ```

ESXi

As part of this, the Deploy administration utility authenticates either to the vCenter server managing the host or directly to the ESXi standalone host.

About this task

Before you register a host that is managed by vCenter, you must add a management server account for the vCenter server. If the host is not managed by vCenter, you can provide the host credential as part of registering the host. You should use this procedure to add each host.

Steps

1. Sign in to the Deploy utility CLI using SSH with the administrator account.
2. If the host is managed by a vCenter server, add the vCenter account credential:

   ```
   credential add -hostname <FQDN|IP> -type vcenter -username
   ```

```
<vcenter_username>
```

Example output:

```
credential add -hostname vc.select.company-demo.com -type vcenter
-username administrator@vsphere.local
```

3. Register the host:
   - Register a standalone host not managed by vCenter:

     ```
     host register -name <FQDN|IP> -hypervisor-type ESX -username
     <esx_username>
     ```

   - Register a host managed by vCenter:

     ```
     host register -name <FQDN|IP> -hypervisor-type ESX -mgmt-server
     <FQDN|IP>
     ```

     Example output:

     ```
     host register -name 10.234.81.14 -hypervisor-type ESX -mgmt
     -server vc.select.company-demo.com
     ```

4. Display the state of the host and confirm it is authenticated.

   ```
   host show -name <FQDN|IP> -detailed
   ```

   Example output:

   ```
   host show -name 10.234.81.14 -detailed
   ```

Step 4: Create and configure an ONTAP Select cluster

You must create and then configure the ONTAP Select cluster. After the cluster is configured, you can configure the individual nodes.

Before you begin

Decide how many nodes the cluster contains and have the associated configuration information.

About this task

When you create an ONTAP Select cluster, the Deploy utility automatically generates the node names based on the cluster name and node count that you provide. Deploy also generates the unique node identifiers.

Steps

1. Sign in to the Deploy utility CLI using SSH with the administrator account.
2. Create the cluster:

   ```
   cluster create -name <cluster_name> -node-count <count>
   ```

   Example output:

   ```
   cluster create -name test-cluster -node-count 1
   ```

3. Configure the cluster:

   ```
   cluster modify -name <cluster_name> -mgmt-ip <IP_address> -netmask
   <netmask> -gateway <IP_address> -dns-servers <FQDN|IP>_LIST -dns-domains
   <domain_list>
   ```

   Example output:

   ```
   cluster modify -name test-cluster -mgmt-ip 10.234.81.20 -netmask
   255.255.255.192
   -gateway 10.234.81.1 -dns-servers 10.221.220.10 -dnsdomains
   select.company-demo.com
   ```

4. Display the configuration and state of the cluster:

   ```
   cluster show -name <cluster_name> -detailed
   ```

Step 5: Configure an ONTAP Select node

You must configure each of the nodes in the ONTAP Select cluster.

Before you begin

- Verify that you have the configuration information for the node.
- Verify that the Capacity Tier or Capacity Pool license file is uploaded and installed at the Deploy utility.

About this task

You should use this procedure to configure each node. A Capacity Tier license is applied to the node in this example.

Steps

1. Sign in to the Deploy utility CLI using SSH with the administrator account.
2. Determine the names assigned to the cluster nodes:

```
node show -cluster-name <cluster_name>
```

3. Select the node and perform basic configuration:

```
node modify -name <node_name> -cluster-name <cluster_name> -host-name
<FQDN|IP> -license-serial-number <number> -instance-type TYPE
-passthrough-disks false
```

Example output:

```
node modify -name test-cluster-01 -cluster-name test-cluster -host-name
10.234.81.14
-license-serial-number 320000nnnn -instance-type small -passthrough
-disks false
```

The RAID configuration for the node is indicated with the *passthrough-disks* parameter. If you are using a local hardware RAID controller, this value must be "false". If you are using software RAID, this value must be "true".

A Capacity Tier license is used for the ONTAP Select node.

4. Display the network configuration available at the host:

```
host network show -host-name <FQDN|IP> -detailed
```

Example output:

```
host network show -host-name 10.234.81.14 -detailed
```

5. Perform network configuration of the node:

ESXi host

```
node modify -name <node_name> -cluster-name <cluster_name> -mgmt-ip
IP -management-networks <network_name> -data-networks <network_name>
-internal-network <network_name>
```

KVM host

```
node modify -name <node_name> -cluster-name <cluster_name> -mgmt-ip
IP -management-vlans <vlan_id> -data-vlans <vlan_id> -internal-vlans
<vlad_id>
```

When deploying a single-node cluster, you do not need an internal network and should remove "-internal-network".

Example output:

```
node modify -name test-cluster-01 -cluster-name test-cluster -mgmt-ip
10.234.81.21
-management-networks sDOT_Network -data-networks sDOT_Network
```

6. Display the configuration of the node:

```
node show -name <node_name> -cluster-name <cluster_name> -detailed
```

Example output:

```
node show -name test-cluster-01 -cluster-name test-cluster -detailed
```

Step 6: Attach storage to the ONTAP Select nodes

Configure the storage used by each node in the ONTAP Select cluster. Every node must always be assigned at least one storage pool. When using software RAID, each node must also be assigned at least one disk drive.

Before you begin

Create the storage pool using VMware vSphere. If you are using software RAID, you also need at least one available disk drive.

About this task

When using a local hardware RAID controller, you need to perform steps 1 through 4. When using software RAID, you need to perform steps 1 through 6.

Steps

1. Sign in to the Deploy utility CLI using SSH with the administrator account credentials.
2. Display the storage pools available at the host:

```
host storage pool show -host-name <FQDN|IP>
```

Example output:

```
host storage pool show -host-name 10.234.81.14
```

You can also obtain the available storage pools through VMware vSphere.

3. Attach an available storage pool to the ONTAP Select node:

```
node storage pool attach -name <pool_name> -cluster-name <cluster_name>
-node-name <node_name> -capacity-limit <limit>
```

If you include the "-capacity-limit" parameter, specify the value as GB or TB.

Example output:

```
node storage pool attach -name sDOT-02 -cluster-name test-cluster -
node-name test-cluster-01 -capacity-limit 500GB
```

4. Display the storage pools attached to the node:

```
node storage pool show -cluster-name <cluster_name> -node-name
<node_name>
```

Example output:

```
node storage pool show -cluster-name test-cluster -node-name
testcluster-01
```

5. If you are using software RAID, attach the available drive or drives:

```
node storage disk attach -node-name <node_name> -cluster-name
<cluster_name> -disks <list_of_drives>
```

Example output:

```
node storage disk attach -node-name NVME_SN-01 -cluster-name NVME_SN
-disks 0000:66:00.0 0000:67:00.0 0000:68:00.0
```

6. If you are using software RAID, display the disks attached to the node:

```
node storage disk show -node-name <node_name> -cluster-name
```

```
<cluster_name>`
```

Example output:

```
node storage disk show -node-name sdot-smicro-009a -cluster-name NVME
```

Step 7: Deploy an ONTAP Select cluster

After the cluster and nodes have been configured, you can deploy the cluster.

Before you begin

Run the network connectivity checker using the web UI or the CLI to confirm the connectivity among the cluster nodes on the internal network.

Steps

1. Sign in to the Deploy utility CLI using SSH with the administrator account.
2. Deploy the ONTAP Select cluster:

   ```
   cluster deploy -name <cluster_name>
   ```

   Example output:

   ```
   cluster deploy -name test-cluster
   ```

   Provide the password to be used for the ONTAP administrator account when prompted.
3. Display the status of the cluster to determine when it has been successfully deployed successfully:

   ```
   cluster show -name <cluster_name>
   ```

What's next?

Back up the ONTAP Select Deploy configuration data.

Secure an ONTAP Select deployment

You can secure an ONTAP Select deployment by changing the password for the Deploy virtual machine administrator account as needed using the command line interface.

Steps

1. Sign in to the Deploy utility CLI using the administrator account.
2. Change the password:
   `password modify`

3. Respond to all prompts as appropriate for your environment.

Confirm network connectivity among ONTAP Select nodes

You can test the network connectivity among two or more ONTAP Select nodes on the internal cluster network. You typically run this test before a multi-node cluster is deployed to detect issues that might cause the operation to fail.

Before you begin

All the ONTAP Select nodes included in the test must be configured and powered on.

About this task

Each time you start a test, a new process run is created in the background and assigned a unique run identifier. Only one run can be active at a time.

The test has two modes that control its operation:

- Quick
  This mode performs a basic non-disruptive test. A PING test is performed, along with a test of the network MTU size and the vSwitch.
- Extended
  This mode performs a more comprehensive test over all the redundant network paths. If you run this on an active ONTAP Select cluster, the performance of the cluster can be impacted.



It is recommended that you always perform a quick test before creating a multi-node cluster. After the quick test completes successfully, you can optionally perform an extended test based on your production requirements.

Steps

1. Sign in to the Deploy utility CLI using the administrator account.
2. Display the current runs of the network connectivity checker and verify that no runs are active:

   ```
   network connectivity-check show
   ```

3. Start the network connectivity checker and note the run identifier in the command output:

   ```
   network connectivity-check start -host-names HOSTNAMES -vswitch-type
   VSWITCH_TYPE-mode MODE
   ```

   Example

   ```
   network connectivity-check start -host-names 10.234.81.14
   10.234.81.15 -vswitch-type StandardVSwitch -mode quick
   ```

4. Monitor the progress of the network connectivity checker based on the run identifier:

   ```
   network connectivity-check show -run-id RUN_ID
   ```

After you finish

The network connectivity checker normally cleans up by removing any temporary ports and IP addresses added to the ONTAP-Internal port group. However, if the connectivity checker fails to remove the temporary ports, you must perform a manual cleanup operation by rerunning the CLI command with the option `-mode cleanup`. If you do not remove the temporary ports from the ONTAP-Internal port group, the ONTAP Select virtual machine may not be created successfully.

Manage ONTAP Select clusters using the CLI

There are several related tasks you can perform to administer an ONTAP Select cluster using the CLI.

Back up the ONTAP Select Deploy configuration data

Backup of the ONTAP Select Deploy configuration data, for example, after deploying a cluster. The data is saved to a single encrypted file that you can download to your local workstation.

The backup file you create captures all the configuration data. This data describes aspects of your deployment environment, including the ONTAP Select clusters.

Before you begin

Make sure that Deploy is not performing any other tasks during the backup operation.

Steps

1. Sign in to the ONTAP Select Deploy utility CLI using SSH with the administrator account.
2. Create a backup of the ONTAP Select Deploy configuration data, which is stored in an internal directory at the ONTAP Select Deploy server:

   ```
   deploy backup create
   ```

3. Provide a password for the backup when prompted.

   The backup file is encrypted based on the password.

4. Display the available backups in the system:

   ```
   deploy backup show -detailed
   ```

5. Select your backup file based on the date in the **Created** field and record the **Download URL** value.

   You can access the backup file through the URL.

6. Using a web browser or utility such as Curl, download the backup file to your local workstation with the URL.

Delete an ONTAP Select cluster

You can delete an ONTAP Select cluster when it is no longer needed.

Before you begin

The cluster must be in the offline state.

Steps

1. Sign in to the Deploy virtual machine CLI using the administrator account.
2. Display the cluster status:

```
cluster show -name <cluster_name>
```

3. If the cluster is not offline, move it to an offline state:

```
cluster offline -name <cluster_name>
```

4. After confirming the cluster is in an offline status, delete the cluster:

```
cluster delete -name <cluster_name>
```

Nodes and hosts

Upgrade to VMware ESXi 8.0 or later for ONTAP Select

If you are running ONTAP Select on VMware ESXi, you can upgrade the ESXi software from an earlier supported version to ESXi 8.0 or later. Before upgrading, you should understand the process and select the appropriate upgrade procedure.

Prepare to upgrade VMware ESXi

Prepare and select the upgrade procedure that is appropriate for your environment before upgrading the ESXi software on the hypervisors hosting an ONTAP Select cluster.

Steps

1. **Become familiar with how to upgrade VMware ESXi**

   Upgrading the ESXi software is a process described and supported by VMware. The hypervisor upgrade process is part of the larger upgrade procedure when using ONTAP Select. Refer to the VMware documentation for more information.

2. **Select an upgrade procedure**

   Several upgrade procedures are available. You should select the applicable procedure based on the following criteria:

   - ONTAP Select cluster size

     Both single-node and multi-node clusters are supported.

   - Use of ONTAP Select Deploy

Upgrade is possible both with and without the Deploy utility.



You should select an upgrade procedure that uses the Deploy administration utility.

Performing an ESXi upgrade using the Deploy administration utility is the more general and resilient option. However, there might be instances when Deploy is unavailable or cannot be used. For example, upgrading to ESXi 8.0 is not supported with earlier versions of ONTAP Select and the Deploy administration utility.

If you are using these earlier versions and attempt an upgrade, you can leave the ONTAP Select virtual machine in a state where it cannot be booted. In this case, you must select an upgrade procedure that does not use Deploy.

3. **Upgrade the Deploy administration utility**

   Before performing an upgrade procedure using the Deploy utility, you might need to upgrade your Deploy instance. In general, you should upgrade to the most recent version of Deploy. The Deploy utility must support the version of ONTAP Select you are using. Refer to the ONTAP Select release notes for more information.

4. **After the update procedure is complete**

   If you select an upgrade procedure that uses the Deploy utility, you should perform a cluster refresh operation using Deploy after all of the nodes have been upgraded. See Refreshing the Deploy cluster configuration for more information.

Upgrade a single-node cluster using Deploy

You can use the Deploy administration utility as part of the procedure to upgrade the VMware ESXi hypervisor hosting an ONTAP Select single-node cluster.

Steps

1. Sign in to the Deploy utility CLI using SSH with the administrator account.
2. Move the node to the offline state:

```
node stop --cluster-name <cluster_name> --node-name <node_name>
```

3. Upgrade the hypervisor host where ONTAP Select is running to ESXi 8.0 or later using the procedure provided by VMware.
4. Move the node to the online state:

```
node start --cluster-name <cluster_name> --node-name <node_name>
```

5. After the node comes up, verify that the cluster is healthy.

   Example:

```
ESX-1N::> cluster show
Node           Health  Eligibility
```

```
-------------- ------- ------------
sdot-d200-011d true    true
```

After you finish

You should perform a cluster refresh operation using the Deploy administration utility.

Upgrade a multi-node cluster using Deploy

You can use the Deploy administration utility as part of the procedure to upgrade the VMware ESXi hypervisors hosting an ONTAP Select multi-node cluster.

About this task

You must perform this upgrade procedure for each of the nodes in the cluster, one node at a time. If the cluster contains four or more nodes, you should upgrade the nodes in each HA pair sequentially before proceeding to the next HA pair.

Steps

1. Sign in to the Deploy utility CLI using SSH with the administrator account.
2. Move the node to the offline state:

```
node stop --cluster-name <cluster_name> --node-name <node_name>
```

3. Upgrade the hypervisor host where ONTAP Select is running to ESXi 8.0 or later using the procedure provided by VMware.

   Refer to Preparing to upgrade VMware ESXi for more information.

4. Move the node to the online state:

```
node start --cluster-name <cluster_name> --node-name <node_name>
```

5. After the node comes up, verify that storage failover is enabled and the cluster is healthy.

   ▼ **Show example**

```
ESX-2N_I2_N11N12::> storage failover show
Takeover
Node Partner Possible State Description
-------------- -------------- -------- -------------------------
sdot-d200-011d sdot-d200-012d true Connected to sdot-d200-012d
sdot-d200-012d sdot-d200-011d true Connected to sdot-d200-011d
2 entries were displayed.
ESX-2N_I2_N11N12::> cluster show
Node Health Eligibility
--------------------- ------- ------------
sdot-d200-011d true true
sdot-d200-012d true true
```

```
2 entries were displayed.
```

After you finish

You must perform the upgrade procedure for each host used in the ONTAP Select cluster. After all the ESXi hosts are upgraded, you should perform a cluster refresh operation using the Deploy administration utility.

Upgrade a single-node cluster without Deploy

You can upgrade the VMware ESXi hypervisor hosting an ONTAP Select single-node cluster without using the Deploy administration utility.

Steps

1. Sign in to the ONTAP command line interface and halt the node.
2. Using VMware vSphere, confirm that the ONTAP Select virtual machine is powered off.
3. Upgrade the hypervisor host where ONTAP Select is running to ESXi 8.0 or later using the procedure provided by VMware.

   Refer to Preparing to upgrade VMware ESXi for more information.

4. Using VMware vSphere, access vCenter and do the following:
   a. Add a floppy drive to the ONTAP Select virtual machine.
   b. Power on the ONTAP Select virtual machine.
   c. Sign in to the ONTAP CLI using SSH with the administrator account.
5. After the node comes up, verify that the cluster is healthy.

   Example:

```
ESX-1N::> cluster show
Node           Health  Eligibility
-------------- ------- ------------
sdot-d200-011d true    true
```

After you finish

You should perform a cluster refresh operation using the Deploy administration utility.

Upgrade a multi-node cluster without Deploy

You can upgrade the VMware ESXi hypervisors hosting an ONTAP Select multi-node cluster without using the Deploy administration utility.

About this task

You must perform this upgrade procedure for each of the nodes in the cluster, one node at a time. If the cluster contains four or more nodes, you should upgrade the nodes in each HA pair sequentially before proceeding to the next HA pair.

Steps

1. Sign in to the ONTAP command line interface and halt the node.

2. Using VMware vSphere, confirm that the ONTAP Select virtual machine is powered off.
3. Upgrade the hypervisor host where ONTAP Select is running to ESXi 8.0 or later using the procedure provided by VMware.
4. Using VMware vSphere, access vCenter and do the following:
   a. Add a floppy drive to the ONTAP Select virtual machine.
   b. Power on the ONTAP Select virtual machine.
   c. Sign in to the ONTAP CLI using SSH with the administrator account.
5. After the node comes up, verify that storage failover is enabled and the cluster is healthy.

▼ Show example

```
ESX-2N_I2_N11N12::> storage failover show
Takeover
Node Partner Possible State Description
-------------- -------------- -------- -------------------------
sdot-d200-011d sdot-d200-012d true Connected to sdot-d200-012d
sdot-d200-012d sdot-d200-011d true Connected to sdot-d200-011d
2 entries were displayed.
ESX-2N_I2_N11N12::> cluster show
Node Health Eligibility
--------------------- ------- ------------
sdot-d200-011d true true
sdot-d200-012d true true
2 entries were displayed.
```

After you finish

You must perform the upgrade procedure for each host used in the ONTAP Select cluster.

Modify a host management server for ONTAP Select Deploy

You can use the `host modify` command to modify a host management server with this instance of ONTAP Select Deploy.

Syntax

`host modify [-help] [-foreground] -name` *`name`* `-mgmt-server` *`management_server`* `[-username` *`username`*`]`

Required parameters

Parameter	Description
`-name` *`name`*	The IP address or FQDN of the host you want to modify.

`-mgmt-server` *`management_server`*	The IP address or FQDN of the host management server to be set to the host. Specify "-" (hyphen) to unset the management server from the host. The credentials for this management server must be added prior to registering this host using the `credential add` command.

Optional parameters

Parameter	Description
`-help`	Displays the help message.
`-foreground`	This parameter controls the behavior of long-running commands. If set, the command will run in the foreground and event messages related to the operation will be displayed as they occur.
`-username` *`username`*	The username that has access to this host. This is required only if the host is not managed by a management server (that is, an ESXi host managed by a vCenter).

Deploy utility

Upgrade an ONTAP Select Deploy instance

Upgrade an existing ONTAP Select Deploy utility virtual machine in-place using the ONTAP Select Deploy utility CLI.

You can upgrade directly to ONTAP Select Deploy 9.19.1 from ONTAP Select Deploy 9.18.1 or 9.17.1. To upgrade from an earlier release, for example, ONTAP Select Deploy 9.16.1, you must first upgrade to ONTAP Select Deploy 9.17.1 or 9.18.1, then upgrade to ONTAP Select Deploy 9.19.1.



If you have an older instance of the ONTAP Select Deploy administration utility installed, you should upgrade to the current release. You upgrade the ONTAP Select node and ONTAP Select Deploy components independently. Refer to Upgrade the ONTAP Select nodes for further details.

Before you begin

Verify that ONTAP Select Deploy is not used to perform any other tasks during the upgrade.

Step 1: Download the upgrade package

To begin the upgrade process, download the ONTAP Select Deploy Upgrade file from the NetApp Support Site. The upgrade package is formatted as a single compressed file.

Steps

1. Access the NetApp Support Site downloads page.
2. Scroll down and select **ONTAP Select Deploy**.
3. Select the desired ONTAP Select version.
4. Review the End User License Agreement (EULA) and select **Accept & Continue**.
5. Select and download the appropriate **ONTAP Select Deploy Upgrade** package. Respond to all prompts as needed.

Step 2: Upload the package to the ONTAP Select Deploy virtual machine

After you download the package, you need to upload the file to the ONTAP Select Deploy virtual machine.

About this task

This task describes one method for uploading the file to the ONTAP Select Deploy virtual machine. There might be other options more suitable for your environment.

Before you begin

- Verify that the upgrade file is available on your local workstation.
- Verify that you have the password for the administrator user account.

Step

In a command shell on your local workstation, use the `scp` (Secure Copy Protocol) utility to upload the image file to the ONTAP Select Deploy virtual machine, as shown in the following example:

```
scp ONTAPdeploy2.12_upgrade.tar.gz admin@10.228.162.221:/home/admin
(provide password when prompted)
```

Result

The upgrade file is stored in the home directory of the admin user.

Step 3: Apply the upgrade package

After you upload the upgrade file to the ONTAP Select Deploy virtual machine, you can apply the upgrade.

Before you begin

- Verify which directory the upgrade file has been placed in at the ONTAP Select Deploy utility virtual machine.
- Verify that ONTAP Select Deploy is not used to perform any other tasks while the upgrade is performed.

Steps

1. Sign in to the ONTAP Select Deploy utility CLI using SSH with the administrator account.
2. Perform the upgrade using the appropriate directory path and file name:

   ```
   deploy upgrade -package-path <file_path>
   ```

   Example command:

   ```
   deploy upgrade -package-path /home/admin/ONTAPdeploy2.12_upgrade.tar.gz
   ```

After you finish

Before the upgrade procedure completes, you are asked to create a backup of the ONTAP Select Deploy virtual machine configuration. Also, you should clear the browser cache so you can view the newly created ONTAP Select Deploy pages.

Migrate an ONTAP Select Deploy instance to a new virtual machine

You can migrate an existing instance of the Deploy administration utility to a new virtual machine using the command line interface.

This procedure is based on creating a new virtual machine that uses the configuration data from the original virtual machine. The new and original virtual machines must run the same version and release of the Deploy utility. You cannot migrate to a different version and release of the Deploy utility.

Step 1:Back up the Deploy configuration data

You must create a backup of the Deploy configuration data as part of migrating the virtual machine. You should also create a backup after deploying an ONTAP Select cluster. The data is saved to a single encrypted file that you can download to your local workstation.

Before you begin

- Make sure that Deploy is not performing any other tasks during the backup operation.
- Save the original Deploy virtual machine image.

> The original Deploy virtual machine image is required later in this procedure when you restore the Deploy configuration data from the original to the new virtual machine.

About this task

The backup file you create captures all the configuration data from the virtual machine. This data describes aspects of your deployment environment, including the ONTAP Select clusters.

Steps

1. Sign in to the Deploy utility CLI using SSH with the administrator account.
2. Create a backup of the Deploy configuration data, which is stored in an internal directory at the Deploy server:

   ```
   deploy backup create
   ```

3. Provide a password for the backup when prompted.

   The backup file is encrypted based on the password.

4. Display the available backups in the system:

   ```
   deploy backup show -detailed
   ```

5. Select your backup file based on the date in the **Created** field and record the **Download URL** value.

   You can access the backup file through the URL.

6. Using a web browser or utility such as Curl, download the backup file to your local workstation with the URL.

Step 2: Install a new instance of the Deploy virtual machine

You must create a new instance of the Deploy virtual machine which you can update with the configuration data from the original virtual machine.

Before you begin

You must be familiar with the procedures used to download and deploy the ONTAP Select Deploy virtual machine in a VMware environment.

About this task

This task is described at a high level.

Steps

1. Create a new instance of the Deploy virtual machine:
   a. Download the virtual machine image.
   b. Deploy the virtual machine and configure the network interface.
   c. Access the Deploy utility using SSH.

Related information

Install ONTAP Select Deploy

Step 3: Restore the Deploy configuration data to the new virtual machine

You must restore the configuration data from the original Deploy utility virtual machine to the new virtual machine. The data is in a single file that you must upload from your local workstation.

Before you begin

You must have the configuration data from a previous backup. The data is contained in a single file and must be available on your local workstation.

Steps

1. In a command shell on your local workstation, use the sftp utility to upload the backup file to the Deploy virtual machine, as shown in the following example:

```
sftp admin@10.234.81.101 (provide password when prompted)
put deploy_backup_20190601162151.tar.gz
exit
```

2. Sign in to the Deploy utility CLI using SSH with the administrator account.
3. Restore the configuration data:

```
deploy backup restore -path <path_name> -filename <file_name>
```

   Example command:

```
deploy backup restore -path /home/admin -filename
deploy_backup_20180601162151.tar.gz
```

Add an ONTAP Select image to Deploy

Add an ONTAP Select image to your instance of the Deploy administration utility. After the image has been installed, you can use it when deploying an ONTAP Select cluster.

Before you begin

Before adding any new ONTAP Select images to Deploy, you should first remove any unneeded images.

> ⓘ You should only add an ONTAP Select image with a version that is earlier than the original version included with your instance of the Deploy utility. Adding later versions of ONTAP Select as they become available from NetApp is not a supported configuration.

Step 1: Download the install image

To begin the process of adding an ONTAP Select image to an instance of the Deploy utility, you must download the install image from the NetApp Support Site. The ONTAP Select install image is formatted as a single compressed file.

Steps

1. Access the NetApp Support Site downloads page.
2. Scroll down and select **ONTAP Select Image**.
3. Select the desired release of the install image.
4. Review the End User License Agreement (EULA) and select **Accept & Continue**.
5. Select and download the appropriate **ONTAP Select Image Install** package. Respond to all prompts as needed.

Step 2: Upload the install image to Deploy

After acquiring the ONTAP Select install image, you must upload the file to the Deploy virtual machine.

Before you begin

Verify that you have the install image file available on your local workstation. You must also have the password for the Deploy administrator user account.

About this task

This task describes one method for uploading the file to the Deploy virtual machine. There may be other options more suitable for your environment.

Step

1. In a command shell on your local workstation, upload the image file to the Deploy virtual machine, as shown in the following examples:

```
scp image_v_93_install_esx.tgz admin@10.234.81.101:/home/admin (provide
password when prompted)
```

```
sftp admin@10.234.81.101 (provide password when prompted)
put image_v_93_install_esx.tgz
```

```
exit
```

Result

The node install file is stored in the home directory of the admin user.

Step 3: Add the install image

Add the ONTAP Select installation image to the Deploy images directory so it is available when deploying a new cluster.

Before you begin

You must know which directory the install image file has been placed in at the Deploy utility virtual machine. It is assumed the file is in the administrator's home directory.

Steps

1. Sign in to the Deploy utility CLI using SSH with the administrator (admin) account.
2. Start the Bash shell:

```
shell bash
```

3. Place the install image file into the images directory, as shown in the following example:

```
tar -xf image_v_93_install_esx.tgz -C /opt/netapp/images/
```

Step 4: Display the available install images

Display the ONTAP Select images that are available when deploying a new cluster.

Steps

1. Access the online documentation web page at the Deploy utility virtual machine and sign in using the administrator (admin) account:

   ```
   http://<FQDN|IP_ADDRESS>/api/ui
   ```

   Use the domain name or IP address of the Deploy virtual machine.

2. Navigate to the bottom of the page and select **Deploy** and then select **GET /images**.
3. Select **Try it out!** to display the available ONTAP Select images.
4. Confirm that the desired image is available.

Remove an ONTAP Select image from Deploy

You can remove ONTAP Select images from your instance of the Deploy administration utility when they are no longer needed.



You should not remove any ONTAP Select images that are in use by a cluster.

About this task

You can remove older ONTAP Select images that are not currently in use by a cluster or planned for use with a future cluster deployment.

Steps

1. Sign in to the Deploy utility CLI using SSH with the administrator (admin) account.
2. Display the clusters managed by Deploy and record the ONTAP images in use:

   ```
   cluster show
   ```

   Note the version number and hypervisor platform in each case.

3. Start the Bash shell:

   ```
   shell bash
   ```

4. Display all of the available ONTAP Select images:

   ```
   ls -lh /opt/netapp/images
   ```

5. Optionally, remove the ONTAP Select image with your hypervisor host.

   ESXi example

   ```
   rm -r /opt/netapp/images/DataONTAPv-9.3RC1-vidconsole-esx.ova
   ```

   KVM example

   ```
   rm -r /opt/netapp/images/DataONTAPv-9.3RC1-serialconsole-kvm.raw.tar
   ```

Recover the ONTAP Select Deploy utility for a two-node cluster

If the ONTAP Select Deploy utility fails or becomes unavailable for some reason, you lose the ability to administer ONTAP Select nodes and clusters. In addition, all two-node clusters lose HA capability because the mediator service included with Deploy is unavailable. If an unrecoverable failure occurs, you must recover the Deploy utility instance to restore administrative and HA functionality.

Prepare to recover the Deploy utility

You need to prepare before attempting to recover an instance of the Deploy utility to ensure success. You must be familiar with several administrative procedures and have the required information.

Steps

1. Verify that you can install a new instance of the ONTAP Select Deploy utility in your hypervisor environment.

   Learn about installing the ONTAP Select Deploy utility

2. Verify that you can log in to the ONTAP Select cluster and access the ONTAP cluster shell (CLI).
3. Determine if you have a backup of the configuration data from the failed Deploy utility instance that contains the ONTAP Select two-node cluster. You might have a backup that does not contain the cluster.
4. Verify that you can restore a backup of the Deploy configuration data, depending on the recovery procedure used.

   Learn about restoring the Deploy configuration data to the new virtual machine

5. You have the IP address of the original Deploy utility virtual machine that failed.
6. Determine whether Capacity Pools or Capacity Tiers licensing is used. If you use Capacity Pools licensing, you must reinstall each Capacity Pool license after recovering or restoring the Deploy instance.
7. Decide which procedure to use when recovering an instance of the ONTAP Select Deploy utility. Your decision is based on whether or not you have a backup of the configuration data from the original failed Deploy utility that contains the ONTAP Select two-node cluster.

Do you have a Deploy backup containing the two-node cluster?	Use the recovery procedure…
Yes	Restore a Deploy utility instance using a configuration backup
No	Reconfigure and recover a Deploy utility instance

Restore a Deploy utility instance using a configuration backup

If you have a backup of the failed Deploy utility instance containing the two-node cluster, you can restore the configuration data to the new Deploy virtual machine instance. You must then complete the recovery by performing additional configuration of the two nodes in the ONTAP Select cluster.

Before you begin

Backup the configuration data from the original failed Deploy virtual machine that contains the two-node cluster. You must be able to sign in to the ONTAP CLI of the two-node cluster and know the ONTAP names of the two nodes.

About this task

Because the configuration backup you restore contains the two-node cluster, the mediator iSCSI targets and mailboxes are recreated in the new Deploy utility virtual machine.

Steps

1. Prepare a new instance of the ONTAP Select Deploy utility:
   a. Install a new Deploy utility virtual machine.
   b. Restore the Deploy configuration from a previous backup to the new virtual machine.

      Refer to the related tasks for more detailed information about the install and restore procedures.

2. Sign in to the ONTAP command line interface of the ONTAP Select two-node cluster.
3. Enter advanced privilege mode:

```
set adv
```

4. If the IP address of the new Deploy virtual machine is different than the original Deploy virtual machine, remove the old mediator iSCSI targets and add new targets:

```
storage iscsi-initiator remove-target -node * -target-type mailbox
```

```
storage iscsi-initiator add-target -node <node1_name> -label mediator
-target-type mailbox -target-portal <ip_address> -target-name <target>
```

```
storage iscsi-initiator add-target -node <node2_name> -label mediator
-target-type mailbox -target-portal <ip_address> -target-name <target>
```

   The `<ip_address>` parameter is the IP address of the new Deploy virtual machine.

   These commands allow the ONTAP Select nodes to discover the mailbox disks on the new Deploy utility virtual machine.

5. Determine the names of the mediator disks:

```
disk show -container-type mediator
```

6. Assign the mailbox disks to the two nodes:

```
disk assign -disk <mediator-disk1-name> -owner <node1-name>

disk assign -disk <mediator-disk2-name> -owner <node2-name>
```

7. Verify that storage failover is enabled:

```
storage failover show
```

After you finish

If you use Capacity Pools licensing, reinstall each Capacity Pool license. See Reinstall a Capacity Pool license for additional details.

Reconfigure and recover a Deploy utility instance

If you do not have a backup of the failed Deploy utility instance containing the two-node cluster, configure the mediator iSCSI target and mailbox in the new Deploy virtual machine. You then complete the recovery by performing additional configuration of the two nodes in the ONTAP Select cluster.

Before you begin

Verify that you have the name of the mediator target for the new Deploy utility instance. You must be able to sign in to the ONTAP CLI of the two-node cluster and know the ONTAP names of the two nodes.

About this task

You can optionally restore a configuration backup to the new Deploy virtual machine even though it does not contain the two-node cluster. Because the two-node cluster is not recreated with the restore, you must manually add the mediator iSCSI target and mailbox to the new Deploy utility instance through the ONTAP Select online documentation web page at the Deploy. You must be able to sign in to the two-node cluster and know the ONTAP names of the two nodes.



The goal of the recovery procedure is to restore the two-node cluster to a healthy state, where normal HA takeover and giveback operations can be performed.

Steps

1. Prepare a new instance of the ONTAP Select Deploy utility:
   a. Install a new Deploy utility virtual machine.
   b. Optionally, restore the Deploy configuration from a previous backup to the new virtual machine.

      If you restore a previous backup, the new Deploy instance will not contain the two-node cluster. Refer to the related information section for more detailed information about the install and restore procedures.
2. Sign in to the ONTAP command line interface of the ONTAP Select two-node cluster.
3. Enter advanced privileged mode:

   ```
   set adv
   ```

4. Get the mediator iSCSI target name:

   ```
   storage iscsi-initiator show -target-type mailbox
   ```

5. Access the online documentation web page at the new Deploy utility virtual machine and sign in using the admin account:

   `http://<ip_address>/api/ui`

   You must use the IP address of your Deploy virtual machine.
6. Select **Mediator** and then **GET /mediators**.
7. Select **Try it out!** to display a list of mediators maintained by Deploy.

   Note the ID of the desired mediator instance.
8. Select **Mediator** and then **POST**.
9. Provide the value for mediator_id.
10. Select the **Model** next to `iscsi_target` and complete the name value.

Use the target name for the iqn_name parameter.

11. Select **Try it out!** to create the mediator iSCSI target.

    If the request is successful, you will receive HTTP status code 200.

12. If the IP address of the new Deploy virtual machine is different than the original Deploy virtual machine, you must use the ONTAP CLI to remove the old mediator iSCSI targets and add new targets:

    ```
    storage iscsi-initiator remove-target -node * -target-type mailbox
    ```

    ```
    storage iscsi-initiator add-target -node <node1_name> -label mediator
    -target-type mailbox -target-portal <ip_address> -target-name <target>
    ```

    ```
    storage iscsi-initiator add-target -node <node2_name> -label mediator-
    target-type mailbox -target-portal <ip_address> -target-name <target>
    ```

    The `<ip_address>` parameter is the IP address of the new Deploy virtual machine.

    These commands allow the ONTAP Select nodes to discover the mailbox disks on the new Deploy utility virtual machine.

13. Determine the names of the mediator disks:

    ```
    disk show -container-type mediator
    ```

14. Assign the mailbox disks to the two nodes:

    ```
    disk assign -disk <mediator-disk1-name> -owner <node1-name>

    disk assign -disk <mediator-disk2-name> -owner <node2-name>
    ```

15. Verify that storage failover is enabled:

    ```
    storage failover show
    ```

After you finish

If you use Capacity Pools licensing, reinstall each Capacity Pool license. See Reinstall a Capacity Pool license for additional details.

Try ONTAP Select

Access ONTAP Select evaluation software

You can access the ONTAP Select evaluation software if you are a current customer of ONTAP Select and you have an active entitlement or a NetApp Partner and want to download a new Capacity Tiers evaluation license.

If you are neither of the above and would like to try ONTAP Select for an evaluation, you can follow these steps to request access to download the evaluation software.

About this task

- This instruction is for NetApp sales, partners, and prospects engaged with NetApp.
- ONTAP Select Capacity Tiers evaluation licenses are limited to 90 days from the day you install them.

> You cannot extend an evaluation license. At the end of the 90 days, the license no longer functions and you cannot access the associated data.

- The evaluation software has all the information that you need to download the ONTAP Select Deploy administration utility and deploy a cluster.

Step 1: Register for an account

Complete the NetApp account registration, if needed, and request prospect access.

Steps

1. Create a new account or sign out of an existing account.

   Create new account

   Register for an account on the NetApp Support site using your business email.

   > Free email accounts such as gmail or yahoo are not acceptable.

   Make sure you are not signed in with the newly created account.

   Existing account

   If you already have an account, go to the NetApp Support site and make sure that you are not signed in.

2. Create a non-technical support case to elevate access levels to "prospect."
   a. Select the **CASES** tab, then select **Create non-technical cases**
   b. Select **User Registration Issue** as the feedback category.
   c. Select **Create case**.

      Fill out all required information and in the comments section, write: "My business email address is <your-email-address>. I would like to get prospect access to download the ONTAP Select evaluation software."

Include the name of your NetApp contact.

d. Select **Submit**.

Step 2: Download ONTAP Select evaluation software

After your support case is reviewed and approved, NetApp support notifies you using the email address supplied.

Step

1. Download the ONTAP Select evaluation software.

   The evaluation license file is located within the zip file.

2. To comply with legal requirements, you need to open a support case before gaining access. This helps NetApp maintain the correct control and documentation while providing you with the production-ready software you need.

> NetApp provides the "production-ready" version of ONTAP Select, which is not an open-source or alternative version. Support is only available if you upgrade to a production license.

What's next?

After the evaluation software download completes, follow the instructions to download the ONTAP Select Deploy administration utility and deploy a cluster.

Related information

- Learn about evaluation licenses for ONTAP Select deployments
- Convert an evaluation license to a production license

Deploy a 90-day evaluation instance of an ONTAP Select cluster

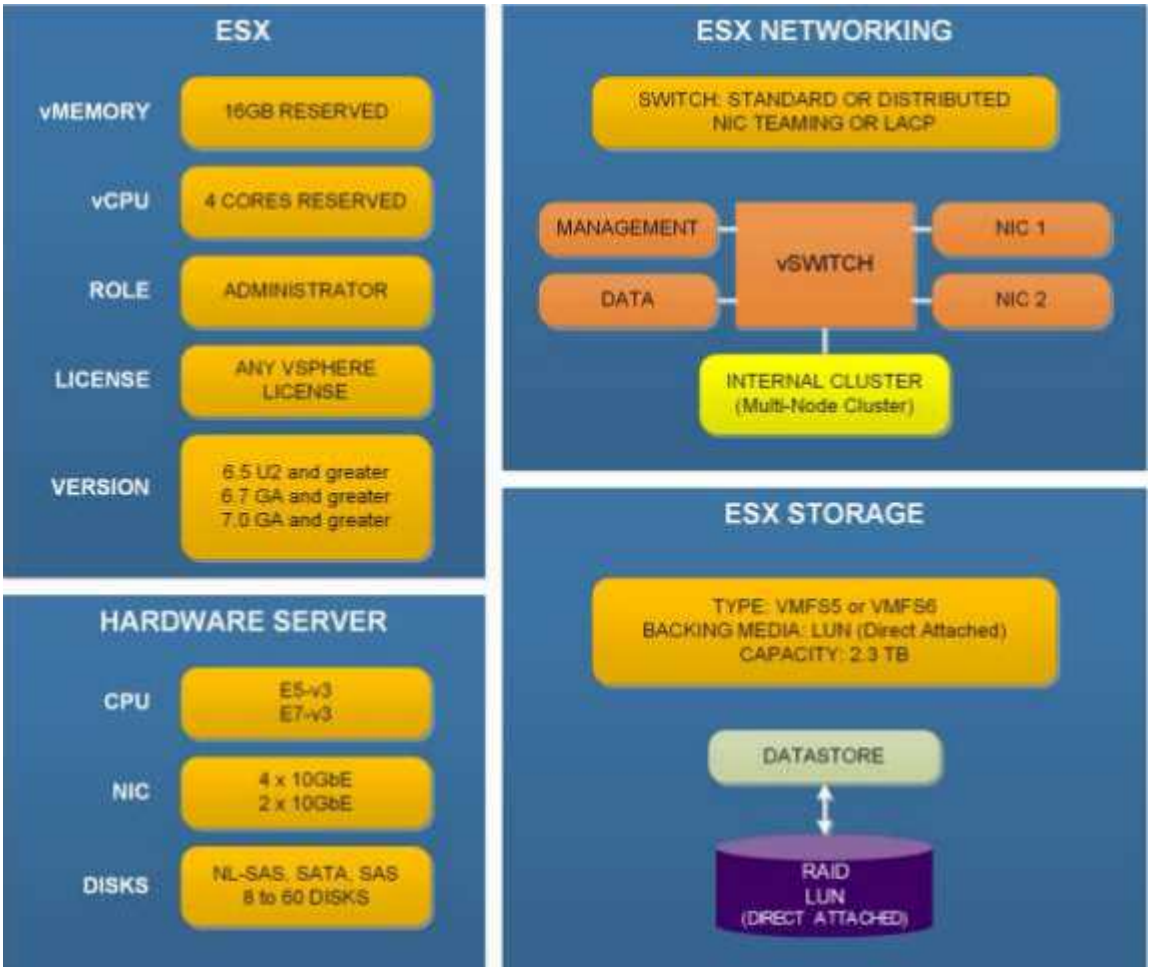
You can deploy a 90-day evaluation instance of an ONTAP select cluster ranging in size from a single-node up to 99 nodes using an OVF template on VMware ESXi. You can download an evaluation license if you are an existing or new customer. For example, you are an existing customer and you want to try a new configuration before purchasing a license.

About this task

- You don't need to obtain a serial number or storage capacity license from NetApp.
- You can allocate the same amount of storage for user data as a purchased license.
- You can only use a vCenter account to deploy the OVF template. Installing directly to an ESXi host is currently not supported.
- You must install the OVF template (contained in an ova file) using the vSphere standalone client or vSphere web client (the only option for ESXi 8.0 and later). You don't use the ONTAP Select Deploy administration utility.

Step 1: Prepare the ONTAP Select cluster host

Use the following requirements to prepare the ESXi host where the ONTAP Select cluster is deployed. The platform description is based on the standard or small instance type configuration with local direct-attached storage (DAS) formatted using the VMFS-5 or VMFS-6 file system.



For additional information about host configuration options, refer to the ONTAP Select Install documentation.

Step 2: Deploy a single-node ONTAP Select cluster using an OVF template

Complete the following steps on a supported 64-bit ESXi host server.

Steps

1. Go to the NetApp Evaluation Products Program page and select **ONTAP Select** to download the ONTAP Select OVF template to your local workstation.
2. Sign in to the VMware vSphere web client using an account with administrative privileges.
3. Select the host using one of the following options:
   - Select **File** > **Deploy OVF Template**.
   - Select **Datacenter**. Then right-select and select **Deploy OVF Template**.
4. Select the ONTAP Select source OVA file on your local workstation, then select **Next**.
5. Review the OVF template details and select **Next**.

6. Review the EULA details and select **Accept**. Then select **Next**.
7. Type the name of the new virtual machine and select **Next**.
8. If more than one datastore is available, select the datastore and select **Next**.
9. Select **Thick Provision Lazy Zeroed**, then select **Next**.
10. Select the data and management networks, then select **Next**.
11. On the **Properties** page, provide all the required values and select **Next**.
12. Review the deployment details and select **Power on after deployment**.
13. Select **Finish** to begin the deployment process.
14. After the ONTAP Select cluster has been deployed, you can configure the cluster using the System Manager or the CLI interface. You must assign the disk using the standard ONTAP **disk assign** operation.

Related information

- Learn about evaluation licenses for ONTAP Select deployments
- Interoperability Matrix Tool

FAQ for ONTAP Select

You can find answers to frequently asked questions about ONTAP Select.

> Beginning with ONTAP Select 9.14.1, support for KVM hypervisor has been reinstated. Previously, support for deploying a new cluster on a KVM hypervisor was removed in ONTAP Select 9.10.1 and support for managing existing KVM clusters and hosts, except to take offline or delete, was removed in ONTAP Select 9.11.1.

General

There are several general questions and answers.

What is the difference between ONTAP Select Deploy and ONTAP Select?

ONTAP Select Deploy is the utility used to create ONTAP Select clusters. Currently ONTAP Select Deploy is the only method available for creating a production cluster. ONTAP Select Deploy can also be used to create an evaluation Select cluster to allow clients to test and document the actual steps of a production deployment. ONTAP Select Deploy can also convert an evaluation cluster to a production cluster using an appropriate Capacity Tier license with sufficient capacity to cover the space consumed during the evaluation.

ONTAP Select Deploy is a virtual machine that contains an image of ONTAP Select. During cluster installation, ONTAP Select Deploy enforces several checks to help make sure that the ONTAP Select minimum requirements are met. The ONTAP Select Deploy VM and Select clusters can be upgraded separately.

How can I troubleshoot a performance issue with ONTAP Select?

Just like ONTAP on FAS, performance data should be collected using the perfstat utility. Here is a sample command:

```
perfstat8 -i N,m -t <sample time in minutes> --verbose --nodes=<filer IP>
--diag-passwd=abcxyz --mode="cluster-mode" > <name of output file>
```

How do I access the Swagger API page for ONTAP Select Deploy?

```
http://<Deploy-IP-Address/api/ui
```

> The API v3 release is not backward compatible with the prior version of the API. A new API procedure is available on the Field Portal.

Can the ONTAP Select VM be backed up with VMware or other third-party snapshots?

No. The ONTAP Select VM uses independent-persistent drives, which are excluded from VMware-based snapshots. The only supported method for backing up ONTAP Select is SnapMirror or SnapVault.

Where do I get clarification for questions not covered in this FAQ?

Contact xref:./ng-ses-ontap-select@netapp.com.

Licenses, installation, upgrades, and reverts

There are several questions and answers dealing with licenses, installation, upgrades, and reverts.

Can ONTAP Select and ONTAP Select Deploy be upgraded separately?

Yes. The ONTAP Select Deploy utility can be upgraded separately from the ONTAP Select cluster. Similarly, the Select cluster can be upgraded separately from the ONTAP Select Deploy utility.

Can ONTAP Select be upgraded using the same procedure as a FAS cluster?

Yes, the upgrade procedure for a Select cluster is identical to the upgrade of a FAS cluster, although the ONTAP Select upgrade binary is a separate download from the ONTAP on FAS upgrade binary.

Can ONTAP Select be reverted using the same procedure as a FAS cluster?

Yes, the revert procedure for an ONTAP Select cluster is almost identical to the revert procedure for a FAS cluster. There are a few differences however:

- Only upgraded instances of ONTAP Select can be reverted, and only up to the original install version. New installs cannot be reverted to an older code release, even if ONTAP Select in general does support that older release.
- For ONTAP Select (KVM) and ONTAP Select (ESXi) using software RAID, it is not possible to revert to a prior version that does not support software RAID. Furthermore, a new installation of ONTAP Select 9.5 or later on ESXi uses VMXNET3 network drivers and, when possible, the vNVMe driver. These new installations cannot be reverted to prior versions of ONTAP Select.
- If the ONTAP Select VM was also upgraded to a Large instance (using the Premium XL license), then reverting to a prior version before 9.6 is not supported, because the Large instance feature is not available in earlier versions.

Does the ONTAP MetroCluster SDS require at a minimum a Premium license?

Yes.

Can the ONTAP Select cluster network configuration be changed after installation?

Changes to the following ONTAP Select cluster properties are recognized by ONTAP Select Deploy using the cluster refresh operation available through the UI, CLI, or REST API:

- Network configuration (IP addresses, DNS, NTP, netmask, and gateway)
- ONTAP Select cluster, node name, and version

The following ONTAP Select VM changes are also recognized:

- ONTAP Select VM name and state changes (for example, online or offline)
- Host network name and storage pool name changes

Upgrading to ONTAP Select Deploy 2.6 enables support for these changes for any ONTAP Select cluster that is already deployed but has not been changed from its original configuration. In other words, if the ONTAP Select cluster properties mentioned above were changed using System Manager or vCenter, then upgrading to ONTAP Select Deploy 2.6 will not fix these inconsistencies. The ONTAP Select property changes must be first rolled back for ONTAP Select Deploy to add its unique metadata to each ONTAP Select VM.

Can the ONTAP Select Deploy network configuration be changed after installation?

Modifying the networking details of the Deploy instance after it's running in an environment is not supported. For more information, refer to the Knowledge Base article Modifying DNS configuration of ONTAP Deploy

instance.

How does Deploy detect that ONTAP Select licenses are renewed?

The method is the same for all licenses, although the specifics vary depending on whether it is a Capacity Tier or Capacity Pool license.

- ONTAP Select Deploy detects if licenses and support contracts are renewed with the purchase of an updated license file from NetApp. The license file (.NLF) includes capacity, start, and end dates; and is generated on the NetApp Support site, and then updated on the Deploy server.

> You can load the NLF into the Deploy server by using the **Add** and **Update** functions. **Add** adds new licenses to the server, and **Update** updates existing files with information such as capacity, node license (standard, premium, premium XL), support start and end dates (Capacity Tier license), or subscription start and end dates (Capacity Pool license).

> Do not attempt to modify the license file. Doing so invalidates the security key and renders the license invalid.

- A **Capacity Tier license** is a per-node permanent license tied to the ONTAP Select node serial number. It is sold with a separate support contract. While the license is permanent, the support contract must be renewed to access ONTAP Select upgrades and to receive assistance from NetApp technical support. A current support contract is also required to change license parameters, such as capacity or node size.

  Purchasing a Capacity Tier license update, parameter change, or support contract renewal, requires the node serial number as part of the order. Capacity Tier node serial numbers are nine digits long, and begin with the number '32'.

  Once the purchase is complete, and the license file generated, it's uploaded to the Deploy server using the **Update** function.

- A **Capacity Pool license** is a subscription for the right to use a specific pool of capacity and node size (standard, premium, premium XL) to deploy one or more clusters. The subscription includes the right to use a license and support for a specified term. The right to use a license and the support contract have specified start and end dates.

How does Deploy detect if the nodes have renewed licenses or support contract?

Purchasing, generating, and uploading an updated license file is how Deploy detects renewed licenses and support contracts.

If a Capacity Tier support contract end date has passed, the node can keep running, but you won't be able to download and install ONTAP updates, or call NetApp technical support for assistance without first bringing the support contract up-to-date.

If a Capacity Pool subscription lapses, the system warns you first but after 30 days, if the system shuts down, it won't reboot until an updated subscription is installed on the Deploy server.

Storage

There are several questions and answers dealing with storage.

Can a single ONTAP Select Deploy instance create clusters on both ESXi and KVM?

Yes. ONTAP Select Deploy can be installed on either KVM or ESXi, and both installations can create ONTAP Select clusters on either hypervisor.

Is vCenter required for ONTAP Select on ESXi?

If the ESXi hosts are properly licensed, then there is no need for the ESXi hosts to be managed by a vCenter Server. However, if the hosts are managed by a vCenter server, then you must configure ONTAP Select Deploy to use that vCenter Server. In other words, you cannot configure ESXi hosts as standalone in ONTAP Select Deploy if they are being actively managed by a vCenter Server. Note that the ONTAP Select Deploy VM relies on vCenter to track all ONTAP Select VM migrations between ESXi hosts due to a vMotion or VMware HA event.

What is Software RAID?

ONTAP Select can use servers without a hardware RAID controller. In this case, the RAID functionality is implemented in software. When using software RAID, both SSD and NVMe drives are supported. The ONTAP Select boot and core disks must still reside inside a virtualized partition (storage pool or datastore). ONTAP Select uses RD2 (root-data-data partitioning) to partition the SSDs. Therefore, the ONTAP Select root partition resides on the same physical spindles that are used for the data aggregates. However, the root aggregate and the boot and core virtualized disks do not count against the capacity license.

All RAID methods available on AFF/FAS are also available to ONTAP Select. This includes RAID 4, RAID DP, and RAID-TEC. The minimum number of SSDs varies depending on the type of RAID configuration chosen. Best practices require the presence of at least one spare. The spare and parity disks do not count toward the capacity license.

How is software RAID different from a hardware RAID configuration?

Software RAID is a layer in the ONTAP software stack. Software RAID provides more administrative control because the physical drives are partitioned and available as raw disks within the ONTAP Select VM. Whereas, with hardware RAID, a single large LUN is usually available that can then be carved out to create VMDISKs seen within ONTAP Select. Software RAID is available as an option and can be used instead of hardware RAID.

Some of the requirements for software RAID are as follows:

- Supported for ESXi and KVM
  - Beginning with ONTAP Select 9.14.1, support for KVM hypervisor has been reinstated. Previously, support for KVM hypervisor was removed in ONTAP Select 9.10.1.
- Size of supported physical disks: 200GB – 32TB
- Only supported on DAS configurations
- Supported with either SSDs or NVMe
- Requires a Premium or Premium XL ONTAP Select license
- The hardware RAID controller should be absent or disabled or it should operate in SAS HBA mode
- An LVM storage pool or datastore based on a dedicated LUN must be used for system disks: core dump, boot/NVRAM, and the Mediator.

Does ONTAP Select for KVM support multiple NIC bonds?

When installing on KVM, you must use a single bond and a single bridge. A host with two or four physical ports should have all the ports in the same bond.

How does ONTAP Select report or alert for a failed physical disk or a NIC in the hypervisor host? Does ONTAP Select retrieve this information from the hypervisor or should monitoring be set at the hypervisor level?

When using a hardware RAID controller, ONTAP Select is largely unaware of underlying server issues. If the server is configured according to our best practices, a certain amount of redundancy should exist. We

recommend RAID 5/6 to survive drive failures. For software RAID configurations, ONTAP is responsible for issuing alerts about disk failure and, if there is a spare drive, initiate the drive rebuild.

You should use a minimum of two physical NICs to avoid a single point of failure at the network layer. NetApp recommends that Data, Mgmt, and Internal port groups have NIC teaming and bonding configured with two or more uplinks in the team or bond. Such configuration ensures that, if there is any uplink failure, the virtual switch moves the traffic from the failed uplink to a healthy uplink in the NIC team. For details about the recommended network configuration, see Summary of best practices: Networking.

All other errors are handled by ONTAP HA in the case of a two-node or four-node cluster. If the hypervisor server needs to be replaced and the ONTAP Select cluster needs to be reconstituted with a new server, refer to *Can I recover an ONTAP Select node?*.

Can I recover an ONTAP Select node?

Contact NetApp Technical Support if you need to recover an ONTAP Select node, for example, if an ONTAP Select VM or host is completely lost or a single-node cluster is accidentally deleted. For more information, refer to the Knowledge Base article ONTAP Select node VM recovery in case of disaster.

What is the maximum datastore size that ONTAP Select supports?

All configurations, including vSAN, support 400TB of storage per ONTAP Select node.

When installing on datastores larger than the supported maximum size, you must use Capacity Cap during product setup.

How can I increase the capacity of an ONTAP Select node?

ONTAP Select Deploy contains a storage add workflow that supports the capacity expansion operation on an ONTAP Select node. You can expand the storage under management by using space from the same datastore (if any space is still available) or add space from a separate datastore. The mixing of local datastores and remote datastores in the same aggregate is not supported.

Storage add also supports software RAID. However, in the case of software RAID, additional physical drives must be added to the ONTAP Select VM. The storage add in this case is similar to managing a FAS or AFF array. RAID group sizes and drive sizes must be considered when adding storage to an ONTAP Select node using software RAID.

Does ONTAP Select support vSAN or external array type datastores?

ONTAP Select Deploy and ONTAP Select for ESXi support the configuration of an ONTAP Select single-node cluster using either a vSAN or an external array type of datastore for its storage pool.

ONTAP Select Deploy and ONTAP Select for KVM support the configuration of an ONTAP Select single-node cluster using a shared logical storage pool type on external arrays. The storage pools can be based on iSCSI or FC/FCoE. Other types of storage pools are not supported.

Multi-node HA clusters on shared storage are supported.

Does ONTAP Select support multi-node clusters on vSAN or other shared external storage including some HCI stacks?

Multi-node clusters using external storage (multi-node vNAS) are supported for both ESXi and KVM. Mixing of hypervisors in the same cluster is not supported. An HA architecture on shared storage still implies that each node in an HA pair has a mirror copy of its partner data. However, a multi-node cluster brings in the benefits of ONTAP nondisruptive operations as opposed to a single-node cluster which relies on VMware HA or KVM Live Motion.

Although ONTAP Select Deploy adds support for multiple ONTAP Select VMs on the same host, it does not

allow those instances to be part of the same ONTAP Select cluster during cluster creation. For ESXi environments, NetApp recommends creating VM anti-affinity rules so that VMware HA does not attempt to migrate multiple ONTAP Select VMs from the same ONTAP Select cluster onto a single ESXi host. Furthermore, if ONTAP Select Deploy detects that an administrative (user-initiated) vMotion or live migration of an ONTAP Select VM has resulted in a violation of our best practice such as two ONTAP Select nodes ending up on the same physical host, ONTAP Select Deploy posts an alert in the Deploy UI and log. The only way that ONTAP Select Deploy becomes aware of the ONTAP Select VM location is as a result of a Cluster Refresh operation, which is a manual operation that the ONTAP Select Deploy administrator must initiate. There is no functionality in ONTAP Select Deploy that enables proactive monitoring, and the alert is only visible through the Deploy UI or log. In other words, this alert cannot be forwarded to a centralized monitoring infrastructure.

Does ONTAP Select support VMware's NSX VXLAN?

NSX-V VXLAN port groups are supported. For multi-node HA, including ONTAP MetroCluster SDS, make sure that you configure the internal network MTU to be between 7500 and 8900 (instead of 9000) to accommodate the VXLAN overhead. The internal network MTU can be configured with ONTAP Select Deploy during cluster deployment.

Does ONTAP Select support KVM live migration?

ONTAP Select VMs that run on external array storage pools support virsh live migrations.

Do I need ONTAP Select Premium for vSAN AF?

No, all versions are supported regardless of whether the external array or vSAN configurations are all flash.

What vSAN FTT/FTM settings are supported?

The Select VM inherits the vSAN datastore storage policy, and there are no restrictions on FTT/FTM settings. However, note that, depending on the FTT/FTM settings, the ONTAP Select VM size can be significantly larger than the capacity configured during its setup. ONTAP Select uses thick-eager, zeroed VMDKs that are created during setup. To avoid affecting other VMs using the same shared datastore, it is important to provide enough free capacity in the datastore to accommodate the true Select VM size as derived from the Select capacity and the FTT/FTM settings.

Can multiple ONTAP Select nodes run on the same host if they are part of different Select clusters?

It is possible to configure multiple ONTAP Select nodes on the same host for vNAS configurations only, as long as these nodes are not part of the same ONTAP Select cluster. This is not supported for DAS configurations because multiple ONTAP Select nodes on the same physical host would compete for access to the RAID controller.

Can you have a host with a single 10GE port run ONTAP Select, and is it available for both ESXi and KVM?

You can use a single 10GE port to connect to the external network. However, NetApp recommends that you use this only in constrained small form-factor environments. This is supported with both ESXi and KVM.

What additional processes do you need to run to do a live migration on KVM?

You must install and run open-source CLVM and pacemaker (pcs) components on each host participating in the live migration. This is required to access the same volume groups on each host.

vCenter

There are several questions and answers dealing with VMware vCenter.

How does ONTAP Select Deploy communicate with vCenter and what firewall ports should be opened?

ONTAP Select Deploy uses the VMware VIX API to communicate with the vCenter, ESXi host, or both. The VMware documentation states that the initial connection to either a vCenter Server or an ESXi host is done

using HTTPS/SOAP on TCP port 443. This is the port for secure HTTP over TLS/SSL. Secondly, a connection to the ESXi host is opened on a socket on TCP port 902. Data going over this connection is encrypted with SSL. Additionally, ONTAP Select Deploy issues a `PING` command to verify that there is an ESXi host responding at the IP address you specified.

ONTAP Select Deploy must also be able to communicate with the ONTAP Select node and cluster management IP addresses as follows:

- Ping
- SSH (port 22)
- SSL (port 443)

For two-node clusters, ONTAP Select Deploy hosts the cluster mailboxes. Each ONTAP Select node must be able to reach ONTAP Select Deploy through iSCSI (port 3260).

For multi-node clusters, the internal network must be fully opened (no NAT or firewalls).

What vCenter rights does ONTAP Select Deploy need to create ONTAP Select clusters?

The list of vCenter rights required is available here: VMware vCenter server.

HA and clusters

There are several questions and answers dealing with HA pairs and clusters.

What is the difference between a four-node, six-node, or eight-node cluster and a two-node ONTAP Select cluster?

Unlike four-node, six-node, and eight-node clusters in which the ONTAP Select Deploy VM is primarily used to create the cluster, a two-node cluster continuously relies on the ONTAP Select Deploy VM for HA quorum. If the ONTAP Select Deploy VM is unavailable, then failover services are disabled.

What is MetroCluster SDS?

MetroCluster SDS is a lower-cost synchronous replication option that falls under the category of the MetroCluster Business Continuity solutions from NetApp. It is available only with ONTAP Select, unlike NetApp MetroCluster that is available on FAS Hybrid Flash, AFF, and NetApp Private Storage for Cloud.

How is the MetroCluster SDS different from NetApp MetroCluster?

MetroCluster SDS provides a synchronous replication solution and falls under NetApp MetroCluster solutions. However, a key difference is the connectivity type (only IP networks are supported rather than FC and IP).

What is the difference between a two-node ONTAP Select cluster and a two-node ONTAP MetroCluster SDS?

The two-node cluster is defined as a cluster for which both nodes are in the same data center within 300m of each other. In general, both nodes have uplinks to the same network switch or set of network switches connected by an Inter-Switch Link.

The two-node MetroCluster SDS is defined as a cluster with nodes that are physically separated (different rooms, different buildings, or different data centers) and each node's uplink connections are connected to separate network switches. Although MetroCluster SDS does not require dedicated hardware, the environment should support the minimum requirements for latency (5ms RTT and 5ms jitter for a max total of 10ms).

MetroCluster SDS is a premium feature and requires the Premium or Premium XL license. A Premium license supports the creation of both Small and Medium VMs as well as HDD and SSD media. All these configurations are supported.

Does the ONTAP MetroCluster SDS require local storage (DAS)?

ONTAP MetroCluster SDS supports all type of storage configurations (DAS and vNAS).

Does ONTAP MetroCluster SDS support software RAID?

Yes, Software RAID is supported with SSD media on both KVM and ESXi.

Does ONTAP MetroCluster SDS support both SSDs and spinning media?

Yes, although a Premium license is required, this license supports both small and medium VMs as well as SSDs and spinning media.

Does ONTAP MetroCluster SDS support four-node and larger cluster sizes?

No, only two-node clusters with a Mediator can be configured as MetroCluster SDS.

What are the requirements for ONTAP MetroCluster SDS?

The requirements are as follows:

- Three data centers (one for the ONTAP Select Deploy Mediator and one for each node)
- 5ms RTT and 5ms jitter for a max total of 10ms
- 125ms RTT and a minimum bandwidth of 5Mbps between the ONTAP Select Deploy Mediator and each ONTAP Select node
- A Premium or Premium XL license

Does ONTAP Select support vMotion or VMware HA?

ONTAP Select VMs that run on vSAN datastores or external array datastores (in other words, vNAS deployments) support vMotion, DRS, and VMware HA functionality.

Does ONTAP Select support Storage vMotion?

Storage vMotion is supported for all configurations, including single-node and multi-node ONTAP Select clusters and the ONTAP Select Deploy VM. Storage vMotion can be used to migrate the ONTAP Select or the ONTAP Select Deploy VM between different VMFS versions (VMFS 5 to VMFS 6 for example), but it is not restricted to this use case. The best practice is to shut down the VM before initiating a Storage vMotion operation. ONTAP Select Deploy must issue the following operation after the storage vMotion operation is completed:

```
cluster refresh
```

A storage vMotion operation between different types of datastores is not supported. In other words, storage vMotion operations between NFS-type datastores and VMFS datastores are not supported. In general, storage vMotion operations between external datastores and DAS datastores are not supported.

Can the HA traffic between ONTAP Select nodes run over a different vSwitch and/or segregated physical ports and/or using point-to-point IP cables between ESXi hosts?

These configurations are not supported. ONTAP Select does not have visibility into the status of the physical network uplinks carrying client traffic. Therefore, ONTAP Select relies on the HA heartbeat to make sure that the VM is accessible to clients and to its peer at the same time. When a loss of physical connectivity occurs, the loss of the HA heartbeat results in an automatic failover to the other node, which is the desired behavior.

Segregating the HA traffic on a separate physical infrastructure can result in a Select VM being able to

communicate with its peer but not with its clients. This prevents the automatic HA process and results in data unavailability until a manual failover is invoked.

Mediator service

There are several questions and answers dealing with the mediator service.

What is the Mediator service?

A two-node cluster continuously relies on the ONTAP Select Deploy VM for HA quorum. An ONTAP Select Deploy VM taking part in a two-node HA quorum negotiation is labeled a Mediator VM.

Can the Mediator service be remote?

Yes. ONTAP Select Deploy acting as a Mediator for a two-node HA pair supports a WAN latency of up to 500ms RTT and requires a minimum bandwidth of 5Mbps.

What protocol does the Mediator service use?

The Mediator traffic is iSCSI, originates on the ONTAP Select node management IP addresses, and terminates on the ONTAP Select Deploy IP address. Note that you cannot use IPv6 for the ONTAP Select node management IP address when using a two-node cluster.

Can I use one Mediator service for multiple two-node HA clusters?

Yes. Each ONTAP Select Deploy VM can serve as a common Mediator service for up to 100 two-node ONTAP Select clusters.

Can the Mediator service location be changed after deployment?

Yes. It is possible to use another ONTAP Select Deploy VM to host the Mediator service.

Does ONTAP Select support stretched clusters with (or without) the Mediator?

Only a two-node cluster with a Mediator is supported in a stretched HA deployment model.

Legal notices for ONTAP Select

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

https://www.netapp.com/company/legal/copyright/

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

https://www.netapp.com/company/legal/trademarks/

Patents

A current list of NetApp owned patents can be found at:

https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf

Privacy policy

https://www.netapp.com/company/legal/privacy-policy/

Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

Notice for ONTAP Select 9.19.1

Copyright information

Copyright © 2026 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.