



ONTAP Select 9.14.1 documentation

ONTAP Select

NetApp
November 11, 2024

Table of Contents

- ONTAP Select 9.14.1 documentation 1
- Release Notes 2
 - ONTAP Select Release Notes 2
 - What’s new in this release 2
- Concepts 6
 - ONTAP Select overview 6
 - ONTAP Select Deploy 8
 - Business use cases 9
 - Terminology and key concepts 14
- Plan 19
 - ONTAP Select installation and deployment workflow 19
 - ONTAP Select 19
 - ONTAP Select Deploy 38
 - Summary of best practices 42
- License 47
 - Options 47
 - Capacity pools licensing model 50
 - Purchase 54
 - ONTAP features 57
- Install 59
 - Pre-installation checklist 59
 - Install ONTAP Select Deploy 74
 - Deploy an ONTAP Select cluster 79
 - Initial state of the cluster after deployment 83
- Administer 84
 - Before you begin administering ONTAP Select 84
 - Upgrade the ONTAP Select nodes 85
 - Diagnostics and support 86
 - Security 88
 - Confirming connectivity among the ONTAP Select nodes 93
 - Administering the Deploy mediator services 94
 - ONTAP Select clusters 94
 - Nodes and hosts 96
 - ONTAP Select licenses 108
- Deep dive 112
 - Storage 112
 - Networking 141
 - High availability architecture 165
 - Performance 173
- Automate with REST 176
 - Concepts 176
 - Access with a browser 183
 - Workflow processes 185

Access with Python	192
Python code samples	194
Automate ONTAP Select deployments with Ansible	220
Roles	220
Example Playbook	221
Use the CLI	223
Sign in to Deploy using SSH	223
Deploy an ONTAP Select cluster using the CLI	223
Security	230
Change the Deploy administrator password	230
Confirm network connectivity among the ONTAP Select nodes	231
ONTAP Select clusters	232
Nodes and hosts	232
Deploy utility	237
Frequently asked questions	249
General	249
Licenses, installation, upgrades, and reverts	250
Storage	251
vCenter	254
HA and clusters	255
Mediator service	257
Legal notices	258
Copyright	258
Trademarks	258
Patents	258
Privacy policy	258
Open source	258

ONTAP Select 9.14.1 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 9.14.1 Release Notes](#) to view the details about the current version.

What's new in this release

NetApp periodically updates ONTAP Select to bring you new features and enhancements.

ONTAP Select 9.14.1

ONTAP Select 9.14.1 includes several new features and improvements.

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, see the [ONTAP Select 9.14.1 Release Notes](#).

Enhanced VMware ESXi support

ONTAP Select 9.14.1 includes support for VMware ESXi 8.0.2.

ONTAP Select 9.13.1

ONTAP Select 9.13.1 includes several new features and improvements.

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, see 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, see the [ONTAP Select 9.14.1 Release Notes](#).

ONTAP Select 9.11.1

ONTAP Select 9.11.1 includes several new features and improvements.

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, see the [EOA Notice](#)

ONTAP Select 9.10.1

ONTAP Select 9.10.1 includes several new features and improvements.

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

ONTAP Select 9.9.1 includes several new features and improvements.

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

There are several new and changed features included in 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

ONTAP Select overview

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 FAS and AFF ONTAP offerings as well as other software-only options such as Cloud Volumes ONTAP.

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, the most valuable IT assets can be separated 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 mobility 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, or eight 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 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.



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.

- Autonomous Ransomware Protection (ARP)
- Cluster IPspace

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

- 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.
- Multi-Tenant Key Manager (MTKM)
- 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 migration
- SnapLock Compliance
- VMware HCX

Related information

- [ONTAP features enabled by default](#)

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 GUI 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

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.

Remote and branch office

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.



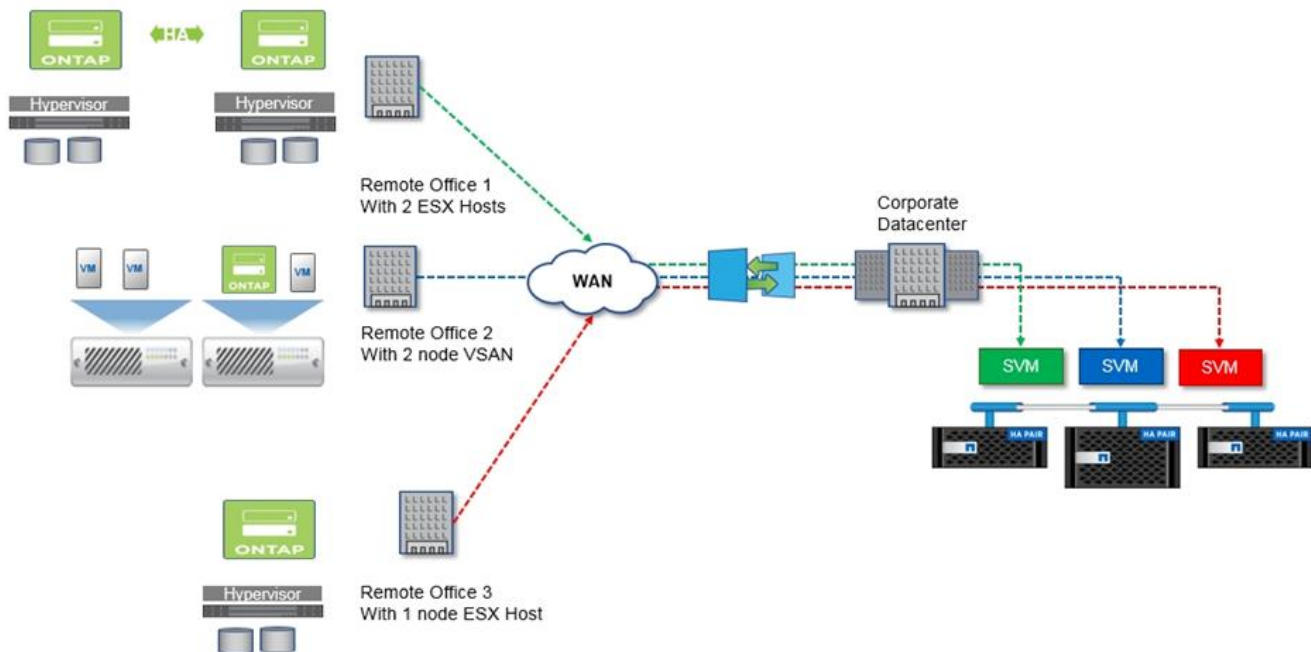
Because all VMware vSphere licenses are supported, you can choose the vSphere Remote Office Branch Office Standard or Advanced license instead of the Enterprise or Enterprise Plus license.

All vSphere and VSAN licenses are now 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



Private cloud and data center


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



MetroCluster software defined storage

ONTAP Select MetroCluster 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. A maximum distance of 10km is also a requirement, although the latency profile is more important. Separation requirements in the market space have more to do with 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.

With this solution, enterprise customers can confidently take advantage of the flexibility of a software-defined storage solution on commodity hardware. They can deploy with peace of mind knowing their data is protected with an enterprise-grade, 0 RPO solution.

ONTAP Select MetroCluster SDS provides the following benefits:

- MetroCluster SDS provides another dimension (data center to data center) of protection for ONTAP Select. Customers 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



Metro/Campus Network:

- 5ms RTT/5ms jitter
- Maximum latency 10ms
- 10KM distance between nodes

For more best practices and other requirements, see the sections [Two-node HA versus multi-node HA](#) and [Two-node stretched HA \(MetroCluster SDS\) best practices](#).

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, or eight 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, see 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. See 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. See 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

Beginning with ONTAP Select Deploy 2.8, 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.



ONTAP Select

General 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, see 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, or eight 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](#)

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 current release notes for your version of ONTAP Select 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

ONTAP Select can be deployed on the following hypervisors:

- KVM on RedHat Enterprise Linux 8.6, 8.7, 8.8, 9.0, 9.1, and 9.2
- KVM on Rocky Linux 8 and 9
- VMware ESXi 7.0 GA (build 15843807 or greater) including 7.0 U1, U2, and U3C
- VMware ESXi 8.0 GA (build 20513097)
- VMware ESXi 8.0 U1 (build 21495797)



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



ESXi 6.5 GA and ESXi 6.7 GA are reaching end of availability status. If you have ONTAP Select clusters with these versions, you must upgrade to the supported versions as per the [Interoperability Matrix Tool \(IMT\)](#).

Upgrade to VMware ESXi 6.5 U2 or later

If you currently have ONTAP Select deployed on VMware ESXi 6.5 U1, you should upgrade to ESXi 6.5 U2 or later as soon as possible. Using ESXi 6.5 U1 can expose you to a virtual machine failure due to a known VMware bug.

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.



Beginning with ONTAP Select 9.9.1 only CPU models based on Intel Xeon Sandy Bridge or later are supported.

Refer to the [NetApp 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 the following:

- Intel Xeon processors for Server (see [Intel Xeon Processors](#) 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
4/6/8 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.

Refer to [Understand the platform license offerings](#) for more information.

Small

- CPU cores
Six physical cores or greater, with four reserved for ONTAP Select.
- Memory
24GB or greater with 16GB reserved for ONTAP Select.
- Required platform license offering
Standard, premium, or premium XL

Medium

- CPU cores
Ten physical cores or greater, with eight reserved for ONTAP Select.
- Memory
72GB or greater with 64GB reserved for ONTAP Select
- Required platform license offering
Premium or premium XL

Large

- CPU cores
Eighteen physical cores or greater, with sixteen reserved for ONTAP Select.
- Memory
136GB or greater with 128GB reserved for ONTAP Select
- Required platform license offering
Premium XL



There are additional disk requirements based on the platform license. See [Storage and RAID](#) for more information.

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 in [ONTAP Select vNAS 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; multiple RAID controllers can be used if needed:
 - Disable the local drive cache for RAID group, which is fundamental to preserving data integrity.
- LUN configuration must be performed based on the following guidelines:
 - If the RAID group size exceeds the maximum LUN size of 64TB, you should 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, you should 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 utilize the software RAID capability provided by ONTAP Select instead of a local hardware RAID controller. There are several requirements and restrictions you must be aware before deploying a cluster using software RAID.

General requirements

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

- VMware ESXi 7.0 GA (build 15843807) 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, see the [Deep dive storage](#) section for additional configuration requirements.

ESXi specific requirements

- VMware ESXi 7.0 GA (build 15843807) 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. See [Preparation of the Linux server](#) 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 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 [NetApp 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. See 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. This storage is not used by ONTAP Select, unless a capacity cap is specified. If a capacity cap is specified, then that amount of storage will be used 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 9.7 or later with a supported Deploy administration utility
- Premium XL platform license offering or a 90-day evaluation license
- VMware ESXi version 6.7 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 requirements

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 [NetApp 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

Consult the available VMware documentation and adhere to 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)
- Assure that all generic external storage arrays use the available redundancy and HA features where possible

KVM requirements

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 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 high-availability (HA) features where possible

Networking considerations

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, see 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

See 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

See 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 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 Deploy utility providing administrative support. With this configuration, the management traffic between the Deploy utility and cluster flows over the WAN. See 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 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 Deploy administration utility. This requirement applies to all Deploy instances that manage one or more ONTAP Select two-node clusters.

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.

Preparing for a MetroCluster SDS deployment

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, however the distance between the two nodes can be up to 10 km. 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.

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.

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.

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:

- 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

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

There are several minimum resource requirements that your ESXi hypervisor host must meet. You should make sure 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

You must make sure 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 there is an ESXi host at the IP address you specify.

Deploy must also 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 language restriction

If you use ONTAP Select Deploy to create a cluster running on ESXi with vCenter on a Windows Server, you must use an English language version. ONTAP Select Deploy does not support vCenter on non-English versions of Windows.

Summary of best practices

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 ESX 6.5 U2 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 ESX 6.5 U2 and higher, ONTAP Select uses the vNVME 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 ESX 6.5 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, see 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 back up the Deploy configuration data on a regular basis, 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.



It's recommended 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 additional 10% of space may be used by the filesystem 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 may have a negative impact on performance.

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 ESX 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

Evaluation licenses

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.

Licensing characteristics

The ONTAP Select evaluation license has the following characteristics:

- A production license with storage capacity is not required
- The node serial number is twenty digits and automatically generated by ONTAP Select Deploy
(you do not 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
- A capacity tier license can be used, however capacity pools licensing is not 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.

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 tiers or capacity pools licensing model as well as the storage capacity for each deployment.

Common licensing characteristics

The *capacity tiers* and *capacity pools* 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.
- 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.
- All platform licensing offerings are supported (standard, premium, 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.
- Both nodes in an HA pair must have the same storage and license capacity.
- 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.
- A nine-digit license serial number is generated by NetApp for each ONTAP Select node.
- The storage allocated to a node is perpetual (no renewal required).
- 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.
- 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 node serial number is twenty digits and is generated by the License Manager based on the capacity pool license serial number.
- Each node automatically leases storage capacity for its local data aggregates from a shared capacity pool.

For more details on the capacity pools licensing model, see the *Capacity pools licensing model* for more information.

Understand the platform license offerings

You can purchase an ONTAP Select capacity tier or capacity pool license in one of three platform capacity levels. These license offerings determine the capabilities of the hosts where you deploy ONTAP Select.

What a platform license offering provides

A specific license offering defines and restricts 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.

Platform license offerings

There are three platform license offerings available.

Standard

The standard offering provides the following capabilities:

- Small instance type only
- Hard disk drives (HDD) only
- Local hardware RAID controller only
- vNAS

Premium

The premium offering provides the following capabilities:

- Small or medium instance type
- Hard disk drives (HDD) or Solid state drives (SSD)
- Local hardware RAID controller or software RAID
- vNAS
- MetroCluster SDS

Premium XL

The premium XL offering provides the following capabilities:

- Small, medium, or large instance type
- HDDs, SSDs, or NVMe drives
- Local hardware RAID controller or software RAID
- vNAS

- MetroCluster SDS



Use of the large instance type or NVMe drives in an SW-RAID configuration is not supported on KVM.

Capacity pools licensing model

Operational details

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**
This page is displayed when you first sign in to Deploy. You can also display the page by clicking the drop down box at the top right of the page and selecting Getting Started. The LLID is displayed in the Add Licenses section.
- **Administration**
Click the **Administration** tab at the top of the page, then click **Systems** and **Settings**.

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. A lease obtained from a previous request that is still valid can be used, or a new lease can be requested 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

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 nnnnnnnnn



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

Each section of the node serial number is described in the following table, 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

Section	Description
nnnnnnnn	Variable eight-digit value generated by the License Manager for each node using the capacity pool



Attention: When opening a case with NetApp support involving a node that uses a capacity pool license, you cannot provide the full twenty-digit node serial number. Instead, you must provide the nine-digit capacity pool license serial number. You can derive the license serial number from the node serial number as shown above. Skip the first three digits of the node serial number ('999') and extract the next nine digits (ppppppppp).

Deployment restrictions for 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 all nodes in a cluster must be managed by the same Deploy instance.

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.

Comparing capacity pools and capacity tiers licensing

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

	Capacity tiers	Capacity pools
License serial number	Nine digits generated by NetApp and assigned to a node	Nine digits generated by NetApp and assigned 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 (renewed 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
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)

Summary of benefits

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.

And because the capacity is anchored in a central pool, 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.

Purchase

Workflow when purchasing a license

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

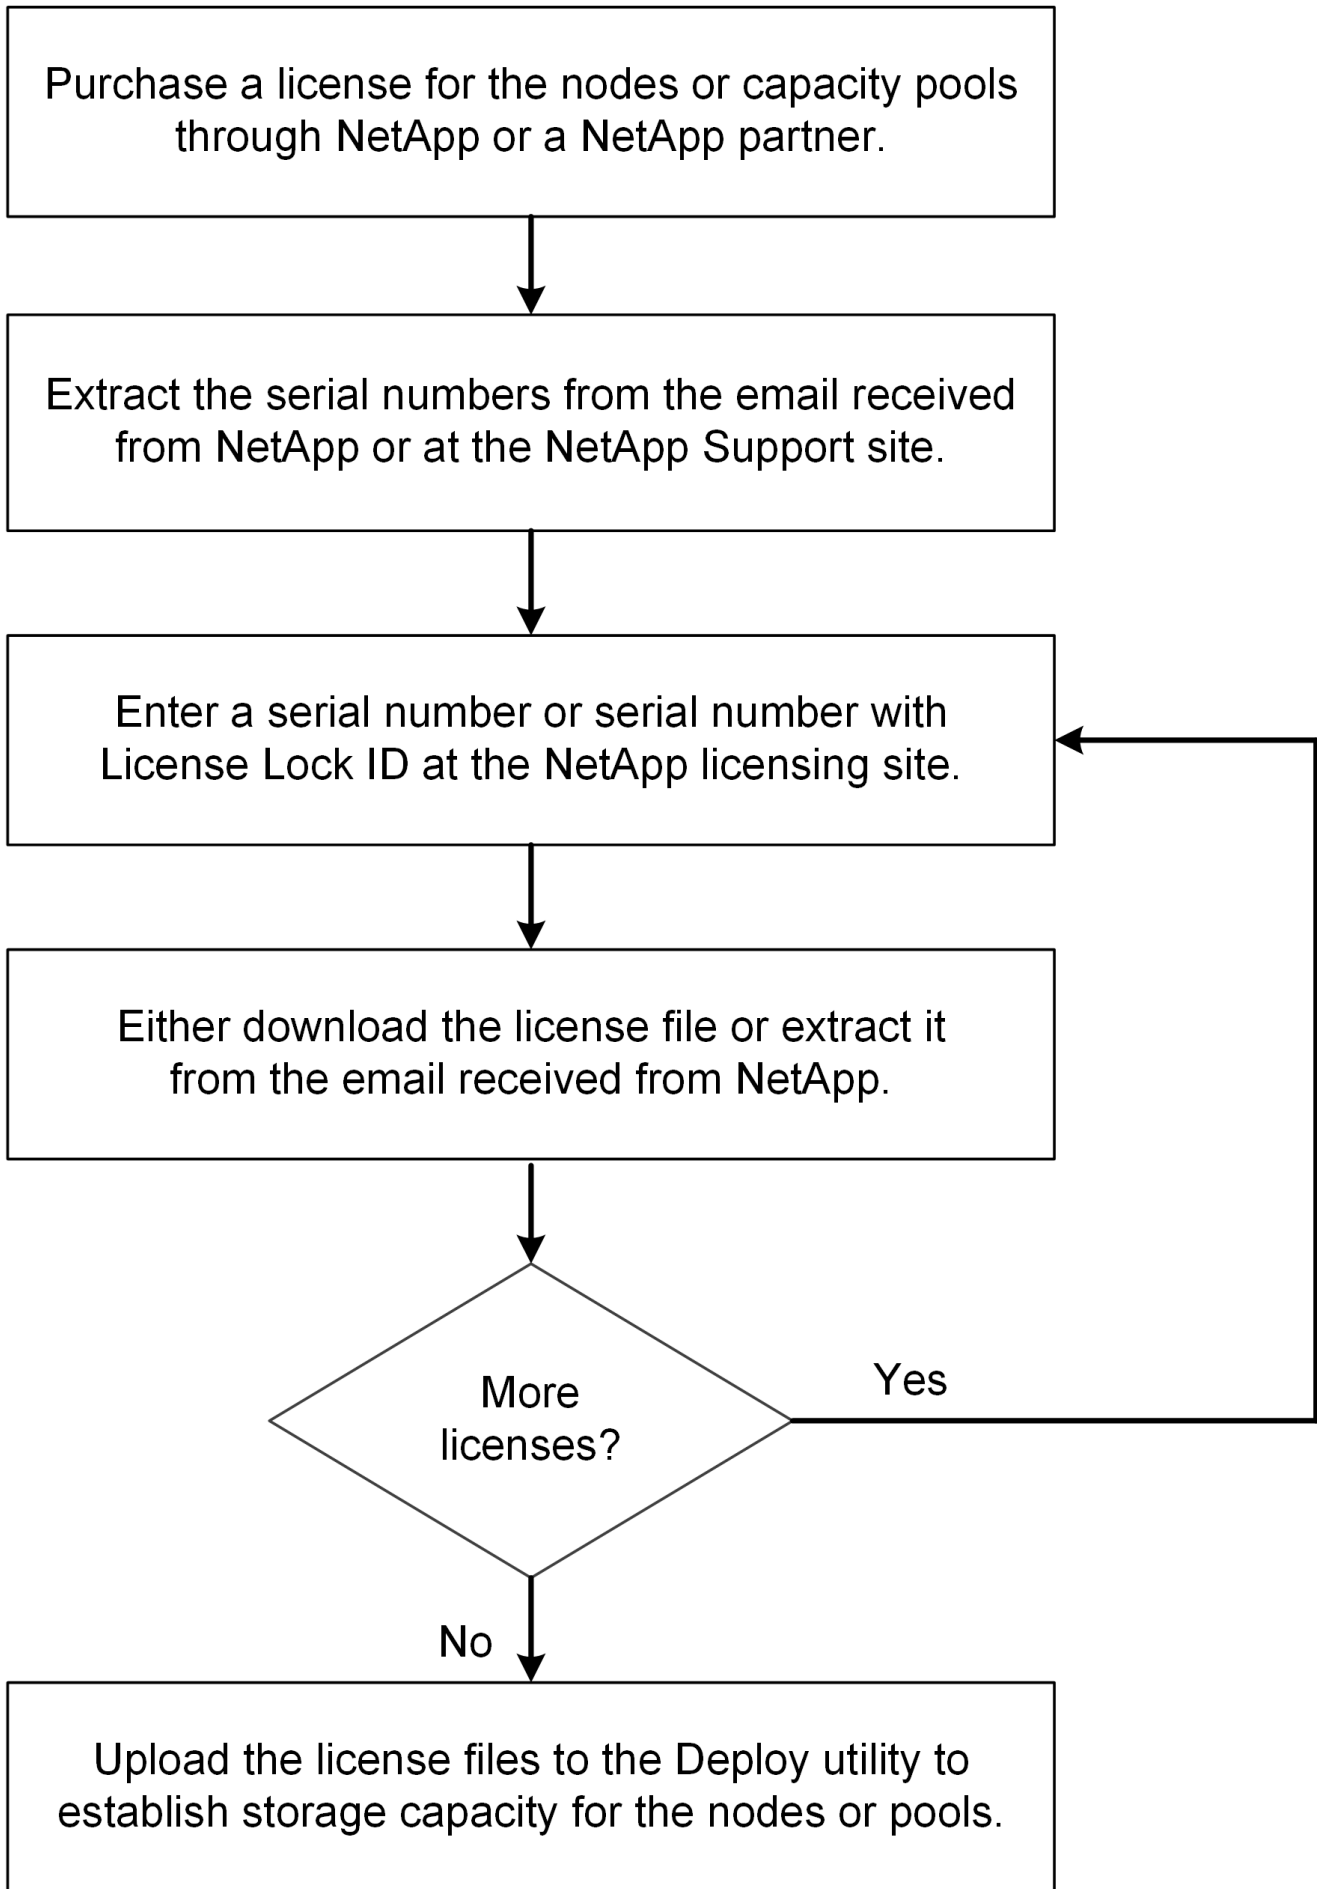
Serial number applies to either a node (capacity tiers) or a storage pool (capacity pools)

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



Acquiring a 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 node by NetApp. Before attempting to acquire a license file, you should wait at least twenty four hours after the shipped date of your purchase order.

About this task

You must perform this task for each ONTAP Select node requiring a capacity tier license.

Steps

1. Access the ONTAP Select license site using a web browser:

<https://register.netapp.com/register/getlicensefile>

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 same page, including the **Product Serial #**, which is the serial number for the ONTAP Select node.
5. Click **Submit**.
6. After the request has been validated, select the delivery method for the license.

You can click 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.

Acquire a 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. Before attempting to acquire a license file, you should wait at least twenty four hours after the shipped date of your purchase order.

About this task

You must perform this task for each capacity pool used by the ONTAP Select nodes.

Steps

1. Access the NetApp Support Site using a web browser and sign in.
2. Click **Systems** at the top and then click **Software Licenses**.
3. Type the license serial number for the capacity pool and click **Go!**
4. On the license details page, navigate to the **Product Details** column.
5. Click **Get NetApp License File** on the appropriate row.
6. Type the License Lock ID for your ONTAP Select Deploy instance and click **Submit**.
7. Select the appropriate delivery method and click **Submit**.
8. Click **OK** on the delivery confirmation window.

After you finish

You must upload the license file to the Deploy administration utility before the capacity pool can be used by an ONTAP Select node.

ONTAP features

ONTAP Select offers full support for most ONTAP functionality. 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 automatically enabled by default

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

- CIFS
- Deduplication and compression
- FlexCache
- FlexClone
- iSCSI
- NDMP
- NetApp Volume Encryption (non-restricted countries only)
- NFS
- NVMe over TCP
- ONTAP multitenancy capability
- ONTAP S3
- SnapMirror
- SnapRestore
- SnapVault
- Storage VM disaster recovery (SVM DR)



ONTAP Select 9.12.1 and later now 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 9.12.1 source can connect to destination ONTAP versions 9.12.1, 9.13.1, or 9.14.1.

ONTAP features that are separately licensed

You must acquire a separate license for any ONTAP feature that is not enabled by default, including:

- Data Availability Service
- FabricPool
- MetroCluster SDS (ONTAP Select premium license offering)
- SnapLock Enterprise (SnapLock Compliance is not supported for ONTAP Select)
 - Tamperproof Snapshot copies
- SyncMirror (zero cost)



You do not need a FabricPool license when utilizing StorageGRID Webscale.

Related information

- [Comparing ONTAP Select and ONTAP 9](#)

Install

Pre-installation checklist

Host configuration and preparation checklist

Prepare each of the hypervisor hosts where an ONTAP Select node is deployed. As part of preparing the hosts, carefully assess the deployment environment to make sure that the hosts are properly configured and ready to support the deployment of an ONTAP Select cluster.



The ONTAP Select Deploy administration utility does not perform the required network and storage configuration of the hypervisor hosts. You must manually prepare each host prior to deploying an ONTAP Select cluster.

General hypervisor preparation

You must prepare the hypervisor hosts.

KVM hypervisor

Prepare the Linux server

You must 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.

Install Red Hat Enterprise Linux

You must install the Red Hat Enterprise Linux (RHEL) operating system using the ISO image. During installation, you should configure the system as follows:

- Select Default as the security policy
- Choose the Virtualized Host software selection
- The destination should be the local boot disk and not a RAID LUN used by ONTAP Select
- 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.

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 you are using. As a first step, verify that the yum repository is available on your server. If it is not available, you can retrieve it 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, see the [link:https://imt.netapp.com/matrix/#welcome](https://imt.netapp.com/matrix/#welcome) [NetApp Interoperability Matrix Tool^].

Additional packages required for RHEL 7.7

Install the same set of packages required for RHEL 7.6.

Additional packages required for RHEL 7.6

Verify that the following packages and dependencies are installed when using RHEL 7.6 or CentOS 7.6. In each case, the package name and version are included.

- qemu-kvm (1.5.3-160)



When using software RAID, you must use version 2.9.0 instead.

- libvirt (4.5.0-10)
- openvswitch (2.7.3)
- virt-install (1.5.0-1)
- lshw (B.02.18-12)
- lsscsi (0.27-6)
- lsof (4.87-6)

If you are using vNAS on KVM (external storage) and plan to migrate virtual machines from one host to another, you should install the following additional packages and dependencies:

- fence-agents-all (4.2.1-11)
- lvm2-cluster (2.02.180-8)
- pacemaker (1.1.19-8)
- pcs (0.9.165-6)

Additional packages required for RHEL 7.5

Verify that the following packages and dependencies are installed when using RHEL 7.5 or CentOS 7.5. In each case, the package name and version are included.

- qemu-kvm (1.5.3-141)



When using software RAID, you must use version 2.9.0 instead.

- libvirt (3.9.0)
- openvswitch (2.7.3)
- virt-install (1.4.1-7)
- lshw (B.02.18-12)
- lsscsi (0.27-6)

- lsof (4.87-5)

If you are using vNAS on KVM (external storage) and plan to migrate virtual machines from one host to another, you should install the following additional packages and dependencies:

- fence-agents-all (4.0.11-86)
- lvm2-cluster (2.02.177-4)
- pacemaker (1.1.18-11)
- pcs (0.9.16205)

Additional packages required for RHEL 7.4

Verify that the following packages and dependencies are installed when using RHEL 7.4 or CentOS 7.4. In each case the package name and version are included.

- qemu-kvm (1.5.3-141)



When using software RAID, you must use version 2.9.0 instead.

- libvirt (3.2.0-14)
- openvswitch (2.7.3)
- virt-install (1.4.1-7)
- lshw (B.02.18-7)
- lsscsi (0.27-6)
- lsof (4.87-4)

If you are using vNAS on KVM (external storage) and plan to migrate virtual machines from one host to another, you should install the following additional packages and dependencies:

- fence-agents-all (4.0.11-66)
- lvm2-cluster (2.02.171-8)
- pacemaker (1.1.16-12)
- pcs (0.9.158-6)

Configuration of 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.

Create a storage pool

You must 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
```

Delete a storage pool

You can delete a storage pool when it is 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_name>
```

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>
```

ESXi hypervisor

Each host must be configured with the following:

- A pre-installed and supported hypervisor
- A VMware vSphere license

Also, the same vCenter server must be able to manage all the hosts where an ONTAP Select node is deployed within the cluster.

In addition, you should make sure 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.

By default, VMware allows access on the following ports:

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

NetApp recommends that the following firewall ports are opened to allow access to vSphere:

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

You should also be familiar with the vCenter rights that are required. See [VMware vCenter server](#) for more information.

ONTAP Select cluster network preparation


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 HA capability.

Illustration of the ONTAP Select networks and nodes

The figures below illustrate the networks used with a single-node cluster and four-node cluster.


Single-node cluster showing one network

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




Four-node cluster showing two networks

The following figure illustrates a four-node cluster. 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.



KVM host

Configure Open vSwitch on a KVM host

You must configure a software-defined switch on each ONTAP Select node using Open vSwitch.

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 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.

1. Bring the network interfaces up:

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

ESXi host


vSwitch configuration on a hypervisor host

The vSwitch is the core hypervisor component used to support the connectivity for the internal and external networks. There are several things you should consider as part of configuring each hypervisor vSwitch.

vSwitch configuration for a host with two physical ports (2x10Gb)

When each host includes two 10Gb ports, you should configure the vSwitch as follows:

- Configure a vSwitch and assign both the ports to the vSwitch. Create a NIC team using the two ports.
- Set the load balancing policy to “Route based on the originating virtual port ID”.
- Mark both adapters as “active” or mark one adapter as “active” and the other as “standby”.
- Set the “Failback” setting to “Yes”.



- Configure the vSwitch to use jumbo frames (9000 MTU).
- Configure a port group on the vSwitch for the internal traffic (ONTAP-internal):
 - The port group is assigned to ONTAP Select virtual network adapters e0c-e0g used for the cluster, HA interconnect, and mirroring traffic.
 - The port group should be on a non-routable VLAN because this network is expected to be private. You should add the appropriate VLAN tag to the port group to take this into account.
 - The load balancing, failback, and failover order settings of the port group should be the same as the vSwitch.
- Configure a port group on the vSwitch for the external traffic (ONTAP-external):
 - The port group is assigned to ONTAP Select virtual network adapters e0a-e0c used for data and management traffic.
 - The port group can be on a routable VLAN. Also, depending on the network environment, you should add an appropriate VLAN tag or configure the port group for VLAN trunking.
 - The load balancing, failback, and failover order settings of the port group should be same as vSwitch.

The above vSwitch configuration is for a host with 2x10Gb ports in a typical network environment.

Required information for 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.
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.

Configuring a host to use NVMe drives

If you plan to use NVMe drives with software RAID, you need to configure the 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.

Before you begin

Make sure your deployment environment meets the following minimum requirements:

- ONTAP Select 9.7 or later with a supported Deploy administration utility
- Premium XL platform license offering or a 90-day evaluation license
- VMware ESXi version 6.7 or later
- NVMe devices conforming to specification 1.0 or later

Follow the [host preparation checklist](#), review the [required information for Deploy utility installation](#), and the [required information for ONTAP Select installation](#) topics for more information.

About this task


This procedure is designed to be performed before creating a new ONTAP Select cluster. You can also perform the procedure to configure additional NVMe drives for an existing SW-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 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 click **Edit** under **Hardware: PCI devices**.
- b. Select the NVMe drives you want to use for ONTAP Select.

Edit PCI Device Availability
sdot-dl380-003.gdl.englab.netapp.com
✕

ID	Status	Vendor Name	Device Name	ESX/ESXi Device
0000:36:01.0	Not Configurable	Intel Corporation	Sky Lake-E PCI Express...	
0000:38:...	Available (pending)	Seagate Technology ...	Nytro Flash Storage	
0000:36:02.0	Not Configurable	Intel Corporation	Sky Lake-E PCI Express...	
0000:39:...	Available (pending)	Seagate Technology ...	Nytro Flash Storage	

No items selected

CANCEL
OK



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. Click **OK**. The selected devices indicate **Available (pending)**.
5. Click **Reboot The Host**.

Configure
Permissions
VMs
Datstores
Networks
Updates

DirectPath I/O PCI Devices Available to VMs

ID	Status	Vendor Name	Device Name
0000:12:00.0	Available (pending)	Seagate Technology PLC	Nytro Flash Storage
0000:13:00.0	Available (pending)	Seagate Technology PLC	Nytro Flash Storage
0000:14:00.0	Available (pending)	Seagate Technology PLC	Nytro Flash Storage
0000:15:00.0	Available (pending)	Seagate Technology PLC	Nytro Flash Storage
0000:37:00.0	Available (pending)	Seagate Technology PLC	Nytro Flash Storage
0000:38:00.0	Available (pending)	Seagate Technology PLC	Nytro Flash Storage

REFRESH
EDIT...

7 devices will become available when this host is rebooted. Reboot This Host

After you finish

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 will detect

the presence of the NVMe drives configured for pass-through and automatically select 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.

ONTAP Select Deploy [Help] [User]

Clusters | Hypervisor Hosts | Administration

Storage

Storage Configuration

RAID Type: Software RAID | Data Disk Type: NVME

System Disk

nvme-snc-01
sdot-dl380-003-nvme(NVME)
Capacity: 1.41 TB



Data Disks for nvme-snc-01

	Device Name	Device Type	Capacity
<input checked="" type="checkbox"/>	0000:12:00.0	NVME	-
<input checked="" type="checkbox"/>	0000:13:00.0	NVME	-
<input checked="" type="checkbox"/>	0000:14:00.0	NVME	-
<input checked="" type="checkbox"/>	0000:15:00.0	NVME	-
<input checked="" type="checkbox"/>	0000:37:00.0	NVME	-
<input checked="" type="checkbox"/>	0000:38:00.0	NVME	-
<input checked="" type="checkbox"/>	0000:39:00.0	NVME	-

Selected Capacity: (7/7 disks)

Done

After the cluster is successfully deployed, ONTAP System Manager allows you to provision the storage according to best practices. ONTAP will automatically enable 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.

Download the virtual machine image

You can download the ONTAP Select package from the NetApp support site.

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`. The VM provides the Deploy server and installation images for ONTAP Select nodes.

Steps

1. Access the [NetApp Support Site](#) using a web browser and sign in.

2. Select **Downloads** from the menu, then select **Downloads** from the drop-down menu.
3. On the Downloads page, under All Products A-Z, select the letter **O**.
4. Scroll down and select **ONTAP Select**.
5. Select the desired release of the package.
6. Review the End User License Agreement (EULA) and select **Accept & Continue**.
7. Select and download the appropriate package, responding to all prompts as needed.

Verify the ONTAP Select Deploy OVA signature

You should 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 <code>csc-prod</code> 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 with the following command:

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

Deploy the virtual machine

You must 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, selecting the options as 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 and power it on if it is not already powered on based on your input to the deployment wizard.
5. If needed, you can configure the Deploy network using the VM console:
 - a. Click 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 press **Enter**.
- d. Wait for the following prompt:

Provide a password for the admin user:

- e. Type the password and press **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 network configuration information as required.

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.

Sign in to the Deploy web interface

You should sign in to the web user interface to confirm the Deploy utility is available and perform 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)

Related information

[Sign in to Deploy using SSH](#)

Deploy an ONTAP Select cluster

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](#).

Before you begin

You should prepare for the deployment to make sure it's successful.


Initial planning

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

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

Host preparation

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:

1. Sign in to the Deploy web user interface.
2. Click  at the top of the page.
3. Click **Prerequisites**.
4. Scroll down to review the requirements and click **OK**.

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.

Deploy installation and account credentials

You must install the Deploy administration utility and perform initial configuration. See [Install ONTAP Select Deploy](#) for more information. You need to have the password for the Deploy administrator account that was configured as part of the installation process.

Installing 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. See [Add an ONTAP Select image to Deploy](#) for more information.

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 clicking the tabs at the top of the page (Clusters, Hypervisor Hosts, Administration).

Network checker

If you will deploy a multi-node cluster, you should be familiar with the network checker.

After deploying a cluster

You should back up the ONTAP Select Deploy configuration data.

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

See [Before you begin](#) to prepare to deploy the cluster. The Deploy administration must be installed and initially configured (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.

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** popup window is displayed, confirm you have met the configuration prerequisites and click **OK**.
3. If the **Getting Started** cluster launch page is not displayed, click  at the top of the page and click **Getting Started**.
4. On the **Getting Started** page, click **Upload** and select a license from your local workstation and click **Open** to upload the license.
5. Click **Refresh** and confirm the license has been added.
6. Click **Next** to add a hypervisor host and then click **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. Click **Refresh** and confirm the **Type** value for the host is **ESX**.

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

8. Click **Next** to begin the cluster creation process.
9. In the **Cluster Details** section, provide all the required information describing the cluster and click **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 click **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 clicking  in the applicable section.


14. Click **Next** and provide the ONTAP administrator password.
15. Click **Create Cluster** to begin the cluster creation process and then click **OK** in the popup 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 click **OK**.
3. If the **Getting Started** cluster launch page is not displayed, click  at the top of the page and click **Getting Started**.
4. On the **Getting Started** page, click **Upload** and select a license from your local workstation and click **Open** to upload the license. Repeat to add a second license.
5. Click **Refresh** and confirm the licenses have been added.
6. Click **Next** to add two hypervisor hosts and then click **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. Click **Refresh** and confirm the **Type** value for the host is **ESX**.

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

8. Click **Next** to begin the cluster creation process.
9. In the **Cluster Details** section, select **2 node cluster** for the **Cluster Size**, provide all the required information describing the clusters, and click **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 click **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 clicking  in the applicable section.

14. Click **Next** and run the Network Precheck by clicking **Run**. This validates that the internal network selected for ONTAP cluster traffic is functioning correctly.
15. Click **Next** and provide the ONTAP administrator password.
16. Click **Create Cluster** to begin the cluster creation process and then click **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.

After you finish

You should 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 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, see the KB 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)

SVMs

Two administrative SVMs are active:

- Default SVM
- Cluster SVM

Aggregates

The root aggregate is created.

Features

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



There are no data SVMs created. Also, the multi-node cluster has an internal network with autogenerated LIFs.

Related information

- [ONTAP features enabled by default](#)

Administer

Before you begin 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 [NetApp 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

There are data spare disks created by the Deploy administration utility 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.

Upgrade the 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.

1. Navigate to downloads page at the NetApp Support Site.

[NetApp Support Downloads](#)

2. Click **ONTAP Select Node Upgrade**.
3. Select and download the appropriate upgrade image responding to all prompts as needed.

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

4. Upgrade the ONTAP Select node using the standard ONTAP upgrade procedures with the ONTAP Select upgrade file. For information on supported upgrade paths, see 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 9.7 is initially installed

You can upgrade the node to version 9.8 and then revert back to version 9.7 if needed.

ONTAP Select 9.8 is initially installed

You cannot revert to version 9.7 because this version is prior to the version that was originally 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](#)

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. Click the **Administration** tab at the top of the page.
3. Click **Settings & AutoSupport** and then click .
4. Provide the configuration data as appropriate for your environment and click **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. Click the **Administration** tab at the top of the page.
3. Click **Events & Jobs** and then click **Events**.
4. Optionally click **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. Click the **Administration** tab at the top of the page.
3. Click **Settings & AutoSupport** and then click .
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. Click the **Administration** tab at the top of the page.

3. Click **Settings & AutoSupport** and then click .
4. Click **Generate**.
5. Select the type and provide a description for the package; you can optionally provide a case number.
6. Click **Generate**.

Each AutoSupport package is assigned a unique sequence identification number.

7. Optionally under **AutoSupport History**, select the correct package and click the download icon to save the AutoSupport file to your local workstation.

Security

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. Click the figure icon at the top right of the page and select **Change Password**.
3. Provide the current and new password as prompted and click **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. Click the **Administration** tab at the top of the page.
3. Click **Management Servers** and then click **Add vCenter**.
4. Enter the following information and click **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. After the new management server is added, you can optionally click  and 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`

Confirming 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. Click the **Administration** tab at the top of the page and click **Network Checker**.
3. Click **Start New Run** and select the hosts and networks for the HA pair

You can add and configure additional HA pairs as needed.

4. Click **Start** to begin the network connectivity test.

Administering the 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. Click the **Administration** tab at the top of the page and click **Mediators**.
3. Optionally click **Filter** to customize your view of the two-node clusters monitored by the mediator service.

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. Click the **Clusters** tab at the top of the page and select the desired cluster from the list.
3. Click  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. Click **Yes** in the popup window to confirm the request.
5. Click **Refresh** occasionally to confirm the cluster is offline.
6. To bring the cluster back online, click  and select **Bring Online**.

7. Click **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. Click the **Clusters** tab at the top of the page and select the desired cluster from the list.
3. Click  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. Click **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 GUI, 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. Click the **Clusters** tab at the top left of the page and select the desired cluster from the list.
3. Click  on the right side of the page and select **Cluster Refresh**.
4. Under **Cluster Credentials**, provide the ONTAP administrator password for the cluster.
5. Click **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.

Nodes and hosts

Access the 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 click the virtual machine and select **Open Console**.

Resize the 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. Click the **Cluster** tab at the top of the page and select the desired cluster from the list.
3. On the cluster details page, click the gear icon at the right of the page and select **Instance Resize**.
4. Select the **Instance Type** and provide the ONTAP credentials then click **Modify**.

After you finish

You must wait for the resize operation to complete.

Replace a failed drive when using SW RAID

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.

Node Details

> **HA Pair 1**

	Node 1 sti-rx2540-345a — 8.73 TB + ⚡	Host 1 sti-rx2540-345 — (Small (4 CPU, 16 GB Memory))
	Node 2 sti-rx2540-346a — 8.73 TB + ⚡	Host 2 sti-rx2540-346 — (Small (4 CPU, 16 GB Memory))

3. Select **+** to expand the storage view.

Edit Node Storage

Node sti-rx2540-345a (Capacity: 135 GB, Licensed 50 TB)

Select License

Storage Disks Details

Edit

Data Disks for sti-rx2540-345a

ONTAP Name	Device Name	Device Type	Adapter	Capacity	Used by
NET-1.1	naa.5002538c40b4e044	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
NET-1.2	naa.5002538c40b4df4b	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
NET-1.3	naa.5002538c40b4e042	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
NET-1.4	naa.5002538c40b4e040	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
NET-1.5	naa.5002538c40b4e041	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
NET-1.6	naa.5002538c40b4df54	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
NET-1.7	naa.5002538c40b4df53	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
NET-1.8	naa.5002538c40b4df4a	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
NET-1.9	naa.5002538c40b4e03e	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
NET-1.10	naa.5002538c40b4e046	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...

4. Select **Edit** to make changes to the attached disks and uncheck the failed drive.

Node sti-rx2540-345a (Capacity: 135 GB, Licensed 50 TB)

Select License

Storage Disks Details

Select Disks for sti-rx2540-345a

	ONTAP Na...	Device Name	Device Type	Adapter	Capacity	Used by
<input checked="" type="checkbox"/>	NET-1.1	naa.5002538c40b4e044	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input checked="" type="checkbox"/>	NET-1.2	naa.5002538c40b4df4b	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input checked="" type="checkbox"/>	NET-1.3	naa.5002538c40b4e042	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input type="checkbox"/>	NET-1.4	naa.5002538c40b4e049	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input checked="" type="checkbox"/>	NET-1.5	naa.5002538c40b4e041	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input checked="" type="checkbox"/>	NET-1.6	naa.5002538c40b4df54	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input checked="" type="checkbox"/>	NET-1.7	naa.5002538c40b4df53	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input checked="" type="checkbox"/>	NET-1.8	naa.5002538c40b4df4a	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input checked="" type="checkbox"/>	NET-1.9	naa.5002538c40b4e03e	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...

Selected Capacity: 7.86 TB (9/10 disks)

5. Provide the cluster credentials and select **Edit Storage**.

Selected Capacity: 8.73 TB (10/10 disks)

ONTAP Credentials

Cluster Username admin

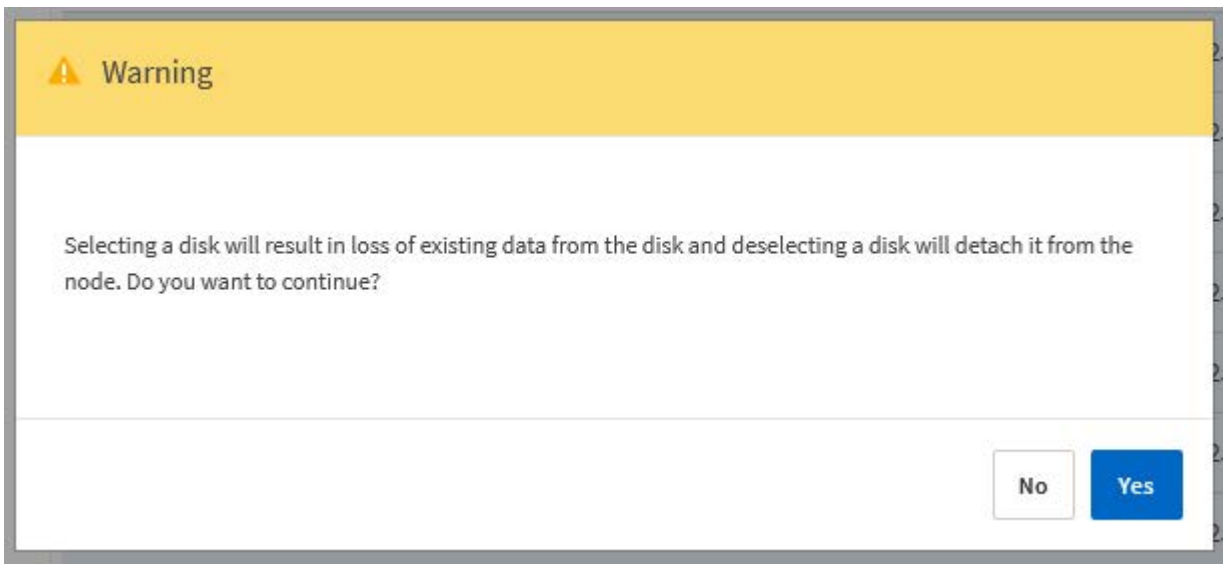
Cluster Password

••••••••

Cancel

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.

Node Details

HA Pair 1



Node 1 sti-rx2540-345a — 8.73 TB + ⚡ **Host 1** sti-rx2540-345 — (Small (4 CPU, 16 GB Memory))
Node 2 sti-rx2540-346a — 8.73 TB + ⚡ **Host 2** sti-rx2540-346 — (Small (4 CPU, 16 GB Memory))

3. Select + to expand the storage view.

Edit Node Storage

Node: sti-rx2540-345a (Capacity: 135 GB, Licensed 50 TB) [Select License](#)

Storage Disks Details [Edit](#)

Data Disks for sti-rx2540-345a

ONTAP Name	Device Name	Device Type	Adapter	Capacity	Used by
NET-1.1	naa.5002538c40b4e044	SSD	vmhba4	894.25 GB	sti-rx2540-345a=>...
NET-1.2	naa.5002538c40b4df4b	SSD	vmhba4	894.25 GB	sti-rx2540-345a=>...
NET-1.3	naa.5002538c40b4e042	SSD	vmhba4	894.25 GB	sti-rx2540-345a=>...
NET-1.4	naa.5002538c40b4e049	SSD	vmhba4	894.25 GB	sti-rx2540-345a=>...
NET-1.5	naa.5002538c40b4e041	SSD	vmhba4	894.25 GB	sti-rx2540-345a=>...
NET-1.6	naa.5002538c40b4df54	SSD	vmhba4	894.25 GB	sti-rx2540-345a=>...
NET-1.7	naa.5002538c40b4df53	SSD	vmhba4	894.25 GB	sti-rx2540-345a=>...
NET-1.8	naa.5002538c40b4df4a	SSD	vmhba4	894.25 GB	sti-rx2540-345a=>...
NET-1.9	naa.5002538c40b4e03e	SSD	vmhba4	894.25 GB	sti-rx2540-345a=>...
NET-1.10	naa.5002538c40b4e046	SSD	vmhba4	894.25 GB	sti-rx2540-345a=>...

4. Select **Edit** and confirm that the new drive is available and select it.

Node: sti-rx2540-345a (Capacity: 135 GB, Licensed 50 TB) [Select License](#)

Storage Disks Details

Select Disks for sti-rx2540-345a

ONTAP Na...	Device Name	Device Type	Adapter	Capacity	Used by	
<input checked="" type="checkbox"/>	naa.5002538c40b4e049	SSD	vmhba4	894.25 GB		
<input checked="" type="checkbox"/>	NET-1.1	naa.5002538c40b4e044	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input checked="" type="checkbox"/>	NET-1.2	naa.5002538c40b4df4b	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input checked="" type="checkbox"/>	NET-1.3	naa.5002538c40b4e042	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input checked="" type="checkbox"/>	NET-1.5	naa.5002538c40b4e041	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input checked="" type="checkbox"/>	NET-1.6	naa.5002538c40b4df54	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input checked="" type="checkbox"/>	NET-1.7	naa.5002538c40b4df53	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input checked="" type="checkbox"/>	NET-1.8	naa.5002538c40b4df4a	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...
<input checked="" type="checkbox"/>	NET-1.9	naa.5002538c40b4e03e	SSD	vmhba4	894.25 GB	sti-rx2540-345a=...

5. Provide the cluster credentials and select **Edit Storage**.

Selected Capacity: 8.73 TB (10/10 disks)

ONTAP Credentials

Cluster Username **admin**

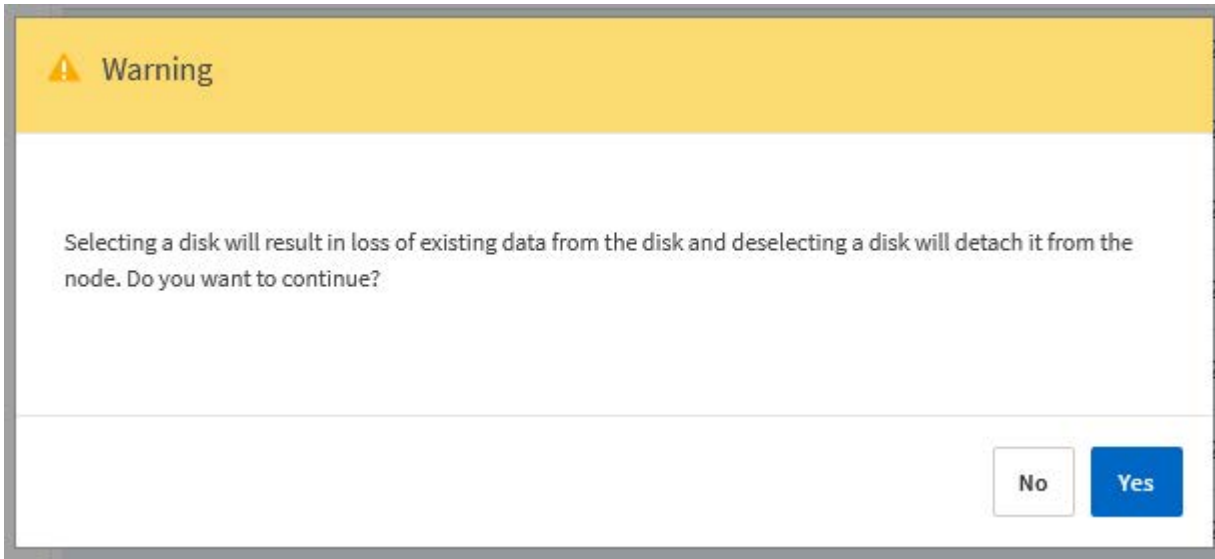
Cluster Password

••••••••

Cancel

Edit Storage


6. Confirm the operation.



Upgrade 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.


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. Click the **Administration** tab at the top of the page.
3. Click **Licenses** and click **Capacity Tier**.
4. Optionally click **Filter** and to limit the licenses displayed.
5. To replace an existing license; select a license, click , and select **Update**.
6. To add a new license, click **Add** at the top of the page and then click **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. Click the **Administration** tab at the top of the page.
3. Click **Licenses** and click **Capacity Pools**.
4. Optionally click **Filter** and to limit the licenses displayed.
5. Optionally select a license and click  to manage an existing license.
6. To add a new license or renew an existing license, click **Add** at the top of the page and then click **Upload License(s)** and select a license file from your local workstation.
7. To see a list of the capacity pools:
 - a. Click **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. Click **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. Click the **Administration** tab at the top of the page.
 - c. Click **Licenses** and then click **Capacity Pool**.
 - d. Click **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. Click the **Clusters** tab at the top of the page and select the desired cluster.
3. At the top of the cluster details page, click **Click here** to modify the cluster license.

You can also click **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 click **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, see 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

Storage: General concepts and characteristics

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 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.

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 ESX 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 ESX 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 nonzero (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, see 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 ESX 6.5 and newer.

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 vNVM driver available in ESX versions of 6.5 or later. This change is most significant for ONTAP Select installations with software RAID, which do not benefit from the RAID controller cache.

Software RAID 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).

Notes

- The terms service disk and system disk are used interchangeably.
 - Service disks are the 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:

- The hardware RAID controller must be disabled such that disks can be presented directly to the system (a JBOD). This change can usually be made 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 could be chosen to enable the JBOD mode. This enables a passthrough, so that the physical drives can be seen as is on the host.

Depending on maximum number of drives supported by the controller, an additional controller may be required. With the SAS HBA mode, ensure that the IO controller (SAS HBA) is supported with a minimum of 6Gb/s 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 – 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 virtual disks (VMDKs) and physical disks Raw Device Mappings [RDMs] for SSDs and passthrough or DirectPath IO devices for NVMeS.

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 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 multinode 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. Since a data disk has a single root partition, that the minimum number of data disks will vary depending on whether the ONTAP Select node is part of an HA pair or not.

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)

VMware ESX does not currently support NVMe disks as Raw Device Maps. For ONTAP Select to take direct control of NVMe disks, the NVMe drives must be configured in ESX as passthrough devices. Please note that configuring an NVMe device as a passthrough device requires support from the server BIOS and it is a disruptive process, requiring an ESX host reboot. Furthermore, the maximum number of passthrough devices per ESX host is 16. However, ONTAP Deploy limits this to 14. This limit of 14 NVMe devices per ONTAP Select node means that an all NVMe configuration will provide a very high IOPs density (IOPs/TB) at the expense of total capacity. Alternatively, if a high performance configuration with larger storage capacity is desired, 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 the section [Increase storage capacity](#).

Similar to FAS and AFF systems, only drives with equal or larger capacities can be added 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 ESX 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 shell, you can enter the following command to blink the LED for a given physical disk (identified by its naa.unique-id).

```
esxcli storage core device set -d <naa_id> -l=locator -L=<seconds>
```

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 proper 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.

```
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

Physical Size	Position	Disk	Type	Size
-	shared	NET-3.2	SSD	-
-	shared	NET-3.3	SSD	-
208.4GB	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

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

Physical Size	Status	Position	Disk	Pool	Type	RPM	Size
447.1GB	(normal)	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.

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 underlying configuration is supported by VMware and should be listed on the respective VMware HCLs.

vNAS architecture

The vNAS nomenclature is used for all setups that do not use DAS. For multinode 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. In other words, 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 de facto 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
```

Note that 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 multinode ONTAP Select clusters on shared storage. ONTAP Deploy enables the configuration of multiple ONTAP Select nodes on the same ESX host as long as these nodes are not part of the same cluster. Note that 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 multinode vNAS cluster does not place multiple ONTAP Select instances from the same cluster on the same host. The following figure shows for an example of a correct deployment of two four-node clusters that intersect on two hosts.

Initial deployment of multinode 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 ESX cluster.

See 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.

The screenshot shows the vSphere configuration interface for VM/Host Rules. The left sidebar contains a navigation tree with the following items: Services (vSphere DRS, vSphere Availability), vSAN (General, Disk Management, Fault Domains & Stretched Cluster, Health and Performance, iSCSI Targets, iSCSI Initiator Groups, Configuration Assist, Updates), Configuration (General, Licensing, VMware EVC, VMHost Groups, **VM/Host Rules**, VM Overrides, Host Options, Profiles, I/O Filters). The main content area is titled 'VM/Host Rules' and includes 'Add...', 'Edit...', and 'Delete' buttons. Below these is a table with columns: Name, Type, Enabled, Conflicts, and Defined By. The table is currently empty, displaying the message 'This list is empty.' At the bottom of the page, the text 'No VM/Host rule selected' is visible.

Name	Type	Enabled	Conflicts	Defined By
This list is empty.				



Two or more ONTAP Select nodes from the same ONTAP Select cluster could potentially be found on the same ESX 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.


Note that 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:



Increase 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 multinode 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.

Note that 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.

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
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

ONTAP Select features	DAS SSD (premium or premium XL ¹)	DAS HDD (all licenses)	vNAS (all licenses)
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

¹ONTAP 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

Networking: General 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.

Single and multiple node network configurations

ONTAP Select supports both single node and multinode 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 multinode 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

ONTAP Select can be deployed as a single-node cluster. The hypervisor host includes a virtual switch that provides access to the external network.

ESXi

The relationship between these ports and the underlying physical adapters can be seen in the following figure, which depicts one ONTAP Select cluster node on the ESX hypervisor.

Network configuration of single-node ONTAP Select cluster



i Even though two adapters are sufficient for a single-node cluster, NIC teaming is still required.

LIF assignment

As explained in the multinode LIF assignment section of this document, IPspaces are used by ONTAP Select 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, no ports are present in the cluster IPspace.

i Cluster and node management LIFs are automatically created during ONTAP Select cluster setup. The remaining LIFs can be created 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)

Multinode network configuration

The multinode 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 in allowing you to build 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- and eight-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 multinode 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, these ports should be placed 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 ESX hypervisor.

Network configuration of a single node that is part of a multinode ONTAP Select cluster



Segregating internal and external traffic across different physical NICs prevents latencies from being introduced into the system due to insufficient access to network resources. Additionally, aggregation through NIC teaming makes sure that failure of a single network adapter does not prevent the ONTAP Select cluster node from accessing the respective network.

Note that 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. The remaining LIFs can be created 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).



These interfaces cannot be assigned static IP addresses, and additional cluster interfaces should not be created.

Cluster network traffic must flow through a low-latency, nonrouted layer-2 network. Due to cluster throughput and latency requirements, the ONTAP Select cluster is expected to be physically located within proximity (for example, multipack, single data center). Building four-node, six-node, or eight-node stretch cluster configurations by separating HA nodes across a WAN or across significant geographical distances is not supported. A stretched two-node configuration with a mediator is supported.

For details, see 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 proper 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 the reassignment of an alternate IP address is not supported. 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.



In a manner similar to 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, the IP address of this interface cannot be modified, and the state of the port cannot be changed. This network port requires the use of jumbo frames (7500 to 9000 MTU).

ONTAP Select internal and external network

Characteristics of ONTAP Select internal and external networks.

ONTAP Select internal network

The internal ONTAP Select network, which is only present in the multinode 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 ESX standard and distributed vSwitches are supported. Other virtual switches or direct connectivity between ESX 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. Proper 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-node, six-node, and eight-node clusters require that the internal ONTAP Select network use 10Gb connectivity; 1Gb NICs are not supported. Tradeoffs 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 multinode 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 multinode 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 GUI) includes the internal network checker as an optional step available during the creation of multinode clusters. Given the important role that the internal network plays in multinode 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 multinode 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. See 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 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
4/6/8 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.



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 that 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 configuring the port groups to be used by ONTAP Select, the following best practices should be followed; 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, it is a NetApp best practice to make sure of hardware redundancy through NIC aggregation.

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 may 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, the desired number of physical ports can be added to the external port group as active adapters.

For multinode 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.

Standard or distributed vSwitch and four physical ports per Node

Four port groups can be assigned to each node in a multinode 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 GUI (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 vmnic 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. Four port groups should be used 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



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 multinode 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 multinode 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 proper 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 that you configure the upstream switch ports as a port channel. Prior to enabling this on the distributed vSwitch, make sure that an LACP-enabled port channel is properly configured.

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 multinode 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 multinode 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 multinode 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



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 multinode 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 (multinode 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

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 multinode 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 multinode clusters (four, six, or eight 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 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 and eight-node clusters are composed of three and four 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 multinode 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 for which nodes 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) and physical distance (a maximum of 10km).

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.

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, see the section [Two-node stretched HA \(MetroCluster SDS\) best practices](#).

Mirrored aggregates

An ONTAP Select cluster is composed of two to eight 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, 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:

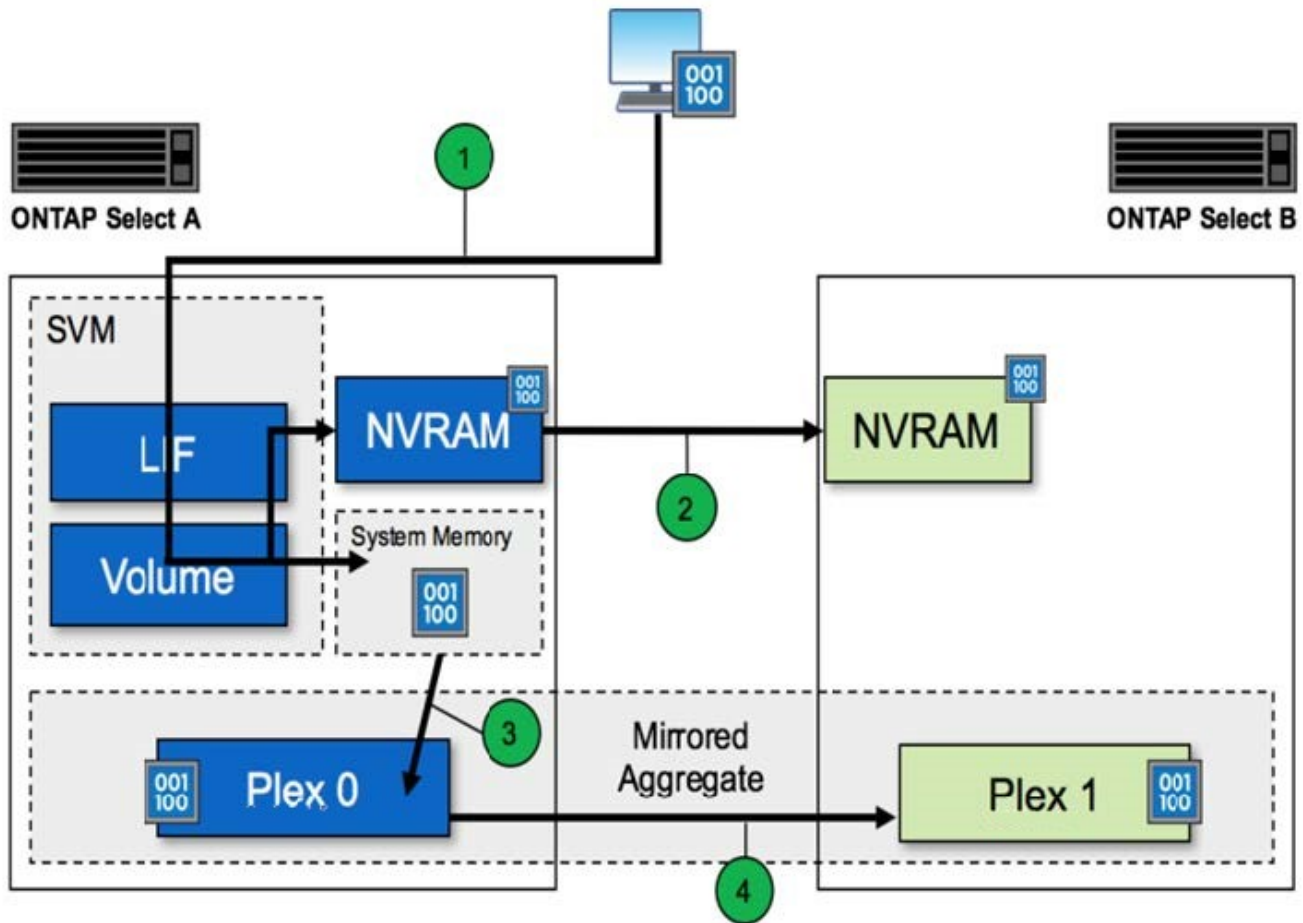
- Acknowledgment
- Destaging

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.

HA additional details

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 multinode 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>
```

Performance

Performance

Performance varies based on hardware configuration.

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 multsocket.** ONTAP Select does not use multsocket features, but the hypervisor overhead for supporting multsocket 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.**

Performance: Premium HA direct-attached SSD storage

Performance information for the reference platform.

Reference platform

ONTAP Select (Premium XL) hardware (per node)

- FUJITSU PRIMERGY RX2540 M4:
 - Intel® Xeon® Gold 6142b CPU at 2.6 GHz
 - 32 physical cores (16 x 2 sockets), 64 logical
 - 256 GB RAM
 - Drives per host: 24 960GB SSD
 - ESX 6.5U1

Client hardware

- 5 x NFSv3 IBM 3550m4 clients

Configuration information

- SW RAID 1 x 9 + 2 RAID-DP (11 drives)
- 22+1 RAID-5 (RAID-0 in ONTAP) / RAID cache NVRAM
- No storage efficiency features in use (compression, deduplication, Snapshot copies, SnapMirror, and so on)

The following table lists the throughput measured against read/write workloads on a high availability (HA) pair of ONTAP Select nodes using both software RAID and hardware RAID. Performance measurements were

taken using the SIO load-generating tool.



These performance numbers are based on ONTAP Select 9.6.

Performance results for a single node (part of a four-node medium instance) ONTAP Select cluster on a direct-attached storage (DAS) SSD, with software RAID and hardware RAID

Description	Sequential Read 64KiB	Sequential Write 64KiB	Random Read 8KiB	Random Write 8KiB	Random WR/RD (50/50) 8KiB
ONTAP Select large instance with DAS (SSD) software RAID	2171 MiBps	559 MiBps	954 MiBps	394 MiBps	564 MiBps
ONTAP Select medium instance with DAS (SSD) software RAID	2090 MiBps	592 MiBps	677 MiBps	335 MiBps	441 3MiBps
ONTAP Select medium instance with DAS (SSD) hardware RAID	2038 MiBps	520 MiBps	578 MiBps	325 MiBps	399 MiBps

64K sequential read

Details:

- SIO direct I/O enabled
- 2 nodes
- 2 x data NIC per node
- 1 x data aggregate per node (2TB hardware RAID), (8TB software RAID)
- 64 SIO procs, 1 thread per proc
- 32 volumes per node
- 1 x files per proc; files are 12000MB each

64K sequential write

Details:

- SIO direct I/O enabled
- 2 nodes
- 2 x data network interface cards (NICs) per node
- 1 x data aggregate per node (2TB hardware RAID), (4TB software RAID)
- 128 SIO procs, 1 thread per proc
- Volumes per node: 32 (hardware RAID), 16 (software RAID)
- 1 x files per proc; files are 30720MB each

8K random read

Details:

- SIO direct I/O enabled
- 2 nodes
- 2 x data NICs per node
- 1 x data aggregate per node (2TB hardware RAID), (4TB software RAID)
- 64 SIO procs, 8 threads per proc
- Volumes per node: 32
- 1 x files per proc; files are 12228MB each

8K random write

Details:

- SIO direct I/O enabled
- 2 nodes
- 2 x data NICs per node
- 1 x data aggregate per node (2TB hardware RAID), (4TB software RAID)
- 64 SIO procs, 8 threads per proc
- Volumes per node: 32
- 1 x files per proc; files are 8192MB each

8K random 50% write 50% read

Details:

- SIO direct I/O enabled
- 2 nodes
- 2 x data NICs per node
- 1 x data aggregate per node (2TB hardware RAID), (4TB software RAID)
- 64 SIO proc208 threads per proc
- Volumes per node: 32
- 1 x files per proc; files are 12228MB each

Automate with REST

Concepts

REST web services foundation

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 improved 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.

How to access the Deploy API

Because of the inherent flexibility of REST web services, the ONTAP Select Deploy API can be accessed in several different ways.

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.

Deploy API versioning

The REST API included with ONTAP Select Deploy is assigned a version number. The API version number is independent of the Deploy release number. You should be aware of the API version included with your release of Deploy and how this might affect your use of the API.

The current release of the Deploy administration utility includes version 3 of the REST API. Past releases of the Deploy utility include the following API versions:

Deploy 2.8 and later

ONTAP Select Deploy 2.8 and all later releases include version 3 of the REST API.

Deploy 2.7.2 and earlier

ONTAP Select Deploy 2.7.2 and all earlier releases include version 2 of the REST API.



Versions 2 and 3 of the REST API are not compatible. If you upgrade to Deploy 2.8 or later from an earlier release that includes version 2 of the API, you must update any existing code that directly accesses the API as well as any scripts using the command line interface.

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

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

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

Before you access the 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 click *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 click *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 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 press **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 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 click **Cluster**.
2. Click **GET /clusters/{cluster_id}** to display the details of the API call used to return information about an ONTAP Select cluster.

Workflow processes

Before you use the 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 click 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
```

Workflow 1: Create a 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
- Cluster is deployed in a production environment with a purchased license
- The KVM hypervisor is used instead of VMware ESXi



- 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.

1. Register vCenter server credential

When deploying to an ESXi host managed by a vCenter server, you must 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

2. Register a hypervisor host

You must 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

3. Create a cluster

When you create an ONTAP Select cluster, the basic cluster configuration is registered and the node names are automatically generated by Deploy.

Category	HTTP verb	Path
Cluster	POST	/clusters

Curl

The query parameter `node_count` should be set 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

4. Configure the cluster

There are several attributes you must provide as part of configuring the cluster.

Category	HTTP verb	Path
Cluster	PATCH	/clusters/{cluster_id}

Curl

You must 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

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 must retrieve the assigned ID.

Category	HTTP verb	Path
Cluster	GET	/clusters/{cluster_id}/nodes

Curl

You must 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

6. Configure the nodes

You must 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

You must 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)

You must 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

7. Retrieve the node networks

You must 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

You must 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

8. Configure the node networking

You must 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

You must 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)

You need to provide the name of the network.

```
{
  "name": "sDOT_Network"
}
```

Processing type

Synchronous

Output

None

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

You must 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

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

You must 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)

You must provide the password for the ONTAP administrator account.

```
{  
  "ontap_credentials": {  
    "password": "mypassword"  
  }  
}
```

Processing type

Asynchronous

Output

- Job object

Access with Python

Before you access the 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 properly:

- 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

Understand the Python scripts

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:



- 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.

- 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 a cluster

You can use the following script to create a cluster based on parameters defined within the script and a JSON input file.

```
1 #!/usr/bin/env python
2 ##-----
3 #
4 # File: 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
10 # merchantability
11 # or fitness of any kind, expressed or implied. Permission to use,
12 # reproduce, modify and create derivatives of the sample code is
13 # granted
14 # solely for the purpose of researching, designing, developing and
15 # testing a software application product for use with NetApp products,
16 # provided that the above copyright notice appears in all copies and
17 # that the software application product is distributed pursuant to
```

```

terms
16 # no less restrictive than those set forth herein.
17 #
18 ##-----
19
20 import traceback
21 import argparse
22 import json
23 import logging
24
25 from deploy_requests import DeployRequests
26
27
28 def add_vcenter_credentials(deploy, config):
29     """ Add credentials for the vcenter if present in the config """
30     log_debug_trace()
31
32     vcenter = config.get('vcenter', None)
33     if vcenter and not deploy.resource_exists('/security/credentials',
34                                             'hostname', vcenter
['hostname']):
35         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',
136     'instance_type',
137     'is_storage_efficiency_enabled'] if k
138     in node}
139     # Optional: Set a serial_number
140     if 'license' in node:
141         data['license'] = {'id': node['license']}
142
143     # Assign the host
144     host_id = deploy.find_resource('/hosts', 'name', node[
145     'host_name'])
146     if not host_id:
147         log_and_exit("Host names must match in the 'hosts' array, and
148         the nodes.host_name property")
149
150     data['host'] = {'id': host_id}
151
152     # Set the correct raid_type
153     is_hw_raid = not node['storage'].get('disks') # The presence of a
154     list of disks indicates sw_raid
155     data['passthrough_disks'] = not is_hw_raid
156
157     # Optionally set a custom node name
158     if 'name' in node:
159         data['name'] = node['name']
160
161     log_info("Node properties: {}".format(data))
162     deploy.patch('/clusters/{}/nodes/{}'.format(cluster_id, node_id),
163     data)
164
165
166 def add_node_networks(deploy, cluster_id, node_id, node):
167     ''' Set the network information for a node '''
168     log_debug_trace()
169
170     log_info("Adding node '{}' network properties".format(node_id))
171
172     num_nodes = deploy.get_num_records('/clusters/{}/nodes'.format
173     (cluster_id))
174
175     for network in node['networks']:
176
177         # 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
176     ('/clusters/{}/nodes/{}/networks'.format(cluster_id, node_id),
177      'purpose', network[
178      'purpose'])
179     data = {"name": network['name']}
180     if 'vlan' in network and network['vlan']:
181         data['vlan_id'] = network['vlan']
182
183     deploy.patch('/clusters/{}/nodes/{}/networks/{}'.format
184                 (cluster_id, node_id, network_id), data)
185
186
187 def add_node_storage(deploy, cluster_id, node_id, node):
188     ''' Set all the storage information on a node '''
189     log_debug_trace()
190
191     log_info("Adding node '{}' storage properties".format(node_id))
192     log_info("Node storage: {}".format(node['storage']['pools']))
193
194     data = {'pool_array': node['storage']['pools']} # use all the
195     json properties
196     deploy.post(
197         '/clusters/{}/nodes/{}/storage/pools'.format(cluster_id,
198         node_id), data)
199
200     if 'disks' in node['storage'] and node['storage']['disks']:
201         data = {'disks': node['storage']['disks']}
202         deploy.post(
203             '/clusters/{}/nodes/{}/storage/disks'.format(cluster_id,
204             node_id), data)
205
206
207
208 def create_cluster_config(deploy, config):
209     ''' Construct a cluster config in the deploy server using the
210     input json data '''
211     log_debug_trace()
212
213     cluster_id = add_cluster_attributes(deploy, config)
214
215     node_ids = get_node_ids(deploy, cluster_id)
216     node_configs = config['nodes']
217

```

```

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']
224 ]['ontap_admin_password']}]
225     deploy.post('/clusters/{}/deploy?inhibit_rollback=true'.format
226 (cluster_id),
227                 data, wait_for_job=True)
228
229
230 def log_debug_trace():
231     stack = traceback.extract_stack()
232     parent_function = stack[-2][2]
233     logging.getLogger('deploy').debug('Calling %s()' %
234 parent_function)
235
236
237 def log_info(msg):
238     logging.getLogger('deploy').info(msg)
239
240
241 def log_and_exit(msg):
242     logging.getLogger('deploy').error(msg)
243     exit(1)
244
245
246 def configure_logging(verbose):
247     FORMAT = '%(asctime)-15s:%(levelname)s:%(name)s: %(message)s'
248     if verbose:
249         logging.basicConfig(level=logging.DEBUG, format=FORMAT)
250     else:
251         logging.basicConfig(level=logging.INFO, format=FORMAT)
252     logging.getLogger('requests.packages.urllib3.connectionpool'
253 ).setLevel(
254     logging.WARNING)

```

```

252
253 def main(args):
254     configure_logging(args.verbose)
255     deploy = DeployRequests(args.deploy, args.password)
256
257     with open(args.config_file) as json_data:
258         config = json.load(json_data)
259
260         add_vcenter_credentials(deploy, config)
261
262         add_standalone_host_credentials(deploy, config)
263
264         register_unknown_hosts(deploy, config)
265
266         cluster_id = create_cluster_config(deploy, config)
267
268         deploy_cluster(deploy, cluster_id, config)
269
270
271 def parseArgs():
272     parser = argparse.ArgumentParser(description='Uses the ONTAP
273     Select Deploy API to construct and deploy a cluster.')
274     parser.add_argument('-d', '--deploy', help='Hostname or IP address
275     of Deploy server')
276     parser.add_argument('-p', '--password', help='Admin password of
277     Deploy server')
278     parser.add_argument('-c', '--config_file', help='Filename of the
279     cluster config')
280     parser.add_argument('-v', '--verbose', help='Display extra
281     debugging messages for seeing exact API calls and responses',
282     action='store_true', default=False)
283     return parser.parse_args()
284
285 if __name__ == '__main__':
286     args = parseArgs()
287     main(args)

```

JSON for script to create a 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

```

1 {
2   "hosts": [
3     {
4       "password": "mypassword1",
5       "name": "host-1234",
6       "type": "ESX",
7       "username": "admin"
8     }
9   ],
10
11  "cluster": {
12    "dns_info": {
13      "domains": ["lab1.company-demo.com", "lab2.company-demo.com",
14        "lab3.company-demo.com", "lab4.company-demo.com"
15      ],
16
17      "dns_ips": ["10.206.80.135", "10.206.80.136"]
18    },
19    "ontap_image_version": "9.7",
20    "gateway": "10.206.80.1",
21    "ip": "10.206.80.115",
22    "name": "mycluster",
23    "ntp_servers": ["10.206.80.183", "10.206.80.142"],
24    "ontap_admin_password": "mypassword2",
25    "netmask": "255.255.254.0"
26  },
27
28  "nodes": [
29    {
30      "serial_number": "3200000nn",
31      "ip": "10.206.80.114",
32      "name": "node-1",
33      "networks": [
34        {
35          "name": "ontap-external",
36          "purpose": "mgmt",
37          "vlan": 1234
38        },
39        {
40          "name": "ontap-external",
41          "purpose": "data",
42          "vlan": null
43        },
44        {
45          "name": "ontap-internal",
46          "purpose": "internal",

```

```

47     "vlan": null
48   }
49 ],
50 "host_name": "host-1234",
51 "is_storage_efficiency_enabled": false,
52 "instance_type": "small",
53 "storage": {
54   "disk": [],
55   "pools": [
56     {
57       "name": "storage-pool-1",
58       "capacity": 4802666790125
59     }
60   ]
61 }
62 }
63 ]
64 }

```

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



- 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.

```
{
  "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 a node license

You can use the following script to add a license for an ONTAP Select node.

```

1 #!/usr/bin/env python
2 ##-----
3 #
4 # File: add_license.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
10 # merchantability
11 # or fitness of any kind, expressed or implied. Permission to use,
12 # reproduce, modify and create derivatives of the sample code is
13 # granted
14 # 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
16 # no less restrictive than those set forth herein.
17 #
18 ##-----
19
20 import argparse
21 import logging
22 import json
23
24 from deploy_requests import DeployRequests
25
26
27 def post_new_license(deploy, license_filename):
28     log_info('Posting a new license: {}'.format(license_filename))
29
30     # Stream the file as multipart/form-data
31     deploy.post('/licensing/licenses', data={},
32                files={'license_file': open(license_filename, 'rb')})
33
34     # Alternative if the NLF license data is converted to a string.
35     # with open(license_filename, 'rb') as f:
36     #     nlf_data = f.read()
37     #     r = deploy.post('/licensing/licenses', data={},
38                          #     files={'license_file': (license_filename,
39 nlf_data)})
40
41 def put_license(deploy, serial_number, data, files):
42     log_info('Adding license for serial number: {}'.format
43             (serial_number))
44
45     deploy.put('/licensing/licenses/{}'.format(serial_number), data
46              =data, files=files)
47
48 def put_used_license(deploy, serial_number, license_filename,
49                     ontap_username, ontap_password):
50     ''' If the license is used by an 'online' cluster, a
51         username/password must be given. '''
52
53     data = {'ontap_username': ontap_username, 'ontap_password':
54            ontap_password}
55     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 a 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
12 # granted
13 # solely for the purpose of researching, designing, developing and
14 # testing a software application product for use with NetApp products,
15 # provided that the above copyright notice appears in all copies and
16 # that the software application product is distributed pursuant to
17 # terms
18 # no less restrictive than those set forth herein.
19 #
20 ##-----
21
20 import argparse
21 import json
22 import logging
23
24 from deploy_requests import DeployRequests
25
26 def find_cluster(deploy, cluster_name):
27     return deploy.find_resource('/clusters', 'name', cluster_name)
28
29
30 def offline_cluster(deploy, cluster_id):
31     # Test that the cluster is online, otherwise do nothing
32     response = deploy.get('/clusters/{}?fields=state'.format(
33         cluster_id))
34     cluster_data = response.json()['record']
35     if cluster_data['state'] == 'powered_on':
```

```

35     log_info("Found the cluster to be online, modifying it to be
powered_off.")
36     deploy.patch('/clusters/{}'.format(cluster_id), {
    'availability': 'powered_off'}, True)
37
38
39 def delete_cluster(deploy, cluster_id):
40     log_info("Deleting the cluster({}).".format(cluster_id))
41     deploy.delete('/clusters/{}'.format(cluster_id), True)
42     pass
43
44
45 def log_info(msg):
46     logging.getLogger('deploy').info(msg)
47
48
49 def configure_logging():
50     FORMAT = '%(asctime)-15s:%(levelname)s:%(name)s: %(message)s'
51     logging.basicConfig(level=logging.INFO, format=FORMAT)
52     logging.getLogger('requests.packages.urllib3.connectionpool'
    ).setLevel(logging.WARNING)
53
54
55 def main(args):
56     configure_logging()
57     deploy = DeployRequests(args.deploy, args.password)
58
59     with open(args.config_file) as json_data:
60         config = json.load(json_data)
61
62         cluster_id = find_cluster(deploy, config['cluster']['name'])
63
64         log_info("Found the cluster {} with id: {}".format(config
    ['cluster']['name'], cluster_id))
65
66         offline_cluster(deploy, cluster_id)
67
68         delete_cluster(deploy, cluster_id)
69
70
71 def parseArgs():
72     parser = argparse.ArgumentParser(description='Uses the ONTAP Select
    Deploy API to delete a cluster')
73     parser.add_argument('-d', '--deploy', required=True, type=str,
    help='Hostname or IP address of Deploy server')
74     parser.add_argument('-p', '--password', required=True, type=str,

```

```

    help='Admin password of Deploy server')
75     parser.add_argument('-c', '--config_file', required=True, type=str,
help='Filename of the cluster json config')
76     return parser.parse_args()
77
78 if __name__ == '__main__':
79     args = parseArgs()
80     main(args)

```

Common support module

All of the Python scripts use a common Python class in a single module.

```

1 #!/usr/bin/env python
2 ##-----
3 #
4 # File: deploy_requests.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
10 # merchantability
11 # or fitness of any kind, expressed or implied. Permission to use,
12 # reproduce, modify and create derivatives of the sample code is
13 # granted
14 # solely for the purpose of researching, designing, developing and
15 # testing a software application product for use with NetApp products,
16 # provided that the above copyright notice appears in all copies and
17 # that the software application product is distributed pursuant to
18 # terms
19 # no less restrictive than those set forth herein.
20 #
21 ##-----
22
23
24 import json
25 import logging
26 import requests
27
28 requests.packages.urllib3.disable_warnings()
29
30 class DeployRequests(object):
31     '''
32     Wrapper class for requests that simplifies the ONTAP Select Deploy
33     path creation and header manipulations for simpler code.

```

```

30     '''
31
32     def __init__(self, ip, admin_password):
33         self.base_url = 'https://{}/api'.format(ip)
34         self.auth = ('admin', admin_password)
35         self.headers = {'Accept': 'application/json'}
36         self.logger = logging.getLogger('deploy')
37
38     def post(self, path, data, files=None, wait_for_job=False):
39         if files:
40             self.logger.debug('POST FILES:')
41             response = requests.post(self.base_url + path,
42                                     auth=self.auth, verify=False,
43                                     files=files)
44         else:
45             self.logger.debug('POST DATA: %s', data)
46             response = requests.post(self.base_url + path,
47                                     auth=self.auth, verify=False,
48                                     json=data,
49                                     headers=self.headers)
50
51             self.logger.debug('HEADERS: %s\nBODY: %s', self.
filter_headers(response), response.text)
52             self.exit_on_errors(response)
53
54             if wait_for_job and response.status_code == 202:
55                 self.wait_for_job(response.json())
56             return response
57
58     def patch(self, path, data, wait_for_job=False):
59         self.logger.debug('PATCH DATA: %s', data)
60         response = requests.patch(self.base_url + path,
61                                  auth=self.auth, verify=False,
62                                  json=data,
63                                  headers=self.headers)
64         self.logger.debug('HEADERS: %s\nBODY: %s', self.
filter_headers(response), response.text)
65         self.exit_on_errors(response)
66
67         if wait_for_job and response.status_code == 202:
68             self.wait_for_job(response.json())
69         return response
70
71     def put(self, path, data, files=None, wait_for_job=False):
72         if files:
73             print('PUT FILES: {}'.format(data))

```



```

74         response = requests.put(self.base_url + path,
75                                 auth=self.auth, verify=False,
76                                 data=data,
77                                 files=files)
78     else:
79         self.logger.debug('PUT DATA:')
80         response = requests.put(self.base_url + path,
81                                 auth=self.auth, verify=False,
82                                 json=data,
83                                 headers=self.headers)
84
85         self.logger.debug('HEADERS: %s\nBODY: %s', self.
filter_headers(response), response.text)
86         self.exit_on_errors(response)
87
88         if wait_for_job and response.status_code == 202:
89             self.wait_for_job(response.json())
90         return response
91
92     def get(self, path):
93         """ Get a resource object from the specified path """
94         response = requests.get(self.base_url + path, auth=self.auth,
verify=False)
95         self.logger.debug('HEADERS: %s\nBODY: %s', self.
filter_headers(response), response.text)
96         self.exit_on_errors(response)
97         return response
98
99     def delete(self, path, wait_for_job=False):
100         """ Delete's a resource from the specified path """
101         response = requests.delete(self.base_url + path, auth=self
.auth, verify=False)
102         self.logger.debug('HEADERS: %s\nBODY: %s', self.
filter_headers(response), response.text)
103         self.exit_on_errors(response)
104
105         if wait_for_job and response.status_code == 202:
106             self.wait_for_job(response.json())
107         return response
108
109     def find_resource(self, path, name, value):
110         ''' Returns the 'id' of the resource if it exists, otherwise
None '''
111         resource = None
112         response = self.get('{path}?{field}={value}'.format(
113             path=path, field=name, value=value))

```

```

114         if response.status_code == 200 and response.json().get
('num_records') >= 1:
115             resource = response.json().get('records')[0].get('id')
116             return resource
117
118     def get_num_records(self, path, query=None):
119         ''' Returns the number of records found in a container, or
None on error '''
120         resource = None
121         query_opt = '{}?{}'.format(path, query) if query else ''
122         response = self.get('{path}{query}'.format(path=path, query
=query_opt))
123         if response.status_code == 200 :
124             return response.json().get('num_records')
125         return None
126
127     def resource_exists(self, path, name, value):
128         return self.find_resource(path, name, value) is not None
129
130     def wait_for_job(self, response, poll_timeout=120):
131         last_modified = response['job']['last_modified']
132         job_id = response['job']['id']
133
134         self.logger.info('Event: ' + response['job']['message'])
135
136         while True:
137             response = self.get('/jobs/{}?fields=state,message&
'poll_timeout={}&last_modified=>={}'
138
.format(
139                 job_id, poll_timeout,
last_modified))
140
141             job_body = response.json().get('record', {})
142
143             # Show interesting message updates
144             message = job_body.get('message', '')
145             self.logger.info('Event: ' + message)
146
147             # Refresh the last modified time for the poll loop
148             last_modified = job_body.get('last_modified')
149
150             # Look for the final states
151             state = job_body.get('state', 'unknown')
152             if state in ['success', 'failure']:
153                 if state == 'failure':
154                     self.logger.error('FAILED background job.\nJOB:

```

```

    %s', job_body)
155             exit(1)    # End the script if a failure occurs
156             break
157
158     def exit_on_errors(self, response):
159         if response.status_code >= 400:
160             self.logger.error('FAILED request to URL: %s\nHEADERS:
%s\nRESPONSE BODY: %s',
161                               response.request.url,
162                               self.filter_headers(response),
163                               response.text)
164             response.raise_for_status()    # Displays the response error,
and exits the script
165
166     @staticmethod
167     def filter_headers(response):
168         ''' Returns a filtered set of the response headers '''
169         return {key: response.headers[key] for key in ['Location',
'request-id'] if key in response.headers}

```

Script to resize 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
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
16 # no less restrictive than those set forth herein.
17 #
18 ##-----

```

```

19
20 import argparse
21 import logging
22 import sys
23
24 from deploy_requests import DeployRequests
25
26
27 def _parse_args():
28     """ Parses the arguments provided on the command line when
29     executing this
30     script and returns the resulting namespace. If all required
31     arguments
32     are not provided, an error message indicating the mismatch is
33     printed and
34     the script will exit.
35     """
36     parser = argparse.ArgumentParser(description=(
37         'Uses the ONTAP Select Deploy API to resize the nodes in the
38         cluster.'
39         ' For example, you might have a small (4 CPU, 16GB RAM per
40         node) 2 node'
41         ' cluster and wish to resize the cluster to medium (8 CPU,
42         64GB RAM per'
43         ' node). This script will take in the cluster details and then
44         perform'
45         ' the operation and wait for it to complete.'
46     ))
47     parser.add_argument('--deploy', required=True, help=(
48         'Hostname or IP of the ONTAP Select Deploy VM.'
49     ))
50     parser.add_argument('--deploy-password', required=True, help=(
51         'The password for the ONTAP Select Deploy admin user.'
52     ))
53     parser.add_argument('--cluster', required=True, help=(
54         'Hostname or IP of the cluster management interface.'
55     ))
56     parser.add_argument('--instance-type', required=True, help=(
57         'The desired instance size of the nodes after the operation is
58         complete.'
59     ))
60     parser.add_argument('--ontap-password', required=True, help=(
61         'The password for the ONTAP administrative user account.'
62     ))
63     parser.add_argument('--ontap-username', default='admin', help=(

```

```

57     'The username for the ONTAP administrative user account.
    Default: admin.'
58     ))
59     parser.add_argument('--nodes', nargs='+', metavar='NODE_NAME',
    help=(
60         'A space separated list of node names for which the resize
    operation'
61         ' should be performed. The default is to apply the resize to
    all nodes in'
62         ' the cluster. If a list of nodes is provided, it must be
    provided in HA'
63         ' pairs. That is, in a 4 node cluster, nodes 1 and 2
    (partners) must be'
64         ' resized in the same operation.'
65     ))
66     return parser.parse_args()
67
68
69 def _get_cluster(deploy, parsed_args):
70     """ Locate the cluster using the arguments provided """
71
72     cluster_id = deploy.find_resource('/clusters', 'ip', parsed_args
    .cluster)
73     if not cluster_id:
74         return None
75     return deploy.get('/clusters/%s?fields=nodes' % cluster_id).
    json()['record']
76
77
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())

```

Automate ONTAP Select deployments with Ansible

Use Ansible roles and playbooks to automate ONTAP Select deployments.

An Ansible role is a logical collection of tasks, templates, handlers, and variables in a standardized file structure. Use these roles to facilitate, reuse, and modularize functionality that can be independently used. Include roles, and the list of hosts where you want them to run, in a playbook for execution. After you install Ansible, update the necessary modules, and understand playbooks, you are ready to download ONTAP Select (OTS) roles from NetApp GitHub to create and run your own playbook to deploy ONTAP Select clusters.

Roles

There are two roles available for ONTAP Select:

na_ots_deploy

This role downloads the installation file for ONTAP Select Deploy onto a vCenter or ESXi host. It then creates and runs the Deploy VM.

This role uses the following input values:

- vCenter or ESXi host credentials
- Deploy VM creation, including information such as:
 - IP address
 - Host name
 - Login
 - Password
 - Datastore
 - Network
- Local path to the OVA file

Make sure the installation file is available before running the playbook including the role).



The simplest way to provide these input values is through a global variables file in YML format. Pass this YML file in the playbook.

The role has a single task that imports the Deploy OVA file onto the host, creates the VM, and runs it. You can access and download this role from the NetApp Ansible GitHub here: [na_ots_deploy](#)

na_ots_cluster

This role uses the ONTAP Select Deploy APIs to create and deploy the ONTAP Select cluster. It assumes that the Deploy VM has been created and is up and running either manually or using the **na_ots_deploy** role. Also, the role assumes that the Select hosts are appropriately configured with the networking and storage required for ONTAP Select cluster.

The role uses the input values for the vCenter or ESXI credentials, and Deploy VM credentials to access the APIs and all the pertaining information required to create the ONTAP Select cluster and the node VM.

The role performs the following tasks:

- Adding (vcenter or esxi) host credentials
- Getting and registering the host IDs
- Validating the internal network (for multi-node clusters)
- Creating the cluster
- Configuring the nodes
- Configuring the network and storage pool
- Deploying the cluster

As a result, the role completes with the cluster VM up and running and the ONTAP Select cluster fully deployed. You can access and download this role from the NetApp Ansible GitHub here: [na_ots_cluster](#)

Example Playbook

The following is an example playbook that calls these two ONTAP Select roles. Note that the input variables are defined in YAML files and passed in via “vars_files”. You can find more details in the README files in each of the roles.


```

-
- name: Create ONTAP Select deploy from OVA (ESXi)

vars_files:

- vars_deploy.yml # All Variables

- vars_deploy_pwd.yml # host_password &
  deploy_password

hosts: "{{ target_vcenter_or_esxi_host }}" # Entry in Ansible 'hosts'
file

gather_facts: false

connection: 'local'

roles:

- na_ots_deploy

- name: Create ONTAP Select Cluster

vars_files:

- vars_cluster_test.yml

- vars_cluster_pwd.yml

hosts: "localhost"

gather_facts: false

roles:

- na_ots_cluster

```

After the ONTAP Select cluster is created and running using the above roles, you can leverage the rich set of ONTAP Ansible roles available to further automate ONTAP features and functionality. The available ONTAP roles can be accessed [here](#). In summary, the ONTAP Select Ansible roles along with ONTAP roles let you fully automate your end-to-end workflow to manage your storage with ONTAP Select.

Use the CLI

Sign in to 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 press **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.

Before you begin

Before creating an ONTAP Select cluster on a hypervisor, you should understand the required preparation.

Prepare to attach storage to the ONTAP Select node

If you use a local hardware RAID controller, you must create at least one storage pool 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.

If you use software RAID, you must create a storage pool 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.

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.

See [Add an ONTAP Select image to Deploy](#) for more information.

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.

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 FILENAME
```

Provide the administrator account password when prompted.

4. Display the licenses in the system to confirm the license was added properly:

```
license show
```

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, 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. 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
```

Creating and configuring 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

You must 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 CLUSTERNAME -node-count NODES
```

Example output

```
cluster create -name test-cluster -node-count 1
```

3. Configure the cluster:

```
cluster modify -name CLUSTERNAME -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 CLUSTERNAME -detailed
```

Configure an ONTAP Select node

You must configure each of the nodes in the ONTAP Select cluster.

Before you begin

You must have the configuration information for the node. The capacity tier license file should be 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 CLUSTERNAME
```

3. Select the node and perform basic configuration:

```
node modify -name NODENAME -cluster-name CLUSTERNAME -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:

```
node modify -name NODENAME -cluster-name CLUSTERNAME -mgmt-ip IP -management  
-networks NETWORK_NAME -data-networks NETWORK_NAME -internal-network  
NETWORK_NAME
```

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 NODENAME -cluster-name CLUSTERNAME -detailed
```

Example output

```
node show -name test-cluster-01 -cluster-name test-cluster -detailed
```

Attach storage to the ONTAP Select nodes

You must 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

You must 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 POOLNAME -cluster-name CLUSTERNAME -node-name NODENAME -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 CLUSTERNAME -node-name NODENAME
```

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 NODENAME -cluster-name CLUSTERNAME -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 NODENAME -cluster-name CLUSTERNAME
```

Example output

```
node storage disk show -node-name sdot-smicro-009a -cluster-name NVME
```

Deploy an ONTAP Select cluster

After the cluster and nodes have been configured, you can deploy the cluster.

Before you begin

Before deploying a multi-node cluster, you should run the network connectivity checker 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 CLUSTERNAME
```

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 CLUSTERNAME
```

After you finish

You should back up the ONTAP Select Deploy configuration data.

Security

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 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 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 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.

ONTAP Select clusters

There are several related tasks you can perform to administer an ONTAP Select cluster.

Delete an ONTAP Select clusters

You can delete an ONTAP Select cluster when it is no longer needed using the command line interface.

About this task

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 CLUSTERNAME
```

3. If the cluster is not offline, move it to an offline state:

```
cluster offline -name CLUSTERNAME
```

4. After confirming the cluster is in an offline status, delete the cluster:

```
cluster delete -name CLUSTERNAME
```

Nodes and hosts

Upgrade VMware ESXi to version 7.0 or later

If you are running ONTAP Select on VMware ESXi, you can upgrade the ESXi software from an earlier supported version to ESXi 7.0 or later. Before upgrading, you should understand the process and select the appropriate upgrade procedure.

Before you begin

Before upgrading the ESXi software on the hypervisors hosting an ONTAP Select cluster, you should prepare and select the upgrade procedure that is appropriate for your environment.



If you choose to upgrade to VMware ESXi 6.5, you should upgrade to ESXi U2 (build 8294253) or greater. Using ESXi 6.5 U1 can expose you to a virtual machine failure due to a known VMware bug.

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.

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 may be instances when Deploy is unavailable or cannot be used. For example, upgrading to ESXi 7.0 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, the ONTAP Select virtual machine can be left in a state where it cannot be booted. In this case, you must select an upgrade procedure that does not use Deploy. Refer to [1172198](#) for more information.

Upgrade the Deploy administration utility

Before performing an upgrade procedure using the Deploy utility, you may 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.

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.

Example

```
node stop --cluster-name <CLUSTERNAME> --node-name <NODENAME>
```

3. Upgrade the hypervisor host where ONTAP Select is running to ESXi 7.0 or later using the procedure provided by VMware.
4. Move the node to the online state.

Example

```
node start --cluster-name <CLUSTERNAME> --node-name <NODENAME>
```

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.

Example

```
node stop --cluster-name <CLUSTERNAME> --node-name <NODENAME>
```

3. Upgrade the hypervisor host where ONTAP Select is running to ESXi 7.0 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.

Example

```
node start --cluster-name <CLUSTERNAME> --node-name <NODENAME>
```

5. After the node comes up, verify that storage failover is enabled and the cluster is healthy.

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 7.0 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 7.0 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.

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

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
<code>-name <i>name</i></code>	The IP address or FQDN of the host you want to modify.
<code>-mgmt-server <i>management_server</i></code>	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 <code>credential add</code> command.

Optional parameters

Parameter	Description
<code>-help</code>	Displays the help message.
<code>-foreground</code>	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.
<code>-username <i>username</i></code>	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 ESX host managed by a vCenter).

Deploy utility

Upgrade a Deploy instance

You can upgrade an existing Deploy utility virtual machine in-place using the command line interface.

Before you begin

Make sure that Deploy is not used to perform any other tasks during the upgrade. You should see the current release notes for information and restrictions about upgrading the Deploy utility.



If you have an older instance of the ONTAP Select Deploy administration utility installed, you should upgrade to the current release. The ONTAP Select node and ONTAP Select Deploy component are upgraded independently. See [Upgrade the ONTAP Select nodes](#) for further details.

Download the upgrade package

To begin the upgrade process, you must download the appropriate Deploy virtual machine upgrade file from the NetApp Support Site. The upgrade package is formatted as a single compressed file.

Steps

1. Access the [NetApp Support Site](#) using a web browser and choose **Downloads** from the Downloads menu.
2. Scroll down and select **ONTAP Select Deploy Upgrade**.
3. Select the desired release of the upgrade package.
4. Review the End User License Agreement (EULA) and select **Accept & Continue**.
5. Select and download the appropriate package, responding to all prompts as needed for your environment.

Upload the package to the Deploy virtual machine

After acquiring the upgrade package, you must upload the file to the Deploy virtual machine.

Before you begin

You must have the upgrade file available on your local workstation. You must also have the password for the 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.

Steps

1. In a command shell on your local workstation, use the scp utility to upload the image file to the Deploy virtual machine.

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.

Apply the upgrade package

After the upgrade file has been uploaded to the Deploy virtual machine, you can apply the upgrade.

Before you begin

You must know which directory the upgrade file has been placed in at the Deploy utility virtual machine. Also, assure that Deploy is not used to perform any other tasks while the upgrade is performed.

Steps

1. Sign in to the 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 FILEPATH
```

Example

```
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 Deploy virtual machine configuration. Also, you should clear the browser cache so you can view the newly created Deploy pages.

Migrate a 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.

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.

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.

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](#)

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.

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 PATHNAME -filename FILENAME
```

Example

```
deploy backup restore -path /home/admin -filename  
deploy_backup_20180601162151.tar.gz
```

Add an ONTAP Select image to Deploy

You can 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

At a high level, the process used to add an ONTAP Select image to an instance of Deploy consists of four steps:

1. Downloading the install image
2. Uploading the install image to the Deploy virtual machine
3. Adding the install image
4. Displaying the available install images

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.

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 using a web browser and click **Support Quick Links**.
2. Click **Download Software** under **Top Tasks** and sign in to the site.
3. Click **Find your product**.
4. Scroll down and click **ONTAP Select**.
5. Under **Other Available Select Software** click **Deploy Upgrade, Node Upgrade, Image Install**.
6. Select the desired release of the upgrade package.
7. Review the End User License Agreement (EULA) and click **Accept & Continue**.
8. Select and download the appropriate package, responding to all prompts as needed for your environment.

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

You must 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.

Example

```
scp image_v_93_install_esx.tgz admin@10.234.81.101:/home/admin (provide password when prompted)
```

Example

```
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.

Add the install image

You can 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.

Example

```
tar -xf image_v_93_install_esx.tgz -C /opt/netapp/images/
```

Display the available install images

You can 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 click **Deploy** and then click **GET /images**.
3. Click **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 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.

Before you begin

You should prepare before attempting to recover an instance of the Deploy utility to assure success.

Required skills and information

You should be familiar with several administrative procedures and have the required information.

Installing the Deploy virtual machine

You must be able to install a new instance of the ONTAP Select Deploy utility in your hypervisor environment.

ONTAP command line interface

You must be able to sign in to the ONTAP CLI of the ONTAP Select cluster and use the shell interface.

Availability of Deploy utility configuration backup

You must 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.

Restoring a backup of the Deploy configuration

You should be able to restore a backup of the Deploy configuration data, depending on the recovery procedure used.

IP address of the original Deploy virtual machine

You must know the IP address of the original Deploy utility virtual machine that failed.

Storage capacity licensing

You must 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.

Deciding which recovery procedure to use

You must 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?	Recovery procedure to use
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

You must have a backup of 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, you must 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, you must reinstall each capacity pool license. See *Reinstalling 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, you must configure the mediator iSCSI target and mailbox in the new Deploy virtual machine. You must then complete the recovery by performing additional configuration of the two nodes in the ONTAP Select cluster.

Before you begin

You must 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. Click **Mediator** and then **GET /mediators**.
7. Click **Try it out!** to display a list of mediators maintained by Deploy.

Note the ID of the desired mediator instance.

8. Click **Mediator** and then **POST**.
9. Provide the value for mediator_id.
10. Click the **Model** next to `iscsi_target` and complete the name value.

Use the target name for the `iqn_name` parameter.

11. Click **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.

1. Determine the names of the mediator disks:

```
disk show -container-type mediator
```

2. 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>
```

3. Verify that storage failover is enabled:

```
storage failover show
```

After you finish

If you use capacity pools licensing, you must reinstall each capacity pool license. See [Reinstalling a capacity pool license](#) for additional details.

Related information

- [Install ONTAP Select Deploy](#)
- [Restore the Deploy configuration data to the new virtual machine](#)
- [Reinstall a capacity pool license](#)

Frequently asked questions

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 (ESX) 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 ESX uses VMXNET3 network drivers and, when possible, the vNMVE 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, since 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 GUI, 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, see 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 ESX and KVM?

Yes. ONTAP Select Deploy can be installed on either KVM or ESX, and both installations can create ONTAP Select clusters on either hypervisor.

Is vCenter required for ONTAP Select on ESX?

If the ESX hosts are properly licensed, then there is no need for the ESX 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 ESX 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 ESX 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, contact NetApp Technical Support.

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 ESX 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.

Multinode HA clusters on shared storage are supported.

Does ONTAP Select support multinode clusters on vSAN or other shared external storage including some HCI stacks?

Multinode clusters using external storage (multinode vNAS) are supported for both ESX 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 multinode cluster brings in the benefits of ONTAP nondisruptive operation 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 ESX 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 ESX 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 GUI 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 GUI 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 multinode 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 ESX 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 ESX 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 and/or the ESX host. The VMware documentation states that the initial connection to either a vCenter Server or an ESX 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 ESX 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 ESX 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 multinode 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, NetApp Private Storage for Cloud, and NetApp FlexArray® technology.

How is the MetroCluster SDS different from NetApp MetroCluster?

MetroCluster SDS provides a synchronous replication solution and falls under NetApp MetroCluster solutions. However, the key differences are in the distances supported (~10km versus 300km), and 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 whose nodes 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 a set of minimum requirements in terms of latency (5ms RTT and 5ms jitter for a max total of 10ms) and physical distance (10km).

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 ESX.

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 and maximum physical distance of 10km between the ONTAP Select nodes.
- 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 multinode 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
```

Please note that 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 ESX 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

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

The notice file provides information about third-party copyright and licenses used in NetApp software.

- [Notice for ONTAP Select 9.14.1](#)

Copyright information

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

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

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

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

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

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

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

Trademark information

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