



Plan

ONTAP Select

NetApp
February 12, 2024


Table of Contents

- Plan 1
 - ONTAP Select installation and deployment workflow 1
 - ONTAP Select 1
 - ONTAP Select Deploy 17
 - Summary of best practices 19

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.

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 hosts

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 9.10.1, the KVM hypervisor is no longer supported when deploying a new cluster. You can still use KVM with an earlier supported release.

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

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.

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:

- VMware ESXi 7.0 GA (build 15843807 or greater) including 7.0 U1 and U2.



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.



An ESXi 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 ESXi 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 ESXi 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 VMware ESXi 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 of before deploying a cluster using software RAID.

General requirements

The environment for a software RAID deployment on VMware ESXi 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.

- 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.

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.

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.

ONTAP Select vNAS requirements

ONTAP Select vNAS is a solution allowing the ONTAP Select datastores to be external to the hypervisor host where the ONTAP Select virtual machine runs. These remote datastores 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

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

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

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

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
Either a vSphere server or ESXi host, depending on how the host is managed in your cluster deployment
- Hypervisor server
ESXi hypervisor 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	Direction	Description
TLS (443)	Deploy to vCenter server (managed) or ESXi (unmanaged)	VMware VIX API
902	Deploy to vCenter server (managed) or ESXi (unmanaged)	VMware VIX API
ICMP	Deploy to hypervisor server	Ping
ICMP	Deploy to each OTS node	Ping
SSH (22)	Admin WS to each OTS node	Administration
TLS (443)	Deploy to OTS nodes and clusters	Access ONTAP
TLS (443)	Each OTS node to Deploy	Access Deploy
iSCSI (3260)	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, you must assure 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 9.10.1, 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.

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

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.

Datstore

- Allocate space
- Browse datstore
- 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.

VMware vCenter plug-in

When deploying ONTAP Select in a VMware environment, you can use the VMware vCenter plug-in instead of the Deploy administration utility.

The VMware vCenter plug-in is packaged with the Deploy administration utility. To use the plug-in, you must install the Deploy utility and then use the Deploy user interface to install the plug-in to your vCenter server. After the plug-in is installed, you can use the vSphere web user interface to deploy and administer ONTAP Select clusters. Aside from managing the plug-in itself (installing, removing, updating), the functionality provided with the plug-in is equivalent to the Deploy utility user interface.

See the [Integrate with vCenter](#) section for the requirements and restrictions as well as information about installing and using the plug-in.

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.

Authentication using the credential store

The ONTAP Select Deploy credential store is a database 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.

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.

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.

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.

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.

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.