



Networking

ONTAP Select

NetApp
January 31, 2025

Table of Contents

- Networking 1
 - Networking: General concepts and characteristics 1
 - Single and multiple node network configurations 3
 - ONTAP Select internal and external network 8
 - Supported network configurations 10
 - VMware vSphere vSwitch Configuration on ESXi 12
 - Physical switch configuration 21
 - Data and management traffic separation 23

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



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.



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.

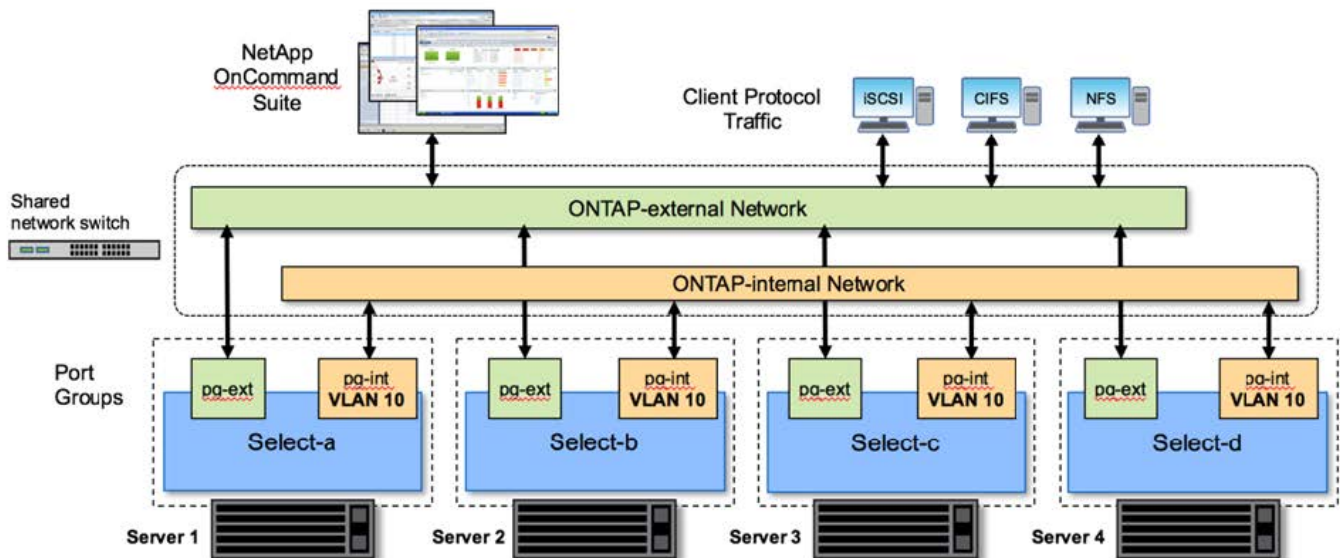
These are an internal network, responsible for providing cluster and internal replication services, and an external network, responsible for providing data access and management services. End-to-end isolation of traffic that flows within these two networks is extremely important in allowing you to build 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.




ONTAP Select supports a maximum of four VMNICs.

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 (vmnics 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.

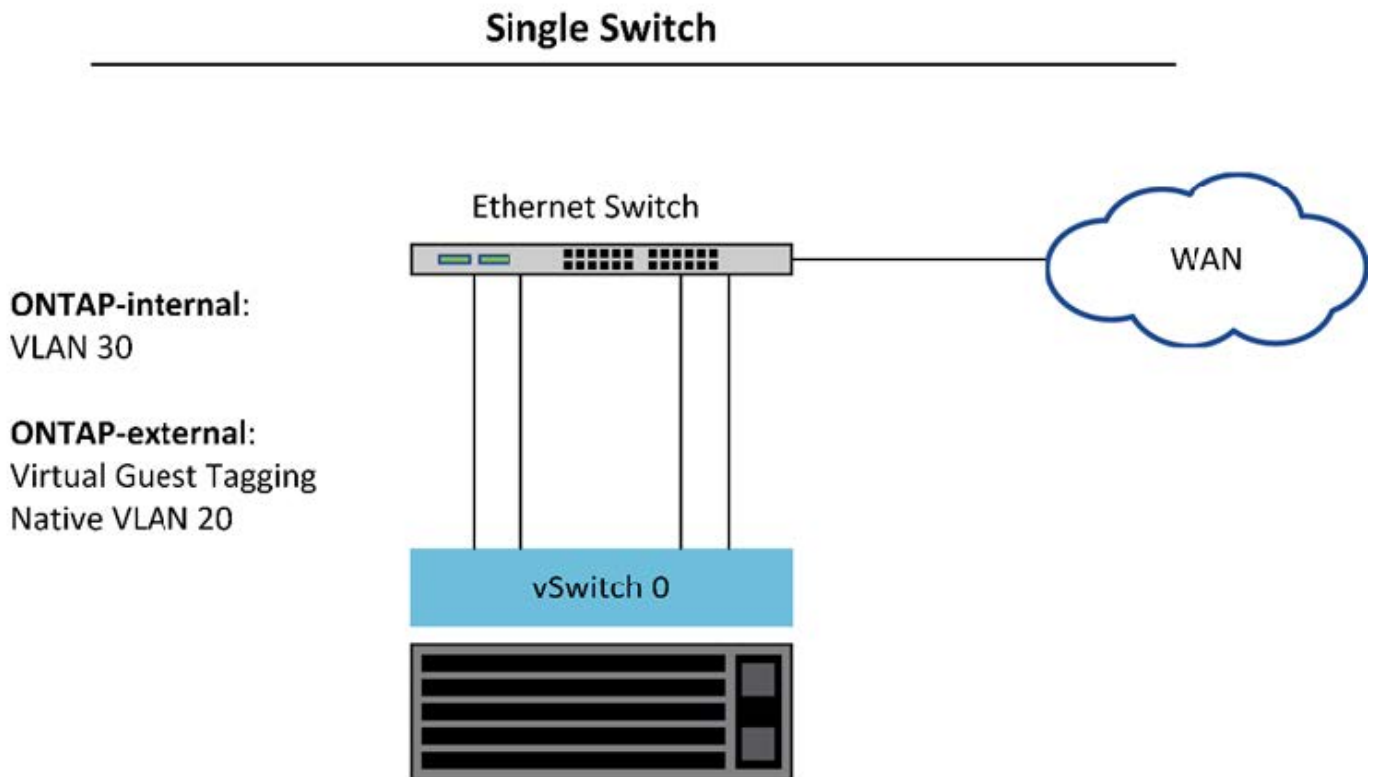
Shared physical switch

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

Ethernet Switch



Port Group 1 – No tagging at Port Group Level
Management traffic

VLAN 10 (VGT)

Data traffic


VLAN 20 (VGT)

Port Group 2

Cluster traffic

VLAN 30 (VST)

Default IPSpace



Copyright information

Copyright © 2025 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.