



Pre-installation checklist

ONTAP Select

NetApp
November 21, 2024

Table of Contents

- Pre-installation checklist 1
- Host configuration and preparation checklist 1
- Required information for Deploy utility installation 11
- Required information for ONTAP Select installation 11
- Configuring a host to use NVMe drives 12

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 Expres...	
<input checked="" type="checkbox"/> 0000:38:...	Available (pending)	Seagate Technology ...	Nytro Flash Storage	
0000:36:02.0	Not Configurable	Intel Corporation	Sky Lake-E PCI Expres...	
<input checked="" type="checkbox"/> 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 REFRESH EDIT...

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

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

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.

ONTAP System Manager (Return to classic version) Search actions, objects, and pages

ots-nvme Version 9.7.0

Health →

✓ All systems are healthy


FDvM300

Capacity →

The system discovered 6 disks. When you prepare the disk for provisioning, the system will group the disks for optimum performance and resiliency.

Prepare Storage

Preparing Local Storage.

 The local storage is being prepared.

DASHBOARD

STORAGE

NETWORK

EVENTS & JOBS

PROTECTION

HOSTS

CLUSTER

ONTAP System Manager (Return to classic version) Search actions, objects, and pages

ots-nvme Version 9.7.0

Health →

✓ All systems are healthy


FDvM300

Capacity →

0 Bytes USED | **4.82 TB** AVAILABLE

0% 20% 40% 60% 80% 100%

1 to 1 Data Reduction

 No cloud tier

DASHBOARD

STORAGE

NETWORK

EVENTS & JOBS

PROTECTION

HOSTS

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.