



Preparation

ONTAP 9

NetApp
February 20, 2026

Table of Contents

- Preparation 1
 - Assess ONTAP NFS physical storage requirements 1
 - Assess ONTAP NFS network configuration requirements 1
 - Learn about ONTAP NFS storage capacity provisioning 3
- ONTAP NFS configuration worksheet 3
 - Configure NFS access to an SVM 3
 - Name service information 6
 - Adding storage capacity to an NFS-enabled SVM 10

Preparation

Assess ONTAP NFS physical storage requirements

Before provisioning NFS storage for clients, you must ensure that there is sufficient space in an existing aggregate for the new volume. If there is not, you can add disks to an existing aggregate or create a new aggregate of the desired type.

Steps

1. Display available space in existing aggregates:

```
storage aggregate show
```

If there is an aggregate with sufficient space, record its name in the worksheet.

```
cluster::> storage aggregate show
Aggregate      Size Available Used% State  #Vols  Nodes  RAID Status
-----
aggr_0        239.0GB   11.13GB   95% online    1 node1  raid_dp,
normal
aggr_1        239.0GB   11.13GB   95% online    1 node1  raid_dp,
normal
aggr_2        239.0GB   11.13GB   95% online    1 node2  raid_dp,
normal
aggr_3        239.0GB   11.13GB   95% online    1 node2  raid_dp,
normal
aggr_4        239.0GB   238.9GB   95% online    5 node3  raid_dp,
normal
aggr_5        239.0GB   239.0GB   95% online    4 node4  raid_dp,
normal

6 entries were displayed.
```

2. If there are no aggregates with sufficient space, add disks to an existing aggregate by using the `storage aggregate add-disks` command, or create a new aggregate by using the `storage aggregate create` command.

Related information

- [Add disks to a local tier \(aggregate\)](#)
- [storage aggregate add-disks](#)
- [storage aggregate create](#)

Assess ONTAP NFS network configuration requirements

Before providing NFS storage to clients, you must verify that networking is correctly configured to meet the NFS provisioning requirements.

Before you begin

The following cluster networking objects must be configured:

- Physical and logical ports
- Broadcast domains
- Subnets (if required)
- IPspaces (as required, in addition to the default IPspace)
- Failover groups (as required, in addition to the default failover group for each broadcast domain)
- External firewalls

Steps

1. Display the available physical and virtual ports:

```
network port show
```

- When possible, you should use the port with the highest speed for the data network.
- All components in the data network must have the same MTU setting for best performance.
- Learn more about `network port show` in the [ONTAP command reference](#).

2. If you are planning to use a subnet name to allocate the IP address and network mask value for a LIF, verify that the subnet exists and has sufficient addresses available:

```
network subnet show
```

Learn more about `network subnet show` in the [ONTAP command reference](#).

Subnets contain a pool of IP addresses that belong to the same layer 3 subnet. Subnets are created by using the `network subnet create` command.

Learn more about `network subnet create` in the [ONTAP command reference](#).

3. Display available IPspaces:

```
network ipspace show
```

You can use the default IPspace or a custom IPspace.

Learn more about `network ipspace show` in the [ONTAP command reference](#).

4. If you want to use IPv6 addresses, verify that IPv6 is enabled on the cluster:

```
network options ipv6 show
```

If required, you can enable IPv6 by using the `network options ipv6 modify` command.

Learn more about `network options ipv6 show` and `network options ipv6 modify` in the [ONTAP command reference](#).

Learn about ONTAP NFS storage capacity provisioning

Before you create a new NFS volume or qtree, you must decide whether to place it in a new or existing SVM, and how much configuration the SVM requires. This decision determines your workflow.

Choices

- If you want to provision a volume or qtree on a new SVM, or on an existing SVM that has NFS enabled but not configured, complete the steps in both "Configuring NFS access to an SVM" and "Adding NFS storage to an NFS-enabled SVM".

[Configure NFS access to an SVM](#)

[Add NFS storage to an NFS-enabled SVM](#)

You might choose to create a new SVM if one of the following is true:

- You are enabling NFS on a cluster for the first time.
- You have existing SVMs in a cluster in which you do not want to enable NFS support.
- You have one or more NFS-enabled SVMs in a cluster, and you want another NFS server in an isolated namespace (multi-tenancy scenario). You should also choose this option to provision storage on an existing SVM that has NFS enabled but not configured. This might be the case if you created the SVM for SAN access or if no protocols were enabled when the SVM was created.

After enabling NFS on the SVM, proceed to provision a volume or qtree.

- If you want to provision a volume or qtree on an existing SVM that is fully configured for NFS access, complete the steps in "Adding NFS storage to an NFS-enabled SVM".

[Adding NFS storage to an NFS-enabled SVM](#)

ONTAP NFS configuration worksheet

The NFS configuration worksheet enables you to collect the required information to set up NFS access for clients.

You should complete one or both sections of the worksheet depending on the decision you made about where to provision storage:

If you are configuring NFS access to an SVM, you should complete both sections.

- Configuring NFS access to an SVM
- Adding storage capacity to an NFS-enabled SVM

If you are adding storage capacity to an NFS-enabled SVM, you should complete only:

- Adding storage capacity to an NFS-enabled SVM

Configure NFS access to an SVM

Parameters for creating an SVM

You supply these values with the `vserver create` command if you are creating a new SVM.

Field	Description	Your value
<code>-vserver</code>	A name you supply for the new SVM that is either a fully qualified domain name (FQDN) or follows another convention that enforces unique SVM names across a cluster.	
<code>-aggregate</code>	The name of an aggregate in the cluster with sufficient space for new NFS storage capacity.	
<code>-rootvolume</code>	A unique name you supply for the SVM root volume.	
<code>-rootvolume-security-style</code>	Use the UNIX security style for the SVM.	<code>unix</code>
<code>-language</code>	Use the default language setting in this workflow.	<code>C.UTF-8</code>
<code>ipspace</code>	IPspaces are distinct IP address spaces in which (storage virtual machines (SVMs)) reside.	

Parameters for creating an NFS server

You supply these values with the `vserver nfs create` command when you create a new NFS server and specify supported NFS versions.

If you are enabling NFSv4 or later, you should use LDAP for improved security.

Field	Description	Your value
<code>-v3, -v4.0, -v4.1, -v4.1-pnfs</code>	<p>Enable NFS versions as needed.</p> <p> v4.2 is also supported in ONTAP 9.8 and later when v4.1 is enabled.</p>	
<code>-v4-id-domain</code>	ID mapping domain name.	
<code>-v4-numeric-ids</code>	Support for numeric owner IDs (enabled or disabled).	

Parameters for creating a LIF

You supply these values with the `network interface create` command when you are creating LIFs. Learn more about `network interface create` in the [ONTAP command reference](#).

If you are using Kerberos, you should enable Kerberos on multiple LIFs.

Field	Description	Your value
<code>-lif</code>	A name you supply for the new LIF.	
<code>-role</code>	Use the data LIF role in this workflow.	<code>data</code>
<code>-data-protocol</code>	Use only the NFS protocol in this workflow.	<code>nfs</code>
<code>-home-node</code>	The node to which the LIF returns when the <code>network interface revert</code> command is run on the LIF. Learn more about <code>network interface revert</code> in the ONTAP command reference .	
<code>-home-port</code>	The port or interface group to which the LIF returns when the <code>network interface revert</code> command is run on the LIF.	
<code>-address</code>	The IPv4 or IPv6 address on the cluster that will be used for data access by the new LIF.	
<code>-netmask</code>	The network mask and gateway for the LIF.	
<code>-subnet</code>	A pool of IP addresses. Used instead of <code>-address</code> and <code>-netmask</code> to assign addresses and netmasks automatically.	
<code>-firewall-policy</code>	Use the default data firewall policy in this workflow.	<code>data</code>

Parameters for DNS host name resolution

You supply these values with the `vserver services name-service dns create` command when you are configuring DNS.

Field	Description	Your value
<code>-domains</code>	Up to five DNS domain names.	
<code>-name-servers</code>	Up to three IP addresses for each DNS name server.	

Name service information

Parameters for creating local users

You supply these values if you are creating local users by using the `vserver services name-service unix-user create` command. If you are configuring local users by loading a file containing UNIX users from a uniform resource identifier (URI), you do not need to specify these values manually.

	User name (<code>-user</code>)	User ID (<code>-id</code>)	Group ID (<code>-primary-gid</code>)	Full name (<code>-full-name</code>)
Example	johnm	123	100	John Miller
1				
2				
3				
...				
n				

Parameters for creating local groups

You supply these values if you are creating local groups by using the `vserver services name-service unix-group create` command. If you are configuring local groups by loading a file containing UNIX groups from a URI, you do not need to specify these values manually.

	Group name (<code>-name</code>)	Group ID (<code>-id</code>)
Example	Engineering	100
1		
2		
3		
...		

n		
---	--	--

Parameters for NIS

You supply these values with the `vserver services name-service nis-domain create` command.



The `-nis-servers` field replaces the `-servers` field. You can use the `-nis-servers` field to specify either a hostname or an IP address for the NIS server.

Field	Description	Your value
<code>-domain</code>	The NIS domain that the SVM will use for name lookups.	
<code>-active</code>	The active NIS domain server.	true or false
<code>-nis-servers</code>	A comma-separated list of IP addresses and hostnames for the NIS servers used by the domain configuration.	

Parameters for LDAP

You supply these values with the `vserver services name-service ldap client create` command.

You will also need a self-signed root CA certificate `.pem` file.

Field	Description	Your value
<code>-vserver</code>	The name of the SVM for which you want to create an LDAP client configuration.	
<code>-client-config</code>	The name you assign for the new LDAP client configuration.	
<code>-ldap-servers</code>	A comma-separated list of IP addresses and hostnames for the LDAP servers.	
<code>-query-timeout</code>	Use the default 3 seconds for this workflow.	3
<code>-min-bind-level</code>	The minimum bind authentication level. The default is anonymous. Must be set to <code>sasl</code> if signing and sealing is configured.	

Field	Description	Your value
-preferred-ad-servers	One or more preferred Active Directory servers by IP address in a comma-delimited list.	
-ad-domain	The Active Directory domain.	
-schema	The schema template to use. You can use a default or custom schema.	
-port	Use the default LDAP server port 389 for this workflow.	389
-bind-dn	The Bind user distinguished name.	
-base-dn	The base distinguished name. The default is "" (root).	
-base-scope	Use the default base search scope subnet for this workflow.	subnet
-session-security	Enables LDAP signing or signing and sealing. The default is none.	
-use-start-tls	Enables LDAP over TLS. The default is false.	

Parameters for Kerberos authentication

You supply these values with the `vserver nfs kerberos realm create` command. Some of the values will differ depending on whether you use Microsoft Active Directory as a Key Distribution Center (KDC) server, or MIT or other UNIX KDC server.

Field	Description	Your value
-vserver	The SVM that will communicate with the KDC.	
-realm	The Kerberos realm.	
-clock-skew	Permitted clock skew between clients and servers.	
-kdc-ip	KDC IP address.	

-kdc-port	KDC port number.	
-adserver-name	Microsoft KDC only: AD server name.	
-adserver-ip	Microsoft KDC only: AD server IP address.	
-adminserver-ip	UNIX KDC only: Admin server IP address.	
-adminserver-port	UNIX KDC only: Admin server port number.	
-passwordserver-ip	UNIX KDC only: Password server IP address.	
-passwordserver-port	UNIX KDC only: Password server port.	
-kdc-vendor	KDC vendor.	{ Microsoft Other }
-comment	Any desired comments.	

You supply these values with the `vserver nfs kerberos interface enable` command.

Field	Description	Your value
-vserver	The name of the SVM for which you want to create a Kerberos configuration.	
-lif	The data LIF on which you will enable Kerberos. You can enable Kerberos on multiple LIFs.	
-spn	The Service Principle Name (SPN)	
-permitted-enc-types	The permitted encryption types for Kerberos over NFS; <code>aes-256</code> is recommended, depending on client capabilities.	
-admin-username	The KDC administrator credentials to retrieve the SPN secret key directly from the KDC. A password is required	

-keytab-uri	The keytab file from the KDC containing the SPN key if you do not have KDC administrator credentials.	
-ou	The organizational unit (OU) under which the Microsoft Active Directory server account will be created when you enable Kerberos using a realm for Microsoft KDC.	

Adding storage capacity to an NFS-enabled SVM

Parameters for creating export policies and rules

You supply these values with the `vserver export-policy create` command.

Field	Description	Your value
-vserver	The name of the SVM that will host the new volume.	
-policyname	A name you supply for a new export policy.	

You supply these values for each rule with the `vserver export-policy rule create` command.

Field	Description	Your value
-clientmatch	Client match specification.	
-ruleindex	Position of export rule in the list of rules.	
-protocol	Use NFS in this workflow.	nfs
-rorule	Authentication method for read-only access.	
-rwrule	Authentication method for read-write access.	
-superuser	Authentication method for superuser access.	
-anon	User ID to which anonymous users are mapped.	

You must create one or more rules for each export policy.

-ruleindex	-clientmatch	-rorule	-rwrule	-superuser	-anon
Examples	0.0.0.0/0,@rootaccess_netgroup	any	krb5	sys	65534
1					
2					
3					
...					
n					

Parameters for creating a volume

You supply these values with the `volume create` command if you are creating a volume instead of a qtree.

Field	Description	Your value
<code>-vserver</code>	The name of a new or existing SVM that will host the new volume.	
<code>-volume</code>	A unique descriptive name you supply for the new volume.	
<code>-aggregate</code>	The name of an aggregate in the cluster with sufficient space for the new NFS volume.	
<code>-size</code>	An integer you supply for the size of the new volume.	
<code>-user</code>	Name or ID of the user that is set as the owner of the volume's root.	
<code>-group</code>	Name or ID of the group that is set as the owner of the volume's root.	
<code>--security-style</code>	Use the UNIX security style for this workflow.	unix
<code>-junction-path</code>	Location under root (/) where the new volume is to be mounted.	

<code>-export-policy</code>	If you are planning to use an existing export policy, you can enter its name when you create the volume.	
-----------------------------	--	--

Parameters for creating a qtree

You supply these values with the `volume qtree create` command if you are creating a qtree instead of a volume.

Field	Description	Your value
<code>-vserver</code>	The name of the SVM on which the volume containing the qtree resides.	
<code>-volume</code>	The name of the volume that will contain the new qtree.	
<code>-qtree</code>	A unique descriptive name you supply for the new qtree, 64 characters or less.	
<code>-qtree-path</code>	The qtree path argument in the format <code>/vol/volume_name/qtree_name\></code> can be specified instead of specifying volume and qtree as separate arguments.	
<code>-unix-permissions</code>	Optional: The UNIX permissions for the qtree.	
<code>-export-policy</code>	If you are planning to use an existing export policy, you can enter its name when you create the qtree.	

Related information

- [ONTAP command reference](#)

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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