



NetApp®

---

## Setting Up a Factory Configured Cluster: Express Guide

# Read

# Me

# First

## Do Not Discard

**This document contains important information the installer needs to complete the setup of this factory configured cluster at the customer's site.**



# Factory configured clusters simplify field installation

Factory configured clusters are assembled, configured, and tested at the NetApp factory. All purchased licenses are added and all nodes are joined to the cluster. You need to finish setting up the hardware, configure the cluster with customer-specific information, and verify that the cluster is set up correctly.

Configuring the cluster in the factory simplifies installation and setup at the customer's site. The checklists in this document provide a high-level list of the hardware and software setup steps completed at the factory and the steps that you still need to complete at the customer's site.

If you have set up clustered Data ONTAP systems before, you can use this document as a checklist reminder of the tasks you need to complete. If you want more detailed instructions, you can use the *Factory Configured Cluster Extended Setup Guide*, which is available from the NetApp Support Site at [https://library.netapp.com/ecm/ecm\\_get\\_file/ECMPI148877](https://library.netapp.com/ecm/ecm_get_file/ECMPI148877) or by scanning the QR code:



## Factory configuration value summary

Factory configured clusters are delivered with user credentials already defined. Many of these default values need to be changed as part of the final setup at the customer's site. Factory configured clusters also have an NFS test configuration on node 1 called "demo".

The following is a summary of the values needed to connect to the cluster. The complete set of values set at the factory is included in this guide and in the *Factory Configured Cluster Extended Setup Guide*.

**User name** admin  
**Password** netapp!123  
**Cluster management port on node 1**

For controller model...	The port is...
32xx and 62xx	e0a
FAS8020	e0e
FAS8040, FAS8060, and FAS8080	e0i

**Cluster management address / gateway** 10.10.10.10 / 255.255.255.0

All factory-defined IP addresses must be changed before connecting the cluster to the customer's network to avoid addressing conflicts.

**Demo (test) aggregate on node 1** demo  
**Demo (test) SVM / volume on node 1** nfs\_server / demo

### Demo (test) data LIF / port on node 1

For controller model...	The port is...
32xx and 62xx	datalif / e0b
FAS8020	datalif / e0f
FAS8040, FAS8060, and FAS8080	datalif / e0j

## Location of components in typical factory configured clusters

Almost any cluster can be ordered as factory configured, so there are no definitive “typical” configurations. However, many factory configured clusters share some common configuration characteristics. Understanding these characteristics helps you complete the installation of the factory configured clusters.

A factory configured cluster can be ordered as a switched or switchless cluster. The types of switches and number of nodes allowed depends on whether it is a switched or switchless cluster.

Cluster type	What is supported
Switched	<ul style="list-style-type: none"><li>• Can be two to eight nodes, with a maximum of four nodes in the main cabinet</li><li>• Must include cluster interconnect switches</li><li>• Optionally can include management switches</li><li>• Optionally can include FC switches (for SAN access)</li></ul>
Switchless	<ul style="list-style-type: none"><li>• Must have two and only two nodes</li></ul> <p><b>Note:</b> Data ONTAP supports switchless clusters starting in Data ONTAP 8.2. Switchless clusters are restricted to two nodes for factory configured clusters.</p> <ul style="list-style-type: none"><li>• Cannot include cluster interconnect or management switches.</li><li>• Optionally can include FC switches (for SAN access)</li></ul>

All factory configured clusters have a main cabinet, which contains nodes and disk shelves. The number of nodes installed in a cabinet depends on the controller model. Any switches are also located in the main cabinet. The following additional cabinets are provided if the number of disk shelves, nodes, or both cannot fit into a single cabinet:

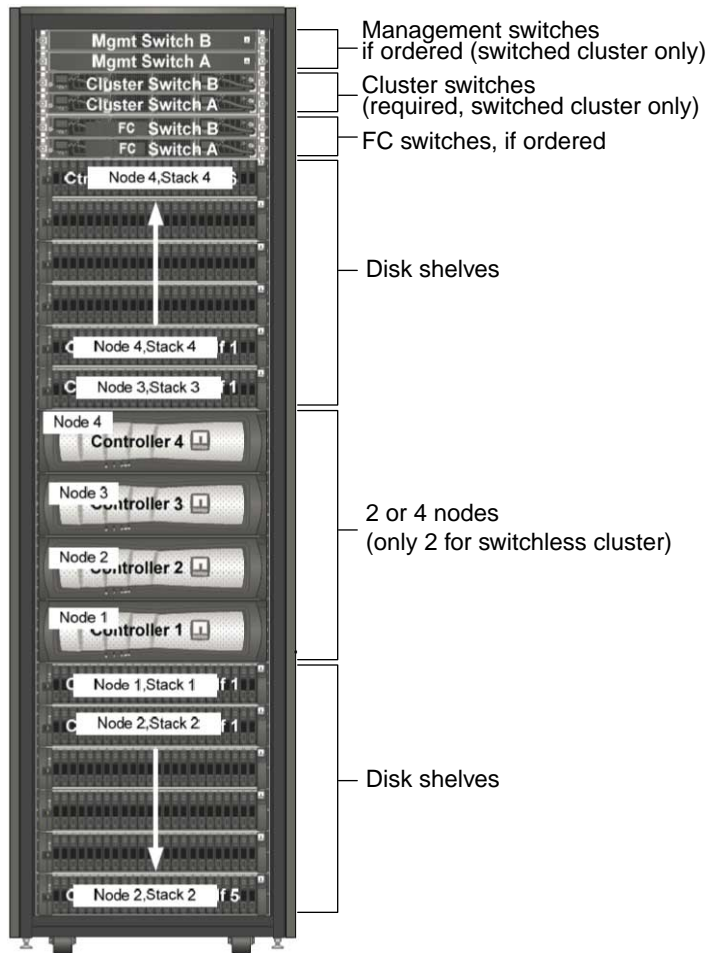
- Shelf-only cabinets
- For switched clusters only, cabinets that contain additional nodes and additional disk shelves.

**Note:** The node and shelf models in a given configuration vary depending on what is ordered. Therefore, the illustrations in this section might not reflect the number of components that can fit in the cabinet for a particular configuration.

### Main cabinet

The following illustration shows an example of the main cabinet for a switched cluster. If the factory configured cluster is switchless, it does not have the management and cluster switches shown in the illustration.

In the following illustration, two of the nodes have only a single disk shelf. The main (and only) cabinet in a single-cabinet configuration could be configured this way with stacks consisting of a single disk shelf. However, the illustration is more likely representative of a multicabinet configuration in which there might be a single disk shelf for a stack in the main cabinet, with the remainder of the stack extending to another cabinet.



Components in the main cabinet are located as follows:

- In a switched cluster, the management switches are at the top of the cabinet, if they are ordered. If management switches are not ordered, the required cluster switches are at the top.
- The FC switches, if ordered, are beneath the cluster switches in a switched configuration and at the top of the cabinet in a switchless configuration.
- For a switched configuration, two or four nodes can be installed in a single cabinet.
- A switchless configuration is limited to two nodes.
- Node 1 is always located so that the bottom of the chassis rests at the 13U mark regardless of the configuration. The first disk shelf in each stack is located in the same cabinet as the associated nodes, as close to them as possible. Additional disk shelves in these stacks are installed in the main cabinet as space permits.

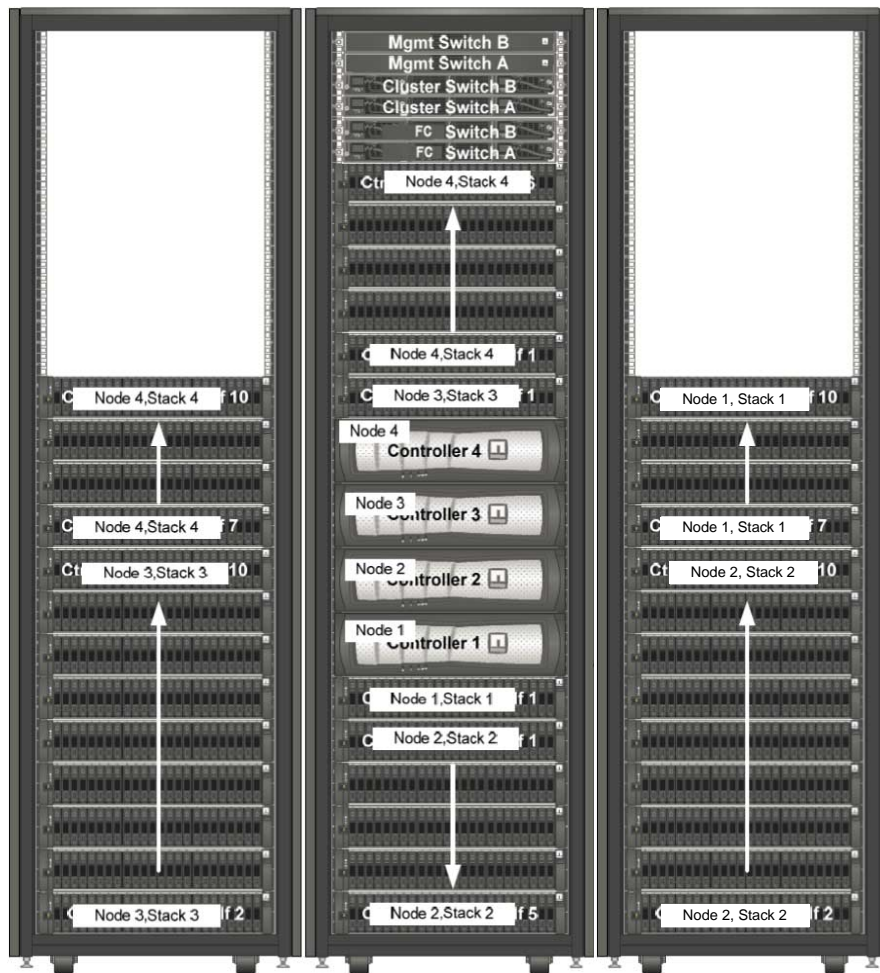
### Single cabinet configuration

If the cluster components fit into a single cabinet, there is not much variation in layout. The nodes and switches are labeled, and the first and last disk shelves are labeled with the node and stack they belong to.

### Multiple cabinet configuration

Configurations that span multiple cabinets have more variation than a single cabinet configuration. The main cabinet is essentially the same as the single cabinet configuration. Additional cabinets contain disk shelves and possibly additional nodes. The node and disk stack labels identify which cabinets and components belong together.

The following illustration shows a multiple cabinet configuration for a switched configuration that includes the optional management switches and FC switches in addition to the required cluster switches:

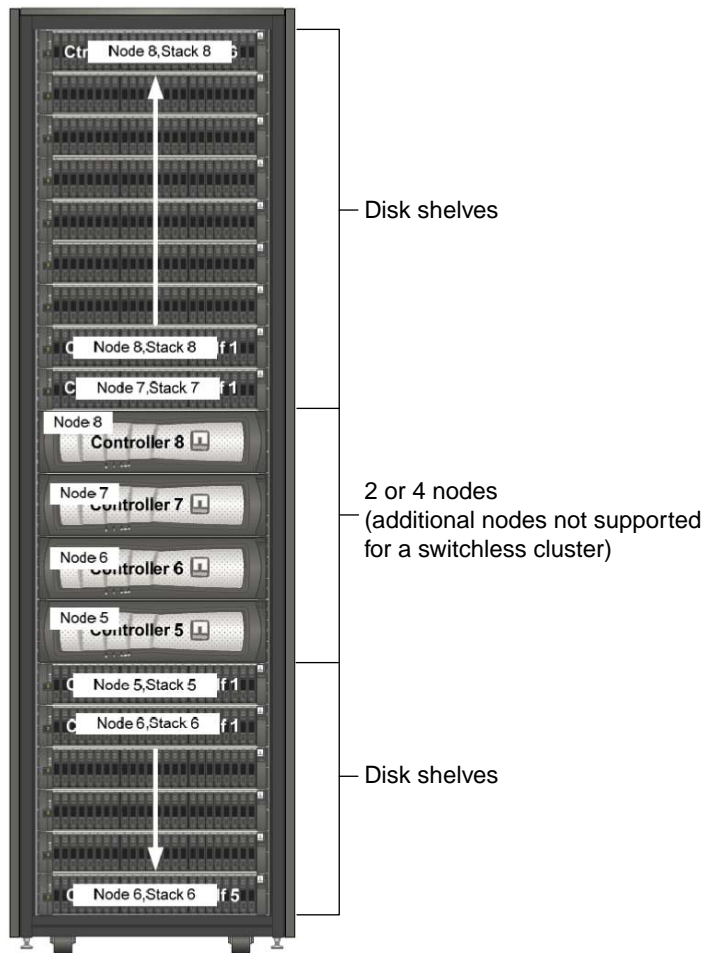


In this example, four nodes are installed in the main cabinet. The additional cabinets contain the disk stacks for those four nodes, but no additional nodes.

The rules for the component location are as follows:

- In a switched cluster, the management switches are at the top of the cabinet, if they are ordered. Otherwise, the required cluster switches are at the top.
- The FC switches, if ordered, are beneath the cluster switches in a switched configuration and at the top of the cabinet in a switchless configuration.
- In the illustration, four nodes are installed in the main cabinets. For a switched configuration, two or four nodes can be installed in a single cabinet. A switchless configuration is limited to two nodes. Node 1 is always located so that the bottom of the chassis rests at the 13U mark regardless of the configuration.
- The first disk shelf in each stack is located in the same cabinet as the associated nodes, as close to them as possible. Additional disk shelves in these stacks are installed as space permits.
- For each controller there is at least one disk shelf in the main cabinet. The first disk shelf of the stack for a node is always located in the same cabinet as the node. However, the main cabinet might not contain all the disk shelves for a node. This is an important concept to understand if your factory configured configuration extends beyond a single cabinet. When cabling, you need to identify the first shelf and the last shelf in the stack; the last shelf in the stack could be in another cabinet that is not directly adjacent to the main cabinet.
- Blank plates are installed to cover any empty spaces in the cabinet.

A multiple cabinet configuration could include a disk-shelf and node cabinet, which would look similar to the following:



As is the case with the nodes in the main cabinet, the first disk shelf of the stack for a node is always located in the same cabinet as the node, as close to it as possible. Disk shelves that are part of the stack for the node might be located in a different cabinet. When connecting the cables between the cabinets, you need to accurately identify the first shelf and the last shelf in the stack when cabling. The factory labels the first and last shelves for you.

## Disk, root aggregate, root volume, and demo configuration layout

The factory configures disk ownership and root aggregate placement for the cluster. The factory also creates a “demo” (test) aggregate, Storage Virtual Machine (SVM, formerly known as Vserver), and volume that can be used to demonstrate that the cluster is set up correctly.

### Disk layout and ownership

The default factory configuration is to distribute shelves evenly between the nodes in an HA pair.

If there are different disk types or sizes, half of each type or size is assigned to each node.

### Root aggregate

The root aggregate for each node is created at the factory using three disks configured as RAID-DP. The specific disks used for the root aggregate are selected by the standard Data ONTAP algorithm and vary depending on the number and types of disks assigned to the node. The algorithm selects three disks of the same type and size for the root aggregate.

SSDs are used for the root aggregate only when SSDs are the only disk type owned by the node.

The root aggregate is named `aggr0` on the first node created in an HA pair, and is named `aggr0node_name` on the other node in each HA pair.

**Note:** For consistency in naming, you might want to rename `aggr0` to `aggr0node_name` when the cluster is being deployed.

## Root volume

A single root volume named `vol0` is created in the root aggregate of each node at the factory.

The root volume can be moved to another aggregate if needed, so that the default disk layout can be changed. Customers must contact technical support for assistance and specific instructions. NetApp and partner personnel should refer to KB article 1013762.

## Demo configuration

The factory creates a demo (test) configuration on node 1 in the cluster that can be used to demonstrate that the cluster is set up correctly and serving data.

The demo test configuration consists of the following items:

- A three-disk aggregate named `demo`
- An SVM named `nfs_server`
- A volume named `/vol/demo` in the aggregate `demo`
- A data LIF named `datalif` on the port of node 1 shown in the following table, with IP address 10.10.9.10, netmask 255.255.255.0

For controller model...	The port is...
32xx and 62xx	<code>datalif / e0b</code>
FAS8020	<code>datalif / e0f</code>
FAS8040, FAS8060, and FAS8080	<code>datalif / e0j</code>

The demo configuration can be expanded, renamed, or both if desired. Or it can be removed if desired. If the demo aggregate is removed, the disks are converted to spares and must be zeroed before they can be added to another aggregate.

## Hardware setup checklist

You need to know which hardware setup steps were completed at the factory and which you need to complete at the customer's site.

Step	Completed at factory	Completed by you
Mount components in one or more cabinets.	Yes	No
Position cabinets in the desired location.	No	Yes Position them in the original order to ensure that the supplied cables are long enough.
Connect multiple cabinets to each other, if applicable.	No	Yes Use the cabinet interconnect kit if it is included in the order. The kit box is labeled.
Secure the cabinets to the floor, if applicable.	No	Yes Use the universal bolt-down kit if it is included in the order. The kit box is labeled.

Step	Completed at factory	Completed by you
Cable the components within the cabinet.	Yes Cables 5 meters and longer are removed for shipping and placed in the accessories box.	No
Connect the cables between cabinets, if applicable.	No	Yes Cables are in the accessories box.
Connect management cables to the customer's network.	No	Yes Connect them directly or through the CN1601 management switches, if present.  <b>Attention:</b> To avoid address conflicts, do not connect management ports to the customer's network until after you change the default IP addresses to the customer's values.
Connect console ports to the customer's terminal server, if applicable.	No	Yes
Connect the customer's data cables to the cluster.	No	Yes
Connect the cabinets to power and power on the components.	No	Yes Power them on in the following order:  1. PDUs 2. Disk shelves 3. Nodes
Assign IP addresses to the management ports of the cluster switches and to the management ports of the management switches if present.	No	Yes for switched clusters only Connect to the serial console port of each switch and log in with user name "admin" with no password. Suggested addresses are 10.10.10.81, 10.10.10.82, 10.10.10.83, and 10.10.10.84.
Verify cabling by running the Config Advisor tool	No	Yes Config Advisor 3.3 or later is required for switchless clusters

## Software setup checklist

You need to know which software setup steps were completed at the factory and which you need to complete at the customer's site.

Step	Completed at factory	Completed by you
Install the clustered Data ONTAP software	Yes	No



Step	Completed at factory	Completed by you
Create the cluster on the first node. This includes the following: <ul style="list-style-type: none"> <li>• Name the cluster.</li> <li>• Set the admin password.</li> <li>• Set up the private cluster interconnect.</li> <li>• Install all purchased license keys.</li> <li>• Create the cluster management interface.</li> <li>• Create the node management interface.</li> </ul>	Yes	No
Join the remaining nodes to the cluster.	Yes	No
Enable storage failover on one node of each HA pair (for clusters with more than two nodes) or configure cluster high availability (two-node clusters only).	Yes	No
Enable the switchless-cluster option on a two-node switchless cluster.	Yes	No
Create a test SVM on node cluster-01.	Yes The SVM is configured for NFS and the volume demo is exported.	No
Configure user credentials and management IP addresses on the management and cluster switches.	Yes, if ordered. User IDs are “admin” with no password.	No
Thoroughly test the cluster.	Yes	No
Complete the cluster setup worksheet.	No	Yes The worksheet is included in both the <i>Setting Up a Factory Configured Cluster: Express Guide</i> and the <i>Factory Configured Cluster Extended Setup Guide</i> .

Step	Completed at factory	Completed by you
Connect to the cluster management port of node cluster1-01 using SSH and log in.	No	<p>Yes</p> <ol style="list-style-type: none"> <li>1. Set your laptop to an unused address in the 10.10.10.x subnet with netmask 255.255.255.0.</li> <li>2. If the optional management switches are installed, connect your laptop to an open port on one of the management switches.</li> <li>3. If there are no management switches, connect directly to the following port of node 1 for your controller: e0a for 32xx and 62xx controllers; e0e for FAS8020 controllers; e0i for FAS8040, FAS8060, and FAS8080 controllers.</li> <li>4. Log in to the cluster as “admin” with a “netapp!123” password. The cluster management IP address is 10.10.10.10 with netmask 255.255.255.0.</li> </ol>
Change the password for the admin account to the customer's value.	No	Yes
Configure each node with the customer's values.	No	<p>Yes</p> <p>Run the Cluster Setup wizard on each node and enter the customer's values from the worksheet.</p>
Discover the cluster in OnCommand System Manager.	No	Yes
Configure an NTP server for the cluster.	No	Yes
Verify the health of the cluster and that the cluster is in quorum.	No	Yes
Verify basic operation of the cluster.	No	<p>Recommended, if the customer allows you to access the network.</p> <ol style="list-style-type: none"> <li>1. Connect an NFS client to the test SVM and verify that you can write and read data.</li> <li>2. If an NFS license was not purchased, install a temporary NFS license for testing.</li> </ol>
Test storage failover.	No	<p>Yes</p> <p>Use the Data ONTAP CLI <code>storage failover takeover</code> and <code>storage failover giveback</code> commands. Be sure to revert all LIFs back to their home nodes after testing.</p>

Step	Completed at factory	Completed by you
Optional: Change the cluster name if desired, for example, to better distinguish the clusters.	No	Yes Use the Data ONTAP CLI <code>cluster identity modify</code> command to
Optional: Change the node name, if desired.	No	Yes Use the Data ONTAP CLI <code>system node rename</code> command to rename the node.
Optional: Rename and expand the demo (test) aggregate, if you decide to keep it.	No	Yes Use the Data ONTAP CLI <code>storage aggregate rename</code> command to rename the demo test aggregate.
Configure AutoSupport.	No	Yes

## Completing the cluster setup worksheet

The worksheet enables you to record the customer values that you need to complete the cluster setup process.

As you record values for the IP addresses, keep in mind that each IP address must be a valid host address that meets the following criteria:

- The host address, in combination with the associated netmask, must not match the polling address or the broadcast address.
- The host address must not be already assigned to another system on the network.
- The host address must be a valid IP address (for example, it cannot be 0.0.0.0 or 256.345.567.257).
- The host address must not be a special use IPv4 address as defined by [RFC 3330](#).

### Switches

Types of information	Factory default	Customer values
Cluster switch user name	admin	
Cluster switch password	No password set	
Cluster switch A management port IP address	No factory default Suggested value for Config Advisor testing: 10.10.10.81	
Cluster switch B management port IP address	No factory default Suggested value for Config Advisor testing: 10.10.10.82	
Cluster switch netmask	No factory default Suggested value for Config Advisor testing: 255.255.255.0	
Management switch user name	admin	
Management switch password	No password set	
Management switch A management port IP address	No factory default Suggested value for Config Advisor testing: 10.10.10.83	

Types of information	Factory default	Customer values
Management switch B management port IP address	No factory default Suggested value for Config Advisor testing: 10.10.10.84	
Management switch netmask	No factory default Suggested value for Config Advisor testing: 255.255.255.0	

### Admin SVM

Types of information	Factory default	Customer values
Cluster management interface port	<ul style="list-style-type: none"> <li>• e0a for 32xx and 62xx controllers</li> <li>• e0e for FAS8020 controllers</li> <li>• e0i for FAS8040, FAS8060, and FAS8080 controllers</li> </ul>	
Cluster management interface IP address	10.10.10.10	
Cluster management interface netmask	255.255.255.0	
Cluster management interface default gateway	10.10.10.1	
DNS domain name	No factory default	
Name server IP addresses	No factory default	

### Node information (for each node in the cluster)

Types of information	Factory default	Customer values
Physical location of the controller	No factory default	Node 1:
		Node 2:
		Node 3:
		Node 4:
		Node 5:
		Node 6:
		Node 7:
		Node 8:

Types of information	Factory default	Customer values
Node management interface port	e0M	Node 1:
		Node 2:
		Node 3:
		Node 4:
		Node 5:
		Node 6:
		Node 7:
		Node 8:
Node management interface IP address	Node 1: 10.10.10.101	Node 1:
	Node 2: 10.10.10.102	Node 2:
	Node 3: 10.10.10.103	Node 3:
	Node 4: 10.10.10.104	Node 4:
	Node 5: 10.10.10.105	Node 5:
	Node 6: 10.10.10.106	Node 6:
	Node 7: 10.10.10.107	Node 7:
	Node 8: 10.10.10.108	Node 8:
Node management interface netmask	255.255.255.0	
Node management interface default gateway	10.10.10.1	

#### Related information

[RFC 3330: tools.ietf.org/html/rfc3330](http://tools.ietf.org/html/rfc3330)

## Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Customer Fitness, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.