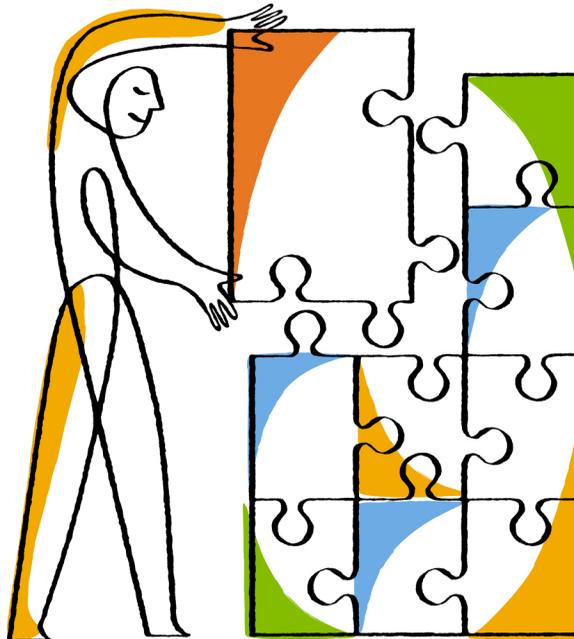




**NetApp®**

## Clustered Data ONTAP® 8.3

### Software Setup Guide



NetApp, Inc.  
495 East Java Drive  
Sunnyvale, CA 94089  
U.S.

Telephone: +1 (408) 822-6000  
Fax: +1 (408) 822-4501  
Support telephone: +1 (888) 463-8277  
Web: [www.netapp.com](http://www.netapp.com)  
Feedback: [doccomments@netapp.com](mailto:doccomments@netapp.com)

Part number: 215-09413\_B0  
January 2015

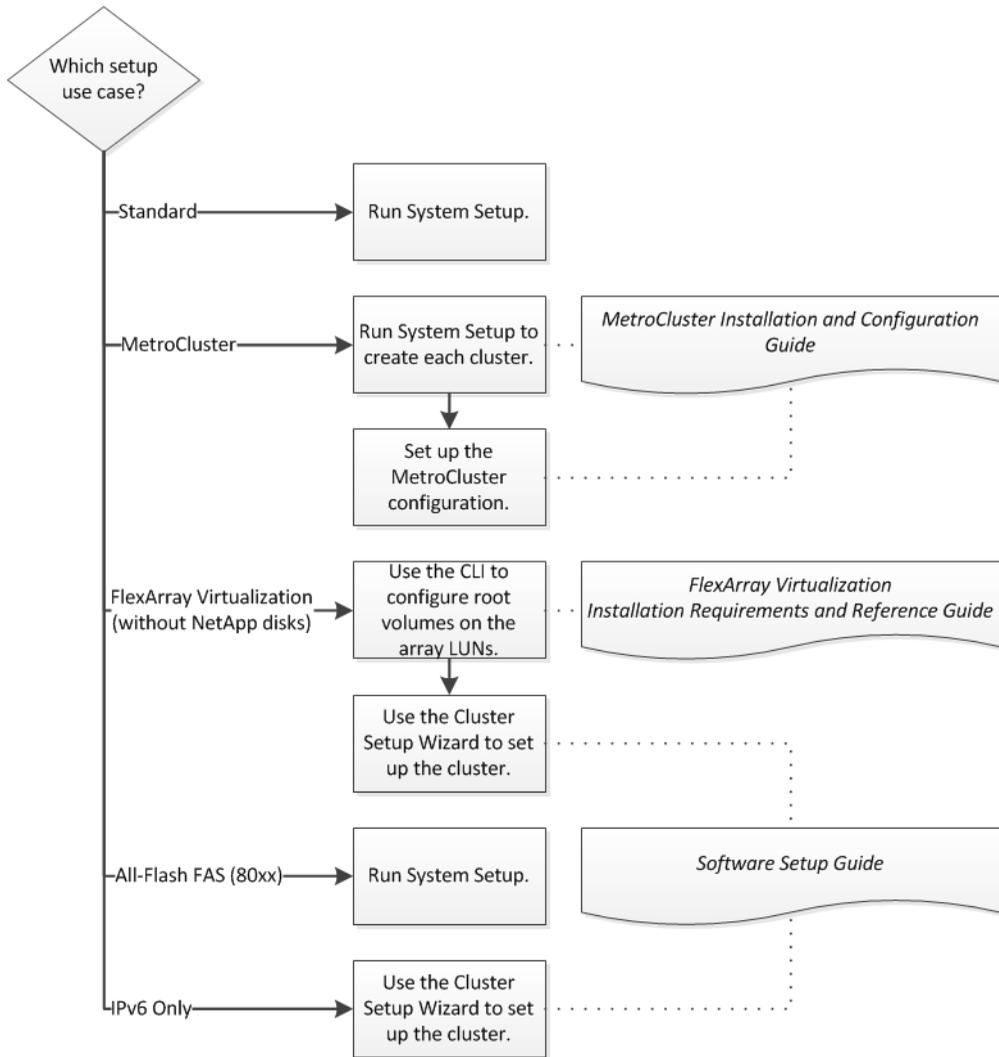


# Contents

<b>Cluster setup workflows .....</b>	<b>4</b>
<b>Setting up the cluster .....</b>	<b>6</b>
Completing the cluster setup worksheet .....	6
Creating the cluster on the first node .....	12
Joining a node to the cluster .....	14
Synchronizing the system time across the cluster .....	16
Deciding where to send high-severity event notifications .....	18
Configuring high-severity EMS events to send email notifications .....	19
Configuring high-severity EMS events to forward notifications to a syslog server .....	20
Configuring SNMP traphosts to receive event notifications .....	21
<b>Additional system configuration tasks to complete .....</b>	<b>22</b>
<b>Copyright information .....</b>	<b>23</b>
<b>Trademark information .....</b>	<b>24</b>
<b>How to send comments about documentation and receive update     notification .....</b>	<b>25</b>
<b>Index .....</b>	<b>26</b>

# Cluster setup workflows

After installing the hardware, you should use System Setup whenever possible to set up the cluster. However, if System Setup does not support your configuration, you can use the CLI.



## Related information

[NetApp Documentation: System Setup \(current releases\)](#)

*Clustered Data ONTAP 8.3 MetroCluster Installation and Configuration Guide*  
*FlexArray Virtualization Installation Requirements and Reference Guide*

## Setting up the cluster

---

Setting up the cluster involves setting up each node, creating the cluster on the first node, and joining any remaining nodes to the cluster.

### Steps

1. [Completing the cluster setup worksheet](#) on page 6  
The cluster setup worksheet enables you to record the values that you need during the cluster setup process. If a default value is provided, you can use that value or else enter your own.
2. [Creating the cluster on the first node](#) on page 12  
You use the Cluster Setup wizard to create the cluster on the first node. The wizard helps you to configure the cluster network that connects the nodes (if the cluster consists of two or more nodes), create the cluster admin Storage Virtual Machine (SVM), add feature license keys, and create the node management interface for the first node.
3. [Joining a node to the cluster](#) on page 14  
After creating a new cluster, you use the Cluster Setup wizard to join each remaining node to the cluster one at a time. The wizard helps you to configure each node's node management interface.
4. [Synchronizing the system time across the cluster](#) on page 16  
Synchronizing the time ensures that every node in the cluster has the same time, and prevents CIFS and Kerberos failures.
5. [Deciding where to send high-severity event notifications](#) on page 18  
Before you configure high-severity EMS event notifications, you need to decide whether to send the notifications to a email address, a syslog server, or an SNMP traphost.

## Completing the cluster setup worksheet

The cluster setup worksheet enables you to record the values that you need during the cluster setup process. If a default value is provided, you can use that value or else enter your own.

### System defaults (for clusters configured to use network switches)

The system defaults are the default values for the private cluster network. It is best to use these default values. However, if they do not meet your requirements, you can use the table to record your own values.

You only need to consider the system defaults for clusters that are connected using network switches. Single-node clusters and two-node switchless clusters do not use a cluster network.

Types of information	Default	Your values
Private cluster network ports	<i>Clustered Data ONTAP 8.3 Network Management Guide</i>	
MTU size for cluster ports Every node in the cluster must have the same MTU size as the cluster interconnect switches.	9000 bytes	
Cluster network netmask	255.255.0.0	
Cluster interface IP addresses (for each cluster network port on each node) The IP addresses for each node must be on the same subnet.	169.254.x.x	

### Cluster information

Types of information	Your values
Cluster name The name must begin with a letter, and it must be fewer than 44 characters. The name can include the following special characters: · - _	
Cluster base license key The cluster base license key is not required to complete setup. However, it must be installed before you can install any feature licenses. To find your cluster base license key: <b>1.</b> Go to the the NetApp Support Site. <i>NetApp Support</i> <b>2.</b> Click <b>My Support &gt; Software Licenses</b> .	

**Feature license keys**

You can find license keys for your initial or add-on software orders at the NetApp Support Site under **My Support > Software Licenses**.

Types of information	Your values
Feature license keys	

**Admin Storage Virtual Machine (SVM)**

Types of information	Your values
<p>Cluster administrator password</p> <p>The password for the admin account that the cluster requires before granting cluster administrator access to the console or through a secure protocol.</p> <p>The default rules for passwords are as follows:</p> <ul style="list-style-type: none"> <li>• A password must be at least eight characters long.</li> <li>• A password must contain at least one letter and one number.</li> </ul>	

<b>Types of information</b>	<b>Your values</b>
<p>Cluster management interface port</p> <p>The physical port that is connected to the data network and enables the cluster administrator to manage the cluster.</p>	
<p>Cluster management interface IP address</p> <p>A unique IPv4 or IPv6 address for the cluster management interface. The cluster administrator uses this address to access the admin SVM and manage the cluster. Typically, this address should be on the data network.</p> <p>You can obtain this IP address from the administrator responsible for assigning IP addresses in your organization.</p> <p>Example: 192.0.2.66</p>	
<p>Cluster management interface netmask (IPv4)</p> <p>The subnet mask that defines the range of valid IPv4 addresses on the cluster management network.</p> <p>Example: 255.255.255.0</p>	
<p>Cluster management interface netmask length (IPv6)</p> <p>If the cluster management interface uses an IPv6 address, then this value represents the prefix length that defines the range of valid IPv6 addresses on the cluster management network.</p> <p>Example: 64</p>	
<p>Cluster management interface default gateway</p> <p>The IP address for the router on the cluster management network.</p>	
<p>DNS domain name</p> <p>The name of your network's DNS domain.</p> <p>The domain name must consist of alphanumeric characters. To enter multiple DNS domain names, separate each name with either a comma or a space.</p>	

Types of information	Your values
<p>Name server IP addresses</p> <p>The IP addresses of the DNS name servers. Separate each address with either a comma or a space.</p>	

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

Types of information	Your values
<p>Physical location of the controller</p> <p>A description of the physical location of the controller. Use a description that identifies where to find this node in the cluster (for example, “Lab 5, Row 7, Rack B”).</p>	
<p>Node management interface port</p> <p>The physical port that is connected to the node management network and enables the cluster administrator to manage the node.</p>	
<p>Node management interface IP address</p> <p>A unique IPv4 or IPv6 address for the node management interface on the management network. If you defined the node management interface port to be a data port, then this IP address should be a unique IP address on the data network.</p> <p>You can obtain this IP address from the administrator responsible for assigning IP addresses in your organization.</p> <p>Example: 192.0.2.66</p>	
<p>Node management interface netmask (IPv4)</p> <p>The subnet mask that defines the range of valid IP addresses on the node management network.</p> <p>If you defined the node management interface port to be a data port, then the netmask should be the subnet mask for the data network.</p> <p>Example: 255.255.255.0</p>	

Types of information	Your values
<p>Node management interface netmask length (IPv6)</p> <p>If the node management interface uses an IPv6 address, then this value represents the prefix length that defines the range of valid IPv6 addresses on the node management network.</p> <p>Example: 64</p>	
<p>Node management interface default gateway</p> <p>The IP address for the router on the node management network.</p>	
<p>System configuration backup destination address (single-node clusters only)</p> <p>The remote URL where the cluster configuration backups will be uploaded. You can specify either an HTTP or FTP address.</p> <p><b>Note:</b> The web server that serves the remote URL must have PUT operations enabled.</p>	
<p>User name for the configuration backup destination address (single-node clusters only)</p> <p>The user name required to log in to the remote URL and upload the configuration backup file.</p>	
<p>Password for the configuration backup destination address (single-node clusters only)</p> <p>The password for the remote URL, if the user name requires a password.</p>	

### NTP server information

Types of information	Your values
<p>NTP server addresses</p> <p>The IP addresses of the Network Time Protocol (NTP) servers at your site. These servers are used to synchronize the time across the cluster.</p>	

## Creating the cluster on the first node

You use the Cluster Setup wizard to create the cluster on the first node. The wizard helps you to configure the cluster network that connects the nodes (if the cluster consists of two or more nodes), create the cluster admin Storage Virtual Machine (SVM), add feature license keys, and create the node management interface for the first node.

### Before you begin

The cluster setup worksheet should be completed, the storage system hardware should be installed and cabled, and the console should be connected to the node on which you intend to create the cluster.

### Steps

1. Power on the first node.

The node boots, and then the Node Setup wizard starts on the console, informing you that AutoSupport will be enabled automatically.

```
Welcome to node setup.

You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the setup wizard.
    Any changes you made before quitting will be saved.

To accept a default or omit a question, do not enter a value.

This system will send event messages and weekly reports to NetApp
Technical
Support. To disable this feature, enter "autosupport modify -support
disable"
within 24 hours. Enabling AutoSupport can significantly speed problem
determination and resolution should a problem occur on your system.
For
further information on AutoSupport, see:
http://support.netapp.com/autosupport/

Type yes to confirm and continue {yes}:
```

2. Exit the **Node Setup** wizard:

**exit**

The Node Setup wizard enables you to configure the node's node management interface for use with System Setup. You do not need to use the Node Setup wizard if you are setting up the cluster in the CLI.

The Node Setup wizard exits, and a login prompt appears, warning that you have not completed the setup tasks:

```
Exiting the node setup wizard. Any changes you made have been saved.

Warning: You have exited the node setup wizard before completing all
of the tasks. The node is not configured. You can complete node setup
by typing
"node setup" in the command line interface.

login:
```

3. Log in to the admin account by using the admin user name.

4. Start the **Cluster Setup** wizard:

```
cluster setup
```

```
::> cluster setup

Welcome to the cluster setup wizard.

You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the cluster setup wizard.
  Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.

Do you want to create a new cluster or join an existing cluster?
{create, join}:
```

5. Create a new cluster:

```
create
```

6. Accept the system defaults or enter your own values:

- To accept the system defaults, press Enter.
- To enter your own values, type **no**, and then press Enter.

7. Follow the prompts to complete the **Cluster Setup** wizard:

- To accept the default value for a prompt, press Enter.

The default values are determined automatically based on your platform and network configuration.

- To enter your own value for the prompt, enter the value, and then press Enter.

8. After the **Cluster Setup** wizard is completed and exits, verify that the cluster is active and the first node is healthy:

```
cluster show
```

### Example

The following example shows a cluster in which the first node (cluster1-01) is healthy and eligible to participate:

```
cluster1::> cluster show
Node                Health  Eligibility
-----
cluster1-01        true   true
```

You can access the Cluster Setup wizard to change any of the values you entered for the admin SVM or node SVM by using the `cluster setup` command.

### After you finish

If the cluster consists of two or more nodes, you should join each remaining node to the cluster.

## Joining a node to the cluster

After creating a new cluster, you use the Cluster Setup wizard to join each remaining node to the cluster one at a time. The wizard helps you to configure each node's node management interface.

### Before you begin

The cluster must be created on the first node.

### About this task

You can only join one node to the cluster at a time. When you start to join a node to the cluster, you must complete the join, and the node must be part of the cluster before you can start to join the next node.

### Steps

1. Power on the node.

The node boots, and then the Node Setup wizard starts on the console, informing you that AutoSupport will be enabled automatically.

```
Welcome to node setup.
```

```
You can enter the following commands at any time:
```

```
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the setup wizard.
Any changes you made before quitting will be saved.
```

```
To accept a default or omit a question, do not enter a value.
```

```
This system will send event messages and weekly reports to NetApp
Technical
Support. To disable this feature, enter "autosupport modify -support
disable"
within 24 hours. Enabling AutoSupport can significantly speed problem
determination and resolution should a problem occur on your system.
For
further information on AutoSupport, see:
http://support.netapp.com/autosupport/
```

```
Type yes to confirm and continue {yes}:
```

## 2. Exit the **Node Setup** wizard:

```
exit
```

The Node Setup wizard exits, and a login prompt appears, warning that you have not completed the setup tasks:

```
Exiting the node setup wizard. Any changes you made have been saved.
```

```
Warning: You have exited the node setup wizard before completing all
of the tasks. The node is not configured. You can complete node setup
by typing
"node setup" in the command line interface.
```

```
login:
```

## 3. Log in to the admin account by using the admin user name.

## 4. Start the **Cluster Setup** wizard:

```
cluster setup
```

```
::> cluster setup
```

```
Welcome to the cluster setup wizard.
```

```
You can enter the following commands at any time:
```

```
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
```

```
"exit" or "quit" - if you want to quit the cluster setup wizard.
  Any changes you made before quitting will be saved.
```

```
You can return to cluster setup at any time by typing "cluster setup".
  To accept a default or omit a question, do not enter a value.
```

```
Do you want to create a new cluster or join an existing cluster?
{create, join}:
```

5. Join the node to the cluster:

```
join
```

6. Follow the prompts to set up the node and join it to the cluster:

- To accept the default value for a prompt, press Enter.
- To enter your own value for the prompt, enter the value, and then press Enter.

7. After the **Cluster Setup** wizard is completed and exits, verify that the node is healthy and eligible to participate in the cluster:

```
cluster show
```

### Example

The following example shows a cluster after the second node (cluster1-02) has been joined to the cluster:

```
cluster1::> cluster show
Node           Health  Eligibility
-----
cluster1-01    true    true
cluster1-02    true    true
```

You can access the Cluster Setup wizard to change any of the values you entered for the admin SVM or node SVM by using the `cluster setup` command.

8. Repeat this task for each remaining node.

## Synchronizing the system time across the cluster

Synchronizing the time ensures that every node in the cluster has the same time, and prevents CIFS and Kerberos failures.

### Before you begin

A Network Time Protocol (NTP) server should be set up at your site.

**About this task**

You synchronize the time across the cluster by associating the cluster with one or more NTP servers.

**Steps**

1. Verify that the system time and time zone is set correctly for each node.

All nodes in the cluster should be set to the same time zone.

- a. Use the `cluster date show` command to display the current date, time, and time zone for each node.

**Example**

```
cluster1::> cluster date show
Node          Date          Time zone
-----
cluster1-01  01/06/2015  09:35:15  America/New_York
cluster1-02  01/06/2015  09:35:15  America/New_York
cluster1-03  01/06/2015  09:35:15  America/New_York
cluster1-04  01/06/2015  09:35:15  America/New_York
4 entries were displayed.
```

- b. Optional: Use the `cluster date modify` command to change the date or time zone for all of the nodes.

**Example**

This example changes the time zone for the cluster to be GMT:

```
cluster1::> cluster date modify -timezone GMT
```

2. Use the `cluster time-service ntp server create` command to associate the cluster with your NTP server.

**Example**

This example assumes that DNS has been configured for the cluster. If you have not configured DNS, you must specify the IP address of the NTP server:

```
cluster1::> cluster time-service ntp server create -server
ntp1.example.com
```

3. Verify that the cluster is associated with an NTP server:

```
cluster time-service ntp server show
```

**Example**

```
cluster1::> cluster time-service ntp server show
Server                               Version
-----
ntp1.example.com                     auto
```

**Related information**

*[Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators](#)*

## Deciding where to send high-severity event notifications

Before you configure high-severity EMS event notifications, you need to decide whether to send the notifications to a email address, a syslog server, or an SNMP traphost.

**About this task**

If your environment already contains a syslog server for aggregating the logged events from other systems, such as servers and applications, then it is easier to use that syslog server also for high-severity event notifications from storage systems.

If your environment does not already contain a syslog server, then it is easier to use email for high-severity event notifications.

If you already forward event notifications to an SNMP traphost, then you might want to monitor that traphost for high-severity events.

**Note:** SNMP supports only a small subset of the high-severity events. The Technical Report *[NetApp Technical Report 4220: SNMP Support in Data ONTAP 8.2.x and Data ONTAP 8.3.x](#)* contains lists of all default events that are supported by SNMP.

**Choices**

- If you want the EMS to send high-severity event notifications to an email address, see *[Configuring high-severity EMS events to send email notifications](#)* on page 19.
- If you want the EMS to forward high-severity event notifications to a syslog server, see *[Configuring high-severity EMS events to forward notifications to a syslog server](#)* on page 20.
- If you want the EMS to forward event notifications to an SNMP traphost, see *[Configuring SNMP traphosts to receive event notifications](#)* on page 21

## Configuring high-severity EMS events to send email notifications

To receive email notifications of the most severe events, you must configure the EMS to send email messages for the top three severity levels (Critical, Alert, Emergency) and a few additional events that signal high-severity activity.

### Before you begin

DNS must be configured on the cluster to resolve the email addresses.

### About this task

Before you configure the events, you must first create a dedicated email destination for the event notifications. Event email notifications are sent to configured email addresses using SMTP.

This task configures all events of severity level Critical, Alert, and Emergency. It also configures a selected group of additional events that also report high-severity activity. You must configure all these events to make sure you are notified of system issues that require prompt attention.

You can perform this task any time the cluster is running by entering the commands on the Data ONTAP command line.

### Steps

1. Configure the event SMTP mail server settings:

```
event config modify -mailserver mailhost@your_domain
-mailfrom cluster_name@your_domain
```

2. Create an email destination for high-severity event notifications:

```
event destination create -name important_events
-mail your_email@your_domain
```

3. Configure all high-severity events to send email notifications.

- a. Configure all events of severity level Critical, Alert, and Emergency to send email notifications to the email address `important_events` that you just created:

```
event route add-destinations {-severity CRITICAL|ALERT|EMERGENCY}
-destinations important_events
```

- b. Configure additional high-severity events to send email notifications to `important_events`:

```
event route add-destinations {csm.sessionFailed| secd.dns*| secd.nis*|
secd.ldap*| callhome.aggr.restricted| callhome.c.fan*|
callhome.carrier.fault| callhome.ch.ps.*| callhome.chassis.*|
callhome.client.app.emerg| callhome.client.app.crit|
callhome.client.app.alert| callhome.cpu*| callhome.hm.alert.*|
callhome.netif.fatal.error| callhome.reboot*|
callhome.sblade.import.susp| callhome.sblade.unavailable|
```

```
callhome.sfo.giveback| callhome.sfo.takeover*| callhome.shlf.overtemp|
callhome.shlf.power.intr| callhome.spm.process.maxexit|
callhome.clus.vol.cre.fail | clam.takeover |
clam.heartbeat.state.change} -destinations important_events
```

## Configuring high-severity EMS events to forward notifications to a syslog server

To log notifications of the most severe events on a syslog server, you must configure the EMS to forward notifications for the top three severity levels (Critical, Alert, Emergency) and a few additional events that signal high-severity activity.

### Before you begin

DNS must be configured on the cluster to resolve the syslog server name.

### About this task

If your environment does not already contain a syslog server for event notifications, you must first create one. If your environment already contains a syslog server for logging events from other systems, then you might want to use that one for high-severity event notifications.

This task configures all events of severity level Critical, Alert, and Emergency. It also configures a selected group of additional events that also report high-severity activity. You must configure all these events to make sure you are notified of system issues that require prompt attention.

You can perform this task any time the cluster is running by entering the commands on the Data ONTAP command line.

### Steps

1. Create a syslog server destination for high-severity events:

```
event destination create -name syslog_ems -syslog ip_address
-syslog-facility default
```

2. Configure all high-severity events to forward notifications to the syslog server.

- a. Configure all events of severity level Critical, Alert, and Emergency to forward notifications to the syslog server that you just created or to your existing syslog server:

```
event route add-destinations {-severity CRITICAL|ALERT|EMERGENCY}
-destinations syslog_ems
```

- b. Configure additional high-severity events to forward notifications to the syslog server:

```
event route add-destinations {csm.sessionFailed| secd.dns*| secd.nis*|
secd.ldap*| callhome.aggr.restricted| callhome.c.fan*|
callhome.carrier.fault| callhome.ch.ps.*| callhome.chassis.*|
callhome.client.app.emerg| callhome.client.app.crit|
callhome.client.app.alert| callhome.cpu*| callhome.hm.alert.*|
```

```
callhome.netif.fatal.error| callhome.reboot*|
callhome.sblade.import.susp| callhome.sblade.unavailable|
callhome.sfo.giveback| callhome.sfo.takeover*| callhome.shlf.overtemp|
callhome.shlf.power.intr| callhome.spm.process.maxexit|
callhome.clus.vol.cre.fail | clam.takeover |
clam.heartbeat.state.change} -destinations syslog_ems
```

## Configuring SNMP traphosts to receive event notifications

To receive event notifications on an SNMP traphost, you must configure a traphost. SNMP supports only a small subset of the top three severity level events (Critical, Alert, Emergency) and none of the additional events that signal high-severity activity.

### Before you begin

- SNMP and SNMP traps must be enabled on the cluster.
  - Note:** SNMP and SNMP traps are enabled by default.
- DNS must be configured on the cluster to resolve the traphost names.

### About this task

If you do not already have an SNMP traphost configured to receive event notifications (SNMP traps), you must add one.

The Technical Report [NetApp Technical Report 4220: SNMP Support in Data ONTAP 8.2.x and Data ONTAP 8.3.x](#) contains lists of all default events that are supported by SNMP.

You can perform this task any time the cluster is running by entering the commands on the Data ONTAP command line.

### Step

1. If your environment does not already have an SNMP traphost configured to receive event notifications, add one:

```
system snmp traphost add -peer-address snmp_traphost_name
```

All event notifications that are supported by SNMP by default are forwarded to the SNMP traphost, which includes both high-severity and low-severity event notifications.

## Additional system configuration tasks to complete

---

After setting up the cluster, you can use either OnCommand System Manager or the Data ONTAP CLI to continue configuring the cluster.

System configuration task	Resource
Configure networking: <ul style="list-style-type: none"><li>• Create broadcast domains</li><li>• Create subnets</li><li>• Create IP spaces</li></ul>	<i>Clustered Data ONTAP 8.3 Network Management Guide</i>
Set up the Service Processor	<i>Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators</i>
Lay out your aggregates	<i>Clustered Data ONTAP 8.3 Physical Storage Management Guide</i>
Create data Storage Virtual Machines (SVMs)	<i>Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators</i>

## Copyright information

---

Copyright © 1994–2015 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.

**RESTRICTED RIGHTS LEGEND:** Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## Trademark information

---

NetApp, the NetApp logo, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

## How to send comments about documentation and receive update notification

---

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email to [doctrcomments@netapp.com](mailto:doctrcomments@netapp.com). To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

# Index

- A**
- alert events
    - configuring EMS to send email notifications [19](#)
    - configuring to forward notifications to syslog server [20](#)
- C**
- cluster setup
    - information to gather for [6](#)
    - introduction to the process for [6](#)
    - workflows for [4](#)
  - Cluster Setup wizard
    - using to create a cluster on the first node [12](#)
    - using to join nodes [14](#)
  - clusters
    - creating on the first node [12](#)
    - joining a node to [14](#)
    - synchronizing the system time [16](#)
    - system configuration tasks to complete after setup [22](#)
  - comments
    - how to send feedback about documentation [25](#)
  - critical events
    - configuring EMS to send email notifications [19](#)
    - configuring to forward notifications to syslog server [20](#)
- D**
- documentation
    - how to receive automatic notification of changes to [25](#)
    - how to send feedback about [25](#)
- E**
- emergency events
    - configuring EMS to send email notifications [19](#)
    - configuring to forward notifications to syslog server [20](#)
  - EMS
    - configuring to forward notifications to syslog server [20](#)
    - configuring to send email notifications [19](#)
    - configuring to send SNMP traps [21](#)
- deciding where to send high-severity event notifications [18](#)
  - event notifications
    - configuring EMS for email [19](#)
    - configuring for syslog server [20](#)
    - forwarding to SNMP traphosts [21](#)
  - events
    - deciding where to send high-severity notifications [18](#)
  - express guides
    - deciding where to send high-severity EMS event notifications [18](#)
- F**
- feedback
    - how to send comments about documentation [25](#)
- H**
- high-severity events
    - configuring EMS for email notification [19](#)
    - configuring for syslog server [20](#)
    - forwarding to SNMP traphosts [21](#)
- I**
- information
    - how to send feedback about improving documentation [25](#)
- N**
- Network Time Protocol (NTP)
    - enabling for the cluster [16](#)
  - nodes
    - joining to a cluster [14](#)
    - synchronizing system time across the cluster [16](#)
    - using the Cluster Setup wizard to create the cluster on the first [12](#)
  - NTP
    - See* Network Time Protocol
- S**
- setup

- completing the worksheet for cluster [6](#)
- workflows for cluster [4](#)
- SNMP traphosts
  - configuring EMS for [21](#)
- suggestions
  - how to send feedback about documentation [25](#)
- system configuration
  - tasks to complete after cluster setup [22](#)
- system time
  - synchronizing across the cluster [16](#)

## T

- time

- system, synchronizing across the cluster [16](#)
- traphosts, SNMP
  - configuring EMS for [21](#)
- twitter
  - how to receive automatic notification of documentation changes [25](#)

## W

- workflows
  - for cluster setup [4](#)
- worksheets
  - completing the cluster setup [6](#)