



Centralized Management Using Clusters

This chapter contains the following sections:

- [Overview of Centralized Management Using Clusters](#) , on page 1
- [Cluster Requirements](#), on page 2
- [Cluster Organization](#), on page 2
- [Creating and Joining a Cluster](#), on page 4
- [Managing Clusters](#), on page 11
- [Administering a Cluster from the GUI](#), on page 16
- [Cluster Communication](#), on page 18
- [Loading a Configuration in Clustered Appliances](#) , on page 23
- [Best Practices and Frequently Asked Questions](#), on page 24

Overview of Centralized Management Using Clusters

The Cisco centralized management feature allows you to manage and configure multiple appliances at the same time, reducing administration time and ensuring a consistent configuration across your network. You do not need to purchase additional hardware for managing multiple appliances . The centralized management feature provides increased reliability, flexibility, and scalability within your network, allowing you to manage globally while complying with local policies.

A *cluster* is defined as a set of machines that share configuration information. Within the cluster, machines (appliances) are divided into *groups* ; every cluster will contain at least one group. A given machine is a member of one and only one group. An administrator user can configure different elements of the system on a cluster-wide, group-wide, or per-machine basis, enabling the segmentation of appliances based on network, geography, business unit, or other logical relationships.

Clusters are implemented as a *peer-to-peer* architecture; there is no primary/secondary relationship within a cluster. You may log into any machine to control and administer the cluster. (Some configuration commands, however, are limited. See [Restricted Commands](#), on page 15.)

The user database is shared across all machines in the cluster. That is, there will be only one set of users and one administrator user (with the associated passphrases) for an entire cluster. All machines that join a cluster will share a single administrator passphrase which is referred to as the *admin passphrase* of the cluster.



Note Having more than 20 appliances in a cluster can cause errors in cluster communication.

Cluster Requirements

- Machines in a cluster must have resolvable hostnames in DNS. Alternatively, you can use IP addresses instead, but you may not mix the two.

See [DNS and Hostname Resolution, on page 19](#). Cluster communication is normally initiated using the DNS hostnames of the machines.

- A cluster must consist entirely of machines running the same version of AsyncOS.

See [Upgrading Machines in a Cluster, on page 13](#) for how to upgrade members of a cluster.

- Machines can either join the cluster via SSH (typically on port 22) *or* via the Cluster Communication Service (CCS).

See [Cluster Communication, on page 18](#).

- Once machines have joined the cluster, they can communicate via SSH or via Cluster Communication Service. The port used is configurable. SSH is typically enabled on port 22, and by default CCS is on port 2222, but you can configure either of these services on a different port.

In addition to the normal firewall ports that must be opened for the appliance, clustered machines communicating via CCS must be able to connect with each other via the CCS port. See [Cluster Communication, on page 18](#).

- You must use the Command Line Interface (CLI) command `clusterconfig` to create, join, or configure clusters of machines.

Once you have created a cluster, you can manage non-cluster configuration settings from either the GUI or the CLI.

See [Creating and Joining a Cluster, on page 4](#) and [Administering a Cluster from the GUI, on page 16](#).

- If you have enabled two-factor authentication on your appliance, you can join it to a cluster machine using pre-shared keys. Use the `clusterconfig > prepjoin` command in the CLI to configure this setting.


Or

Disable two-factor authentication on your appliance, before you create or join a cluster. For more information, see [Disabling Two-Factor Authentication](#).

Cluster Organization

Within a cluster, configuration information is divided into 3 groupings or *levels*. The top level describes cluster settings; the middle level describes group settings; and the lowest level describes machine-specific settings.

Figure 1: Cluster Level Hierarchy



Within each level there will be one or more specific members for which settings may be configured; these are referred to as modes. A mode refers to a named member at a specified level. For example, the group “usa” represents one of two group modes in the diagram. While levels are a general term, modes are specific; modes are always referred to by name. The cluster depicted in the above figure has six modes.

Although settings are configured at a given level, they are always configured for a specific mode. It is not necessary to configure settings for all modes within a level. The cluster mode is a special case. Because there can only be one cluster, all settings configured for the cluster mode can be said to be configured at the cluster level.

You should normally configure most settings at the cluster level. However, settings that have been specifically configured at lower levels will override settings configured at higher levels. Thus, you can override cluster-mode settings with group-mode or machine-mode settings.

For example, you might start by configuring the Good Neighbor Table in cluster mode; all machines in the cluster would use that configuration. Then, you might also configure this table in machine mode for machine newyork . In this case, all other machines in the cluster will still use the good neighbor table defined at the cluster level, but the machine newyork will override the cluster settings with its individual machine mode settings.

The ability to override cluster settings for specific groups or machines gives you a lot of flexibility. However, if you find yourself configuring many settings individually in machine mode, you will lose much of the ease of administration that clusters were intended to provide.

Initial Configuration Settings

For most features, when you begin to configure settings for a new mode, those settings will initially be empty by default. There is a distinction between empty settings and having no settings in a mode. As an example, consider a very simple cluster composed of one group and one machine. Imagine that you have an LDAP query configured at the cluster level. There are no settings configured at the group or machine levels:

Cluster	(ldap queries: a, b, c)
Group	
Machine	

Now, imagine that you create new LDAP query settings for the group. The result will be something like this:

Cluster	(ldap queries: a, b, c)
Group	(ldap queries: None)
Machine	

The group-level settings now override the cluster-level setting; however, the new group settings are initially empty. The group mode does not actually have any LDAP queries of its own configured. Note that a machine within this group will inherit this “empty” set of LDAP queries from the group.

Next, you can add an LDAP query to the group, for example:

Cluster	(ldap queries: a, b, c)
Group	(ldap queries: d)
Machine	

Now the cluster level has one set of queries configured while the group has another set of queries. The machine will inherit its queries from the group.

Creating and Joining a Cluster

You cannot create or join a cluster from the Graphical User Interface (GUI). You must use the Command Line Interface (CLI) to create, join, or configure clusters of machines. Once you have created a cluster, you can change configuration settings from either the GUI or the CLI.



Caution If you have enabled two-factor authentication on your appliance, you can join it to a cluster machine using pre-shared keys. Use the `clusterconfig > prepjoin` command in the CLI to configure this setting.

Or

Disable two-factor authentication on your appliance, before you create or join a cluster. For more information, see [Disabling Two-Factor Authentication](#).

The clusterconfig Command

A machine can create or join a cluster only via the `clusterconfig` command.

- When a new cluster is *created*, all of that cluster’s initial settings will be inherited from the machine that creates the cluster. If a machine was previously configured in “standalone” mode, its standalone settings are used when creating the cluster.
- When a machine *joins* a cluster, all of that machine’s clusterable settings will be inherited from the cluster level. In other words, everything except certain machine-specific settings (IP addresses, etc) will be lost and will be replaced with the settings from the cluster and/or the group selected for that machine to join. If a machine was previously configured in “standalone” mode, its standalone settings are used when creating the cluster, and no settings at the machine level are maintained.

If the current machine is not already part of a cluster, issuing the `clusterconfig` command presents the option to join an existing cluster or create a new one.

At this point you can add machines to the new cluster. Those machines can communicate via SSH or CCS.

```
newyork.example.com> clusterconfig
```

```
Do you want to join or create a cluster?
```

```
1. No, configure as standalone.
2. Create a new cluster.
3. Join an existing cluster over SSH.
4. Join an existing cluster over CCS.

[1]> 2

Enter the name of the new cluster.

[ ]> americas

New cluster committed: Wed Jun 22 10:02:04 2005 PDT

Creating a cluster takes effect immediately, there is no need to commit.

Cluster americas

Choose the operation you want to perform:

- ADDGROUP - Add a cluster group.
- SETGROUP - Set the group that machines are a member of.
- RENAMEGROUP - Rename a cluster group.
- DELETEGROUP - Remove a cluster group.
- REMOVEMACHINE - Remove a machine from the cluster.
- SETNAME - Set the cluster name.
- LIST - List the machines in the cluster.
- LISTDETAIL - List the machines in the cluster with detail.
- DISCONNECT - Temporarily detach machines from the cluster.
- RECONNECT - Restore connections with machines that were previously detached.
- PREPJOIN - Prepare the addition of a new machine over CCS.

[ ]>
```

Joining an Existing Cluster

From the host you want to add to the cluster, issue the `clusterconfig` command to join the existing cluster. You can choose to join the cluster over SSH or over CCS (cluster communication service).

In order to join a host to an existing cluster, you must:

- be able to validate the SSH host key of a machine in the cluster
- know the IP address of a machine in the cluster and be able to connect to this machine in the cluster (for example, via SSH or CCS)
- know the administrator passphrase for the admin user on a machine belonging to the cluster

Joining an Existing Cluster over SSH

The following table demonstrates adding the machine losangeles.example.com to the cluster using the SSH option.

```
losangeles.example.com> clusterconfig
```

```
Do you want to join or create a cluster?
```

1. No, configure as standalone.
2. Create a new cluster.
3. Join an existing cluster over SSH.
4. Join an existing cluster over CCS.

```
[1]> 3
```

While joining a cluster, you will need to validate the SSH host key of the remote machine to which you are joining. To get the public host key

```
fingerprint of the remote host, connect to the cluster and run: logconfig ->
hostkeyconfig -> fingerprint.
```

WARNING: All non-network settings will be lost. System will inherit the values set at the group or cluster mode for the non-network settings. Ensure that the cluster settings are compatible with your network settings (e.g. dnsconfig settings)

```
Do you want to enable the Cluster Communication Service on
losangeles.example.com? [N]> n
```

```
Enter the IP address of a machine in the cluster.
```

```
[ ]> IP address is entered
```

```
Enter the remote port to connect to. The must be the normal admin ssh
port, not the CCS port.
```

```
[22]> 22
```

```
Enter the admin passphrase for the cluster.
The administrator passphrase for the clustered machine is entered
```

```
Please verify the SSH host key for IP address:
```

```
Public host key fingerprint: xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx
```

```
Is this a valid key for this host? [Y]> y
```

```
Joining cluster group Main_Group.
```

```
Joining a cluster takes effect immediately, there is no need to commit.
```

```
Cluster americas
```

```
Choose the operation you want to perform:
```

- ADDGROUP - Add a cluster group.
- SETGROUP - Set the group that machines are a member of.

```
- RENAMEGROUP - Rename a cluster group.
- DELETEGROUP - Remove a cluster group.
- REMOVEMACHINE - Remove a machine from the cluster.
- SETNAME - Set the cluster name.
- LIST - List the machines in the cluster.
- LISTDETAIL - List the machines in the cluster with detail.
- DISCONNECT - Temporarily detach machines from the cluster.
- RECONNECT - Restore connections with machines that were previously detached.
- PREPJOIN - Prepare the addition of a new machine over CCS.

[ ]>

(Cluster americas)>
```

Joining an Existing Cluster over CCS

Use CCS instead of SSH if you cannot use SSH. The only advantage of CCS is that only cluster communication happens over that port (no user logins, SCP, etc). To add another machine to an existing cluster via CCS, use the `prepjoin` subcommand of `clusterconfig` to prepare the machine to be added to the cluster. In this example, the `prepjoin` command is issued on the machine `newyork` to prepare the machine `losangeles` to be added to the cluster.

The `prepjoin` command involves obtaining the user key of the host you want to add to the cluster by typing `clusterconfig prepjoin print` in the CLI of that host, and then copying the key into the command line of the host that is currently in the cluster.

Once a machine is already part of a cluster, the `clusterconfig` command allows you to configure various settings for the cluster.

Choose the operation you want to perform:

```
- ADDGROUP - Add a cluster group.
- SETGROUP - Set the group that machines are a member of.
- RENAMEGROUP - Rename a cluster group.
- DELETEGROUP - Remove a cluster group.
- REMOVEMACHINE - Remove a machine from the cluster.
- SETNAME - Set the cluster name.
- LIST - List the machines in the cluster.
- LISTDETAIL - List the machines in the cluster with detail.
- DISCONNECT - Temporarily detach machines from the cluster.
- RECONNECT - Restore connections with machines that were previously detached.
```

- PREPJOIN - Prepare the addition of a new machine over CCS.

```
[ ]> prepjoin
```

Prepare Cluster Join Over CCS

No host entries waiting to be added to the cluster.

Choose the operation you want to perform:

- NEW - Add a new host that will join the cluster.

```
[ ]> new
```

Enter the hostname of the system you want to add.

```
[ ]> losangeles.example.com
```

Enter the serial number of the host mail3.example.com.

```
[ ]> unique serial number is added
```

Enter the user key of the host losangeles.example.com. This can be obtained by typing "clusterconfig prepjoin print" in the CLI on mail3.example.com. Press enter on a blank line to finish.

unique user key from output of prepjoin print is pasted

Host losangeles.example.com added.

Prepare Cluster Join Over CCS

1. losangeles.example.com (serial-number)

Choose the operation you want to perform:

- NEW - Add a new host that will join the cluster.

- DELETE - Remove a host from the pending join list.

```
[ ]>
```

```
(Cluster Americas)> clusterconfig
```

Cluster americas

Choose the operation you want to perform:

- ADDGROUP - Add a cluster group.

- SETGROUP - Set the group that machines are a member of.

- RENAMEGROUP - Rename a cluster group.

- DELETEDGROUP - Remove a cluster group.

- REMOVEMACHINE - Remove a machine from the cluster.

- SETNAME - Set the cluster name.

- LIST - List the machines in the cluster.

- LISTDETAIL - List the machines in the cluster with detail.


```

- DISCONNECT - Temporarily detach machines from the cluster.
- RECONNECT - Restore connections with machines that were previously detached.
- PREPJOIN - Prepare the addition of a new machine over CCS.

[ ]>

```

Joining an Existing Cluster over SSH with Pre-Shared Keys

The following table demonstrates how to join the machine (testmachine.example.com) to the cluster (test_cluster) over SSH using pre-shared keys.

```
testmachine.example.com> clusterconfig
```

```
Do you want to join or create a cluster?
```

1. No, configure as standalone.
2. Create a new cluster.
3. Join an existing cluster over SSH.
4. Join an existing cluster over CCS.

```
[1]> 3
```

While joining a cluster, you will need to validate the SSH host key of the remote machine to which you are joining. To get the public host key

```
fingerprint of the remote host, connect to the cluster and run: logconfig ->
hostkeyconfig -> fingerprint.
```

WARNING: All non-network settings will be lost. System will inherit the values set at the group or cluster mode for the non-network settings. Ensure that the cluster settings are compatible with your network settings (e.g. dnsconfig settings)

```
Do you want to enable the Cluster Communication Service on
testmachine.example.com? [N]>
```

```
Enter the IP address of a machine in the cluster.
```

```
[ ]> IP address entered
```

```
Enter the remote port to connect to. The must be the normal admin ssh
port, not the CCS port.
```

```
[22]>
```

```
Would you like to join this appliance to a cluster using pre-shared keys?
Use this option if you have enabled two-factor authentication on the appliance.) [Y]> yes
```

To join this appliance to a cluster using pre-shared keys, log in to the cluster machine, run the clusterconfig > prepjoin > command, enter the following details, and commit your changes.

```
Host: pod1226-esa07.ibesa
Serial Number: 42291A18D741EDB4C601-BC14E5579F34
User Key:
```

```
ssh-dss
AAAAB3NzaC1kc3MAAACBAJ6Xm+ja4aau9n4D0cJs/gGwEDEUWgERYchhgWApKt6IW+s58I7knGM81rQgQbNdnCO58D
EqaVGmP0Vyb0Ttpgvh6f0mr80OuTgWh9bqg4uiOJvbKvlTvDt0o7//mTkml159zr2KT/qFH+9L5i+8iIMX62R5y+a
6E8JV0BrJCNAAAFAQCmK+Wou9HSribsC0f/5dVoAddxEwAAAIA5p7NR74rlSrs0JWWYItNatE1SamAN+gqCOdUWGPpHT
qdrtBilPQ9tfFoThZElqY4Tx81ku9laasoRLruQ2Z36R3bQGzIn4jzQqujvbxTvLK9eLoSr8yFbEE3ZvuUo0+vhDn
LIDX2N65AQSQsTaOrKX+yQZ8yAVt48CscptsDrgAAAIAVROGlWoSl8g3FFm2eRTa+/oZ+cMjv+pSZiZoiUCoaIlouc
ulZDpN413QBnf6p/3D8wVD8m5uo8O4N/HXasAMektZvGoP4Sf+shItPuISrv3lrMTEYsD0sqVcMc7vIXUeD2jpk7MB
ooVktZB/rdTbNMfXrhDkNJ2IAPQQiUKVnw==
```

Before you proceed to the next step, make sure you add the 'Host', Serial Number' and 'User Key' details to the cluster machine.

Would you like to continue? [Y]> **yes**

Joining cluster group Main_Group.

Joining a cluster takes effect immediately, there is no need to commit.

Cluster **test_cluster**

Choose the operation you want to perform:

- ADDGROUP - Add a cluster group.
- SETGROUP - Set the group that machines are a member of.
- RENAMEGROUP - Rename a cluster group.
- DELETEDGROUP - Remove a cluster group.
- REMOVEMACHINE - Remove a machine from the cluster.
- SETNAME - Set the cluster name.
- LIST - List the machines in the cluster.
- LISTDETAIL - List the machines in the cluster with detail.
- DISCONNECT - Temporarily detach machines from the cluster.
- RECONNECT - Restore connections with machines that were previously detached.
- PREPJOIN - Prepare the addition of a new machine over CCS.

[]>

(Cluster test_cluster)>

Adding Groups

All clusters must contain at least one group. When you create a new cluster, a default group called **Main_Group** is created automatically. However, you may decide to create additional groups within your cluster. This example shows how to create additional groups within an existing cluster and assign machines to the new group(s).

Procedure

-
- Step 1** Issue the `clusterconfig` command.
 - Step 2** Choose the `addgroup` subcommand and enter the name of the new group.
 - Step 3** Use the `setgroup` subcommand to choose machines for the new group.
-

Managing Clusters

Administering a Cluster from the CLI

For machines that are part of a cluster, the CLI can be switched into different modes. Recall that a mode refers to a specific, named, member of a level.

The CLI mode determines precisely where a configuration setting will be modified. The default is “machine” mode for the machine the user logged into, the “login host.”

Use the `clustermode` command to switch between different modes.

Table 1: Administering Clusters

Command Example	Description
<code>clustermode</code>	Prompt to switch cluster mode
<code>clustermode group northamerica</code>	Switch to group mode for the group “northamerica”
<code>clustermode machine losangeles.example.com</code>	Switch to machine mode for the machine “losangeles”

The prompt in the CLI changes to indicate your current mode.

```
(Cluster Americas)>
```

or

```
(Machine losangeles.example.com)>
```

In machine mode, the prompt will include the fully qualified domain name of the machine.

Copying and Moving Settings

All non-restricted (see [Restricted Commands, on page 15](#)) commands have new operations: **CLUSTERSHOW** and **CLUSTERSET**. **CLUSTERSHOW** is used to show in which modes a command is configured (see [New Operation Added, on page 14](#)). The **CLUSTERSET** operation allows you to move or copy the current settings

(configurable with the current command) from one mode to another or between levels (e.g. from a machine to a group).

A *copy* retains the settings for the current mode. A *move* resets (clears) the configuration of the current mode; i.e., following a move, no settings will be configured for the current mode.

For example, if you have configured Good Neighbor Table settings (the **destconfig** command) for group **northamerica**, and you decide that you want the entire cluster to have these settings, you can use the **clusterset** operation from within the **destconfig** command to copy (or move) the current settings to the cluster mode. (See [Experimenting with New Configurations, on page 12.](#))



Caution Exercise caution when moving or copying configuration settings to avoid inconsistent dependencies. For example, if you move or copy listeners with disclaimer stamping configured to another machine, and that new machine does not have the same disclaimers configured, disclaimer stamping will not be enabled on the new machine.

Experimenting with New Configurations

One of the most advantageous ways to use clusters is to experiment with new configuration settings. First you make changes at the machine mode, in an isolated environment. Then, when you are satisfied with your configuration, you move those configuration changes up to the cluster mode to make them available on all machines.

The following example shows the steps to change a listener setting on one machine and then publish the setting to the rest of the cluster when ready. Because listeners are normally configured at the cluster level, the example starts by pulling the configuration down to machine mode on one machine before making and testing the changes. You should test experimental changes of this type on one machine before making the change to the other machines in the cluster.

Procedure

- Step 1** Use the **clustermode cluster** command to change to the cluster mode.
Remember: the **clustermode** command is the CLI command you use to change modes to the cluster, group, and machine levels.
- Step 2** Type **listenerconfig** to see the listener settings configured for the cluster.
- Step 3** Choose the machine you want to experiment with, then use the **clusterset** command to copy settings from the cluster “down” to machine mode.
- Step 4** Use the **clustermode** command to navigate to machine mode for the experimental machine, e.g.:
clustermode machine newyork.example.com
- Step 5** In machine mode, on the experimental machine, issue the **listenerconfig** command to make changes specifically for the experimental machine.
- Step 6** Commit the changes.
- Step 7** Continue to experiment with the configuration changes on the experimental machine, remembering to commit the changes.

- Step 8** When you are ready to apply your new settings to all the other machines, use the `clusterset` command to move the settings up to the cluster mode.
- Step 9** Commit the changes.
-

Leaving a Cluster Permanently (Removal)

You use the `REMOVEMACHINE` operation of `clusterconfig` to remove a machine permanently from a cluster. When a machine is permanently removed from a cluster, its configuration is “flattened” such that it will work the same as it did when it was part of the cluster. For example, if there is only a cluster-mode Global Unsubscribe table, the Global Unsubscribe table data will be copied to the machine’s local configuration when the machine is removed from the cluster.

Upgrading Machines in a Cluster

A cluster does not allow the connected machines to have different versions of AsyncOS.

Before you install an AsyncOS upgrade, you need to disconnect each machine in the cluster via the `clusterconfig` command. After you upgrade all the machines, the cluster can be reconnected via the `clusterconfig` command. You can have two separate clusters running while you upgrade machines to the same version. You can also upgrade clustered machines on the GUI Upgrades page.

You can download the upgrade in the background so that you do not need to disconnect the cluster machines until you are ready to install the upgrade.



Note If you use the upgrade command before disconnecting the individual machine from the cluster, AsyncOS disconnects all the machines in the cluster. Cisco Systems recommends that you disconnect each machine from the cluster before upgrading it. Then, other machines can continue working as a cluster until each is disconnected and upgraded.

Procedure

- Step 1** On a machine in the cluster, use the `disconnect` operation of `clusterconfig`. For example, to disconnect the machine `losangeles.example.com`, type `clusterconfig disconnect losangeles .example.com`. No commit is necessary.
- Step 2** Optionally, use the `suspendlistener` command to halt acceptance of new connections and messages during the upgrade process.
- Step 3** Issue the upgrade command to upgrade AsyncOS to a newer version.
- Note** Disregard any warnings or confirmation prompts about disconnecting all of the machines in the cluster. Because you have disconnected the machine, AsyncOS does not disconnect the other machines in the cluster at this point.
- Step 4** Select the version of AsyncOS for the machine. The machine will reboot after the upgrade is complete.
- Step 5** Use the `resume` command on the upgraded machine to begin accepting new messages.
- Step 6** Repeat steps 1 - 5 for each machine in the cluster.

Note After you disconnect a machine from the cluster, you cannot use it to change the configurations of other machines. Although you can still modify the cluster configuration, do not change it while machines are disconnected because settings can become unsynchronized.

Step 7 After you have upgraded all the machines, use the reconnect operation of clusterconfig for each upgraded machine to reconnect it. For example, to reconnect the machine losangeles.example.com , type clusterconfig reconnect losangeles .example.com. Note that you can only connect a machine to a cluster that is running the same version of AsyncOS.

CLI Command Support

All Commands Are Cluster-aware

All CLI commands in AsyncOS are now cluster-aware. The behavior of some commands will change slightly when issued in a cluster mode. For example, the behavior of the following commands changes when issued on a machine that is part of a cluster:

The commit and clearchanges Commands

commit

The commit command commits all changes for all three levels of the cluster, regardless of which mode you are currently in.

commitdetail

The commitdetail command provides details about configuration changes as they are propagated to all machines within a cluster.

clearchanges

The clearchanges (clear) command clears all changes for all three levels of the cluster, regardless of which mode you are currently in.

New Operation Added

CLUSTERSHOW

Within each command, there is now a CLUSTERSHOW operation that allows you to see in which modes a command is configured.

When you enter a CLI command to perform an action that will be overridden by existing settings at a lower level, you will be presented with a notification. For example, if you are in cluster mode and enter a command, you may see a notification like this:

Note: Changes to these settings will not affect the following groups and machines because they are overriding the cluster-wide settings:

```
East_Coast, West_Coast
```

```
facilities_A, facilities_B, receiving_A
```

A similar message would be printed if you are editing settings for a group mode.

Restricted Commands

Most CLI commands and their corresponding GUI pages can be run in any mode (cluster, group, or machine). However, some commands and pages are restricted to one mode only.

The system interface (either the GUI and the CLI) will always will make it clear that a command is restricted and how it is restricted. It is easy to switch to the appropriate mode for configuring the command.

- In the GUI, use the “Change Mode” menu or the “Settings for this features are currently defined at:” links to switch modes.
- In the CLI, use the `clustermode` command to switch modes.

Table 2: Commands Restricted to Cluster Mode

<code>clusterconfig</code>	<code>sshconfig</code>
<code>clustercheck</code>	<code>userconfig</code>
<code>passwd</code>	

If a you try to run one of these commands in group or machine mode, you will be given a warning message and the opportunity to switch to the appropriate mode.



Note The `passwd` command is a special case because it needs to be usable by guest users. If a guest user issues the `passwd` command on a machine in a cluster, it will not print the warning message but will instead just silently operate on the cluster level data without changing the user’s mode. All other users will get the above written behavior (consistent with the other restricted configuration commands).

The following commands are restricted to *machine mode* :

<code>antispamstatus</code>	<code>etherconfig</code>	<code>resume</code>	<code>suspenddel</code>
<code>antispamupdate</code>	<code>featurekey</code>	<code>resumedel</code>	<code>suspendlistener</code>
<code>antivirusstatus</code>	<code>hostrate</code>	<code>resumelister</code>	<code>techsupport</code>
<code>antivirusupdate</code>	<code>hoststatus</code>	<code>rollovernow</code>	<code>tophosts</code>
<code>bouncerecipients</code>	<code>interfaceconfig</code>	<code>routeconfig</code>	<code>topin</code>
<code>deleterecipients</code>	<code>ldapflush</code>	<code>sbstatus</code>	<code>trace</code>
<code>delivernow</code>	<code>ldaptest</code>	<code>setgateway</code>	<code>version</code>
<code>diagnostic</code>	<code>nslookup</code>	<code>sethostname</code>	<code>vofflush</code>

dnsflush	quarantineconfig	settime	vofstatus
dnslistflush	rate	shutdown	workqueue
dnslisttest	reboot	status	
dnsstatus	resetcounters	suspend	

If a you try to run one of the commands above in cluster or group mode, you will be given a warning message and the opportunity to switch to an appropriate mode.

The following commands are further restricted to the *login host* (i.e., the specific machine you are logged into). These commands require access to the local file system.

Table 3: Commands Restricted to Login Host Mode

last	resetconfig	tail	upgrade
ping	supportrequest	telnet	who

Administering a Cluster from the GUI

Although you cannot create or join clusters or administer cluster specific settings from the GUI (the equivalent of the **clusterconfig** command), you can browse machines in the cluster, create, delete, copy, and move settings among the cluster, groups, and machines (that is, perform the equivalent of the **clustermode** and **clusterset** commands) from within the GUI.

The Incoming Mail Overview page is an example of a command that is restricted to the login host, because the Mail Flow Monitoring data you are viewing is stored on the local machine. To view the Incoming Mail Overview reports for another machine, you must log into the GUI for that machine.

Note the URL in the browser's address field when clustering has been enabled on an appliance . The URL will contain the word **machine** , **group** , or **cluster** as appropriate. For example, when you first log in, the URL of the Incoming Mail Overview page will appear as:

https:// hostnamemachine/serial_number /monitor/incoming_mail_overview



Note The Incoming Mail Overview and Incoming Mail Details pages on the Monitor menu are restricted to the login machine.

The Mail Policies, Security Services, Network, and System Administration tabs contain pages that are not restricted to the local machine. If you click the Mail Policies tab, the centralized management information in the GUI changes.

Figure 2: Centralized Management Feature in the GUI: No Settings Defined

The screenshot shows the 'Incoming Mail Policies' interface. At the top, the mode is set to 'Machine: example.com'. Below this, it indicates that settings are being inherited from the cluster mode 'americas'. A section titled 'Find Policies' includes an 'Email Address' search field and radio buttons for 'Recipient' and 'Sender'. The main area displays a table of policies with columns for Order, Policy Name, Anti-Spam, Anti-Virus, Virus Outbreak Filters, Content Filters, and Delete. A key at the bottom indicates that light grey cells represent default settings, yellow cells represent custom settings, and greyed-out cells represent disabled settings.

Order	Policy Name	Anti-Spam	Anti-Virus	Virus Outbreak Filters	Content Filters	Delete
	Default Policy	IronPort Positive: Deliver Suspected: Disabled	Repaired: Deliver Encrypted: Deliver Unscannable: Deliver Virus Positive: Drop	Enabled	Disabled	

In the above figure, the machine is inheriting all of its configuration settings for the current feature from the cluster mode. The settings being inherited in a light grey (preview). You can retain these settings or change them, overriding the cluster level settings for this machine.



Note The inherited settings (preview display) will always show the settings inherited from the cluster. Use caution when enabling or disabling dependent services among group and cluster levels. For more information, see [Copying and Moving Settings, on page 11](#).

If you click the Override Settings link, you are taken to a new page for that feature. This page allows you to create new configuration settings for machine mode. You may begin with the default settings, or, if you've already configured settings in another mode, you can copy those settings to this machine.

Figure 3: Centralized Management Feature in the GUI: Create New Settings

The screenshot shows the 'Creating New Settings for Machine: example.com' dialog. It includes a note stating that creating new settings will override the settings currently inherited from the cluster mode 'americas'. There are two radio button options: 'Start with default settings' (which is selected) and 'Copy from: Cluster: americas'. 'Cancel' and 'Submit' buttons are also visible.

Alternatively, as shown in *Figure Centralized Management Feature in the GUI: No Settings Defined*, you can also navigate to modes where this configuration setting is already defined. The modes are listed in the lower half of the centralized management box, under “Settings for this feature are currently defined at:”. Only those modes where the settings are actually defined will be listed here. When you view a page for settings that are defined in (and inherited from) another mode, the page will display those settings for you.

If you click on one of the listed modes (for example, the Cluster: Americas link as shown in *Figure Centralized Management Feature in the GUI: No Settings Defined*), you will be taken to a new page that allows you to view and manage the settings for that mode.

Figure 4: Centralized Management Feature in GUI: Settings Defined

The screenshot shows the top portion of the GUI for the 'Cluster: americas' mode. It includes a 'Change Mode...' dropdown menu and a breadcrumb trail for 'Centralized Management Options'.

When settings are defined for a given mode, the centralized management box is displayed on every page in a minimized state. Click the “Centralized Management Options” link to expand the box to show a list of options available for the current mode with respect to the current page. Clicking the “Manage Settings” button allows you to copy or move the current settings to a different mode or to delete those settings completely.

For example, in the following figure, the Centralized Management Options link has been clicked to present the available options.


Figure 5: Centralized Management Feature in GUI: Manage Settings



On the right side of the box is the “Change Mode” menu. This menu displays your current mode and provides the ability to navigate to any other mode (cluster, group, or machine) at any time.

Figure 6: The Change Mode Menu

Incoming Mail Policies



When you navigate to a page that represents a different mode, the “Mode —” text on the left side of the centralized management box will flash yellow, briefly, to alert you that your mode has changed.

Some pages within certain tabs are restricted to machine mode. However, unlike the Incoming Mail Overview page (which is restricted to the current login host), these pages can be used for any machine in the cluster.

Figure 7: Centralized Management Feature: Machine Restricted



Choose which machine to administer from the Change Mode menu. You will see a brief flashing of the text to remind you that you have changed modes.

Cluster Communication

Machines within a cluster communicate with each other using a mesh network. By default, all machines connect to all other machines. If one link goes down, other machines will not be prevented from receiving updates.

By default, all intra-cluster communication is secured with SSH. Each machine keeps an in-memory copy of the route table and makes in-memory changes as necessary if links go down or up. Each machine also performs a periodic “ping” (every 1 minute) of every other machine in the cluster. This ensures up-to-date link status and maintains the connections in case a router or NAT has a timeout.



Note If your appliances are in a cluster mode, and you plan to access data (not related to configuration, for example, viewing messages present in the quarantine or refreshing reports at a fast rate) of another appliance remotely; there will be cluster reconnection attempts that can generate alerts and errors. The appliances automatically will reconnect and manual intervention is not required..

DNS and Hostname Resolution

DNS is required to connect a machine to the cluster. Cluster communication is normally initiated using the DNS hostnames of the machines (not the hostname of an interface on the machine). A machine with an unresolvable hostname would be unable to actually communicate with any other machines in the cluster, even though it is technically part of the cluster.

Your DNS must be configured to have the hostname point to the correct IP interface on the appliance that has SSH or CCS enabled. This is very important. If DNS points to another IP address that does not have SSH or CCS enabled it will not find the host. Note that centralized management uses the “main hostname,” as set with the `sethostname` command, not the per-interface hostname.

If you use an IP address to connect to another machine in the cluster, the machine you connect to must be able to make a reverse look up of the connecting IP address. If the reverse look up times out because the IP address isn't in the DNS, the machine cannot connect to the cluster.

Clustering, Fully Qualified Domain Names, and Upgrading

DNS changes can cause a loss of connectivity after upgrading AsyncOS. Please note that if you need to change the fully qualified domain name of a machine in the cluster (not the hostname of an interface on a machine in the cluster), you must change the hostname settings via `sethostname` and update the DNS record for that machine *prior* to upgrading AsyncOS.

Cluster Communication Security

Cluster Communication Security (CCS) is a secure shell service similar to a regular SSH service. Cisco implemented CCS in response to concerns regarding using regular SSH for cluster communication. SSH communication between two machines opens regular logins (admin, etc.) on the same port. Many administrators prefer not to open regular logins on their clustered machines.

Tip: never enable Cluster Communication Services, even though it is the default, unless you have firewalls blocking port 22 between some of your clustered machines. Clustering uses a full mesh of SSH tunnels (on port 22) between all machines. If you have already answered Yes to enabling CCS on any machine, remove all machines from the cluster and start again. Removing the last machine in the cluster removes the cluster.

CCS provides an enhancement where the administrator can open up cluster communication, but not CLI logins. By default, the service is disabled. You will be prompted to enable CCS from the `interfaceconfig` command when you are prompted to enable other services. For example:

```
Do you want to enable SSH on this interface? [Y]>
```

```
Which port do you want to use for SSH?
```

```
[22]>
```

```
Do you want to enable Cluster Communication Service on this interface?
```

```
[N]> y
```

```
Which port do you want to use for Cluster Communication Service?
```

```
[2222]>
```

The default port number for CCS is 2222. You may change this to another open, unused, port number if you prefer. After the join is complete and the joining machine has all the configuration data from the cluster, the following question is presented:

```
Do you want to enable Cluster Communication Service on this interface? [N]> y
```

```
Which port do you want to use for Cluster Communication Service?
```

```
[2222]>
```

Cluster Consistency

The machines that are “cluster aware” will continually verify network connections to other machines within the cluster. This verification is done by periodic “pings” sent to other machines in the cluster.

If all attempts to communicate with a particular machine fail, then the machine that has been trying to communicate will log a message saying that the remote host has disconnected. The system will send an alert to the administrator that the remote host went down.

Even if a machine is down, the verification pings will continue to be sent. When a machine rejoins the cluster network, a synchronization command will be issued so that any previously offline machines can download any updates. The synchronization command will also determine if there have been any changes on one side but not the other. If so, then the previously down machine will silently download the updates.

Disconnect/Reconnect

A machine may be disconnected from a cluster. Occasionally, you may intend to deliberately disconnect the machine, for example, because you are upgrading the machine. A disconnect could also occur by accident, for example, due to a power failure or other software or hardware error. A disconnect can also occur if one appliance attempts to open more than the maximum number of SSH connections allowed in a session. A machine that is disconnected from a cluster can still be accessed directly and configured; however, any changes made will not be propagated to other machines within the cluster until the disconnected machine becomes reconnected.

When a machine reconnects to the cluster, it tries to reconnect to all machines at once.

In theory, two machines in a cluster that are disconnected could commit a similar change to their local databases at the same time. When the machines are reconnected to the cluster, an attempt will be made to synchronize these changes. If there is a conflict, the most recent change is recorded (supersedes any other changes).

During a commit, the appliance checks every variable that is being changed. The commit data includes version information, sequence identification numbers, and other information that can be compared. If the data you are about to change is found to be in conflict with previous changes, you will be given the option to discard your changes. For example, you might see something like this:

```
(Machine mail3.example.com)> clustercheck
```

```
This command is restricted to "cluster" mode. Would you like to switch to "cluster"
```

```
mode? [Y]> y
Checking Listeners (including HAT, RAT, bounce profiles)...
Inconsistency found!
Listeners (including HAT, RAT, bounce profiles) at Cluster enterprise:
mail3.example.com was updated Mon Sep 12 10:59:17 2005 PDT by 'admin' on
mail3.example.com
test.example.com was updated Mon Sep 12 10:59:17 2005 PDT by 'admin' on
mail3.example.com
How do you want to resolve this inconsistency?
1. Force entire cluster to use test.example.com version.
2. Force entire cluster to use mail3.example.com version.
3. Ignore.
[1]>
```

If you choose not to discard your changes, they are still intact (but uncommitted). You can review your changes against the current settings and decide how to proceed.

You can also use the `clustercheck` command at any time to verify that the cluster is operating correctly.

```
losangeles> clustercheck
Do you want to check the config consistency across all machines in the cluster? [Y]> y
Checking losangeles...
Checking newyork...
No inconsistencies found.
```

Interdependent Settings

It is recommended that you avoid configuring the following settings on the appliance .

In a centrally managed environment, some interdependent settings are configured in different modes. The flexibility of the configuration model allows you to configure settings at multiple modes, and the laws of inheritance govern which settings will be used on a per-machine basis. However, some settings have dependencies on other settings, and the availability of the dependent settings' configuration is not limited to settings at the same mode. Thus, it is possible to configure a setting for one level that references a setting that is configured for a specific machine at a different level.

The most common example of an interdependent setting involves a select field on a page that pulls data from a different cluster section. For example, the following features can be configured in different modes:

- using LDAP queries
- using dictionaries or text resources
- using bounce or SMTP authentication profiles.

Within centralized management, there are restricted and non-restricted commands. (See [Restricted Commands, on page 15](#).) Non-restricted commands are generally configuration commands that can be shared across the cluster.

The `listenerconfig` command is an example of a command that can be configured for all machines in a cluster. Non-restricted commands represent commands that can be mirrored on all machines in a cluster, and do not require machine-specific data to be modified.

Restricted commands, on the other hand, are commands that only apply to a specific mode. For example, users cannot be configured for specific machines — there must be only one user set across the whole cluster. (Otherwise, it would be impossible to login to remote machines with the same login.) Likewise, since the Mail Flow Monitor data, System Overview counters, and log files are only maintained on a per-machine basis, these commands and pages must be restricted to a machine.

You will notice that while Scheduled Reports may be configured identically across the whole cluster, the viewing of reports is machine-specific. Therefore, within a single Scheduled Reports page in the GUI, configuration must be performed at the cluster mode, but viewing of reports must be done at the machine mode.

The System Time pages encompass the `settz`, `ntpconfig`, and `settime` commands, and thus represents a mixture of restricted and non-restricted commands. In this case, `settime` must be restricted to machine-only modes (since time settings are specific for machine), while `settz` and `ntpconfig` may be configured at cluster or group modes.

Figure 8: Example of Interdependent Settings

The screenshot shows the 'Edit Listener' configuration page. At the top, the mode is 'Cluster: americas'. Below this, the 'Listener Settings' section is visible. The 'Name' field is 'IncomingMail'. The 'Type of Listener' is 'Public'. The 'Interface' is 'Data 1' and the 'TCP Port' is '25'. The 'Bounce Profile' is 'Default'. The 'Disclaimer Above' is 'None'. The 'Disclaimer Below' dropdown menu is open, showing 'None' and 'disclaimer (- Unavailable on Machine: buttercup.run)'. The 'SMTP Authentication Profile' is 'test'. The 'Certificate' is 'test'. There are three expandable sections: 'SMTP Address Parsing Options', 'Advanced', and 'LDAP Queries'. The 'SMTP Call-Ahead Profile' is 'SMTP_Call_Ahead'. At the bottom, there are 'Cancel' and 'Submit' buttons.

In this representation, the listener “IncomingMail” is referencing a footer named “disclaimer” that has been configured at the machine level only. The drop-down list of available footer resources shows that the footer is not available on the machine “buttercup.run” which is also available in the cluster. There are two solutions to this dilemma:

- promote the footer “disclaimer” from the machine level to the cluster level
- demote the listener to the machine level to remove the interdependency

In order to fully maximize the features of a centrally managed system, the former solution is preferred. Be aware of interdependencies among settings as you tailor the configuration of your clustered machines.

Loading a Configuration in Clustered Appliances

AsyncOS allows you to load a cluster configuration in clustered appliances . You can load the cluster configuration in the following scenarios:

- If you are migrating from an on-premise environment to a hosted environment and you want to migrate the on-premise cluster configuration to the hosted environment.
- If an appliance in a cluster is down or needs to be retired and you want to load the configuration from this appliance to a new appliance that you plan to add to the cluster.
- If you are adding more appliances to your cluster and you want to load the configuration from one of the existing appliances in the cluster to the newly added appliances .
- If you want to load a backed-up configuration to a cluster.

Depending on your requirements, you can load a cluster configuration or appliance configuration from a valid cluster configuration file.



Note You cannot load the configuration of a standalone appliance on a clustered appliance .

Before You Begin

- Make sure that you have a valid and complete XML configuration. See [Loading a Configuration File](#).
- Create a backup of the current configuration of the appliance to which you plan to load the configuration. See [Saving and Exporting the Current Configuration File](#).
- Create a cluster setup with all the appliances that you plan to have in your setup. See [Creating and Joining a Cluster, on page 4](#).



Note You can have all the appliances under one group. Ensure that the interfaces for cluster communication in your setup have same names and SSH and CCS settings as in the XML configuration.

Procedure

Step 1 Click **System Administration > Configuration File**.

Step 2 Choose the cluster from the **Mode** drop-down menu.

Step 3 Depending on whether you want to a load cluster or appliance configuration, do one of the following:

- **Load Cluster Configuration**

- a. In the Load Configuration section, choose **Cluster** from the drop-down list.
- b. Load the cluster configuration, and click **Load**. See [Loading a Configuration File](#).
- c. Assign groups from the loaded configuration to the appliancea in the cluster, and copy appliance configuration from the appliances in the selected group to the respective appliances . Use the **Group Configuration** and **Appliance Configuration** drop-down lists.

If you do not want to copy an appliance configuration, choose **Don't Copy** from the **Appliance Configuration** drop-down list.

1. Review the configuration. Click **Review**.
2. Click **Confirm**.
3. Click **Continue**.

• **Load Appliance Configuration**

- a. In the Load Configuration section, choose **Appliance in cluster** from the drop-down list.
- b. Load the configuration, and click **Load**. See [Loading a Configuration File](#). Note that you cannot load the configuration of a standalone appliance on a clustered appliance .
- c. Choose the appliance configuration from the loaded configuration and the intended appliance in the cluster to which you want to load the configuration. Use the drop-down lists.
- d. Click **OK**.
- e. Click **Continue**.
- f. To load the appliance configuration to more appliances , repeat Step a through Step e.

Step 4 Review the network settings of the clustered appliances , and commit your changes.

Best Practices and Frequently Asked Questions

Best Practices

When you create the cluster, the machine you happen to be logged into is automatically added to the cluster as the first machine, and also added to the Main_Group. Its machine level settings effectively get moved to the cluster level as much as possible. There are no settings at the group level, and the only settings left at the machine level are those which do not make sense at the cluster level, and cannot be clustered. Examples are IP addresses, featurekeys, etc.

Leave as many settings at the cluster level as possible. If only one machine in the cluster needs a different setting, copy that cluster setting to the machine level for that machine. Do not move that setting. If you move a setting which has no factory default (e.g. HAT table, SMTPROUTES table, LDAP server profile, etc.), the systems inheriting the cluster settings will have blank tables and will probably not process email.

To have that machine re-inherit the cluster setting, manage the CM settings and delete the machine setting. You will only know if a machine is overriding the cluster setting when you see this display:

Settings are defined:

To inherit settings from a higher level: Delete Settings for this feature at this mode.

You can also Manage Settings.

Settings for this feature are also defined at:

Cluster: xxx

Or this display:

Delete settings from:

Cluster: xxx

Machine: yyyy.domain.com

Copy vs Move

When to copy: when you want the cluster to have a setting, and a group or machine to also have no settings or to have different settings.

When to move: when you want the cluster to have no setting at all, and for the group or machine to have the settings.

Good CM Design Practices

When you LIST your CM machines, you want to see something like this:

```
cluster = CompanyName
```

```
Group Main_Group:
```

```
Machine lab1.example.com (Serial #: XXXXXXXXXXXXXXX-XXXXXXX)
```

```
Machine lab2.example.com (Serial #: XXXXXXXXXXXXXXX-XXXXXXX)
```

```
Group Paris:
```

```
Machine lab3.example.com (Serial #: XXXXXXXXXXXXXXX-XXXXXXX)
```

```
Machine lab4.example.com (Serial #: XXXXXXXXXXXXXXX-XXXXXXX)
```

```
Group Rome:
```

```
Machine lab5.example.com (Serial #: XXXXXXXXXXXXXXX-XXXXXXX)
```

```
Machine lab6.example.com (Serial #: XXXXXXXXXXXXXXX-XXXXXXX)
```

Be careful not to lose track of the level at which you are making changes. For example, if you have changed the name of your Main_Group (using RENAMEGROUP) to London, it will look like this:

```
cluster = CompanyName
```

```
Group London:
```

```
Machine lab1.cable.nu (Serial #: 000F1FF7B3F0-CF2SX51)
```

```
...
```

However, this configuration tends to confuse many administrators, because they begin making changes to the London systems at the group level, and they stop using the Cluster level as the normal configuration level for basic settings.

Tip: it is not a good practice to have a group with the same name as the cluster, e.g. cluster London, group London. If you are using site names for group names, it is not good practice to have a cluster name that refers to a location.

The correct method, as explained above, is to leave as many settings at the cluster level as possible. In most cases you should leave your primary site or main collection of machines in the Main_Group, and use groups for your additional sites. This is true even if you consider that both sites are “equal.” Remember, CM has no primary/secondary servers — all clustered machines are peers.

Tip: if you will be using extra groups you can easily prepare the groups before those extra machines are joined to the cluster.

Best Practices for Accessing Spam or Policy Quarantines in Cluster Setup

Accessing spam or policy quarantines of other appliances in a cluster from the logged-in appliance may cause excessive CPU usage on the logged-in appliance. To avoid this scenario, you can access the spam or policy quarantines by logging into the required appliances.

Procedures: Configuring an Example Cluster

To configure this example cluster, log out of all GUIs on all machines before running `clusterconfig`. Run `clusterconfig` on any one of the primary site machines. You will then join to this cluster only the other local and remote machines that need the maximum possible shared settings (allowing for the machine only-settings like IP address). The `clusterconfig` command cannot be used to join a remote machine to the cluster — you must use the CLI on the remote machine and run `clusterconfig` (“join an existing cluster”).

In our example above we log in to `lab1`, run `clusterconfig` and create a cluster called `CompanyName`. We have only one machine with identical requirements, so we log in to `lab2`, and `saveconfig` the existing configuration (it will be drastically altered when it inherits most of `lab1` settings.) On `lab2` we can then use `clusterconfig` to join an existing cluster. Repeat if you have additional machines at this site needing similar policies and settings.

Run `CONNSTATUS` to confirm that DNS resolves correctly. As machines are joined to the cluster, the new machines inherit almost all of their settings from `lab1` and their older settings are lost. If they are production machines you will need to anticipate if mail will still be processed using the new configuration instead of their previous configuration. If you remove them from the cluster, they will not revert to their old, private configs.

Next, we count the number of exceptional machines. If there is only one, it should receive a few extra machine level settings and you will not need to create an extra group for it. Join it to the cluster and begin copying settings down to the machine level. If this machine is an existing production machine you must back up the configuration and consider the changes to mail processing as above.

If there are two or more, as in our example, decide if those two will share any settings with each other that are not shared with the cluster. In that case, you will be creating one or more groups for them. Otherwise, you will make machine level settings for each, and do not need to have extra groups.

In our case we want to run `clusterconfig` from the CLI on any of the machines already in the cluster, and select `ADDGROUP`. We will do this twice, once for `Paris` and once for `Rome`.

Now you can begin using the GUI and CLI to build configuration settings for the cluster and for ALL the groups, even if the groups have no machines in them yet. You will only be able to create machine specific settings for machines *after* they have joined the cluster.

The best way to create your override or exceptional settings is to copy the settings from the higher (e.g. cluster) level down to a lower (e.g. group) level.

For example, after creating the cluster our `dnsconfig` settings initially looked like this:

Configured at mode:

Cluster: Yes

Group `Main_Group`: No

Group `Paris`: No

Group `Rome`: No

Machine `lab2.cable.nu`: No

If we "Copy to Group" the DNS settings, it will look like this:

Configured at mode:

Cluster: Yes

Group Main_Group: No

Group Paris: Yes

Group Rome: No

Machine lab2.cable.nu: No

Now you can edit the Paris group-level DNS settings, and other machines in the Paris group will inherit them. Non-Paris machines will inherit the cluster settings, unless they have machine-specific settings. Besides DNS settings, it is common to create group level settings for SMTPROUTES.



Tip When using the CLI CLUSTERSET function in various menus, you can use a special option to copy settings to All Groups, which is not available through the GUI.

Complete listeners will be automatically inherited from the group or cluster, and you normally only create these on the first system in the cluster. This reduces administration considerably. However, for this to work *you must name the Interfaces identically throughout your group or cluster*.

Once the settings are defined correctly at the group level, you can join machines to the cluster and make them part of this group. This requires two steps:

First, to join our remaining 4 systems to the cluster, we run clusterconfig on each. The larger and more complex the cluster, the longer it takes to join, and this can take several minutes. You can monitor the joining progress with the LIST and CONNSTATUS sub-commands. After the joins are complete you can use SETGROUP to move the machines from the Main_Group into Paris and Rome. There is no way to avoid the fact that initially, all machines added to the cluster inherit the Main_Group settings, not the Paris and Rome settings. This could affect mail flow traffic if the new systems are already in production.



Tip Do not make your lab machines part of the same cluster as your production machines. Use a new cluster name for lab systems. This provides an added layer of protection against unexpected changes (someone changing a lab system and accidentally losing production mail, for example).

Summary of GUI Options for Using CM Settings Other Than the Cluster Default

Override settings, and start with default settings. For example, the default settings for the SMTPROUTES configuration is a blank table, which you can then build from scratch.

Override settings, but start with a copy of the settings currently inherited from Cluster xxx, or group yyy. For example, you may want a new copy of the SMTPROUTES table at the group level which is initially identical to the cluster table. All appliances that are contained in that same group (SETGROUP) will get this table. Machines not in the group will still use the cluster level settings. Changing the SMTPROUTES on this independent copy of the table will not affect other groups, machines inheriting the cluster settings, or machines where the setting is defined at the individual machine level. This is the most common selection.

Manage settings, a sub-menu of Centralized Management Options. From this menu you can copy as above, but you can also move or delete settings. If you move the SMTPROUTES to a group or machine level, then the routes table will be blank at the cluster level but will exist at the more specific level.

Manage settings. Continuing our SMTPROUTES example, using the delete option will also result in a blank SMTPROUTES table for the cluster. This is fine if you previously configured definitions for SMTPROUTES

at the group level or machine levels. It is not a best practice to delete the cluster level settings and rely only on group or machine settings. The cluster-wide settings are useful as defaults on newly added machines, and keeping them reduces the number of group or site settings you have to maintain by one.

Setup and Configuration Questions

Q. I have a previously configured standalone machine and I join an existing cluster. What happens to my settings?

A. When a machine joins a cluster, all of that machine's clusterable settings will be inherited from the cluster level. Upon joining a cluster, all locally configured non-network settings will be lost, overwritten with the settings of the cluster and any associated groups. (This includes the user/passphrase table; passphrases and users are shared within a cluster).

Q. I have a clustered machine and I remove it (permanently) from the cluster. What happens to my settings?

A. When a machine is permanently removed from a cluster, its configuration hierarchy is “flattened” such that the machine will continue to work the same as it did when it was part of the cluster. All settings that the machine has been inheriting will be applied to the machine in the standalone setting.

For example, if there is only a cluster-mode Global Unsubscribe table, that Global Unsubscribe table data will be copied to the machine's local configuration when the machine is removed from the cluster.

General Questions

Q. Are log files aggregated within centrally managed machines?

A. No. Log files are still retained for each individual machines. The Security Management appliance can be used to aggregate mail logs from multiple machines for the purposes of tracking and reporting.

Q. How does User Access work?

A. The appliance share one database for the entire cluster. In particular, there is only admin account (and passphrase) for the entire cluster.

Q. How should I cluster a data center?

A. Ideally, a data center would be a “group” within a cluster, not its own cluster. However, if the data centers do not share much between themselves, you may have better results with separate clusters for each data center.

Q. What happens if systems are offline and they reconnect?

A. Systems attempt to synchronize upon reconnecting to the cluster.

Network Questions

Q. Is the centralized management feature a “peer-to-peer” architecture or a “primary/secondary” architecture?

A. Because every machine has all of the data for all of the machines (including all machine-specific settings that it will never use), the centralized management feature can be considered a peer-to-peer architecture.

Q. How do I set up a box so it is not a peer? I want a “secondary” system.

A. Creating a true “secondary” machine is not possible with this architecture. However, you can disable the HTTP (GUI) and SSH (CLI) access at the machine level. In this manner, a machine without GUI or CLI

access *only* be configured by clusterconfig commands (that is, it can never be a login host). This is similar to having a secondary machine, but the configuration can be defeated by turning on login access again.

Q. Can I create multiple, segmented clusters?

A. Isolated “islands” of clusters are possible; in fact, there may be situations where creating them may be beneficial, for example, for performance reasons.

Q. I would like to reconfigure the IP address and hostname on one of my clustered appliances . If I do this, will I lose my GUI/CLI session before being able to run the reboot command?

A. Follow these steps:

1. Add the new IP address
2. Move the listener onto the new address
3. Leave the cluster
4. Change the hostname
5. Make sure that oldmachinename does not appear in the clusterconfig connections list when viewed from any machine
6. Make sure that all GUI sessions are logged out
7. Make sure that CCS is not enabled on any interface (check via interfaceconfig or Network > Listeners)
8. Add the machine back into the cluster

Q. Can the Destination Controls function be applied at the cluster level, or is it local machine level only?

A. It may be set at a cluster level; however, the limits are on a per-machine basis. So if you limit to 50 connections, that is the limit set for each machine in the cluster.

Planning and Configuration

Q. What can I do to maximize efficiency and minimize problems when setting up a cluster?

1. Initial Planning

- Try to configure as many things as possible at the cluster level.
- Manage by machines only for the exceptions.
- If you have multiple data centers, for example, use groups to share traits that are neither cluster-wide nor necessarily machine-specific.
- Use the same name for Interfaces and Listeners on each of the appliances .

2. Be aware of restricted commands.

3. Pay attention to interdependencies among settings.

For example, the listenerconfig command (even at the cluster level) depends on interfaces that only exist at a machine level. If the interface does not exist at the machine level on all machines in the cluster, that listener will be disabled.

Note that deleting an interface would also affect listenerconfig .

4. Pay attention to your settings!

Remember that previously-configured machines will lose their independent settings upon joining a cluster. If you want to re-apply some of these previously configured settings at the machine level, be sure to take note of all settings before joining the cluster.

Remember that a “disconnected” machine is still part of the cluster. When it is reconnected, any changes you made while it was offline will be synchronized with the rest of the cluster.

Remember that if you permanently remove a machine from a cluster, it will retain all of the settings it had as part of that cluster. However, if you change your mind and re-join the cluster, the machine will lose all standalone settings.

Use the `saveconfig` command to keep records of settings.