
Cluster Install for Single-Node Configuration

For Red Hat Cluster Suite Integration

Zimbra Collaboration Suite (ZCS) can be integrated with Red Hat® Enterprise Linux® Cluster Suite version 4, update 3 to provide high availability.

In a single-node cluster implementation, all Zimbra servers are part of a cluster under the control of the Red Hat Cluster Manager.

Note: *Red Hat Cluster Suite consists of Red Hat Cluster Manager and Linux Virtual Server Cluster. For ZCS, only Red Hat Cluster Manager is used. In this guide, Red Hat Cluster Suite refers only to Cluster Manager.*

This chapter describes configuring one active node and one standby node in a cluster environment. In the example commands in this guide, both the service name and the domain name are **mail.example.com**.

Pre-configuration Requirements

Both servers must meet the requirements described in the Zimbra Collaboration Suite Quick Start Guide for single server installations, in addition to the requirements described here. For the latest copy of this guide, go to <http://www.zimbra.com/products/documentation.html>.

Go to the Red Hat Cluster Suite website, <https://www.redhat.com/software/rha/cluster> to view specific system requirements for cluster configurations using Red Hat Cluster Suite. If you are not familiar with the Red Hat Cluster Suite, read the documentation to understand how each of the components works to provide high availability.

ZCS License Requirements

A Zimbra license is required in order to create accounts in ZCS. See the Quick Start Guide for more details.

The license file should be saved to a directory on the active node. During the installation you will be asked to enter the file name and location.

Hardware for the Cluster Environment

For Red Hat Cluster Suite integration, the following hardware is required.

- SAN (shared disk storage device) to store the data for each of the Zimbra servers. The size of the shared storage device depends on your expected site capacity.
- Network power control switch to connect cluster nodes. The power control switch is used as the fence device for I/O fencing during a failover. ZCS recommends that you use either a APC or a WTI network power switch. If you do not use one of these power switches, you will need to manually edit the configuration after the installation is complete.

Configure the network power control switch according to the manufacturer's requirements.

Software Requirements For Clustering

- The Red Hat Enterprise Linux 4, Update 3 operating system installed on each server node configured with the same netmask and broadcast address.
- To use the Red Hat Cluster Configuration Tool GUI, install X Window and a desktop environment such as GNOME or KDE.
- Red Hat Cluster Suite, Update 3 or later on each server node.

Preparing the SAN

Note: *You can place all service data on a single volume or choose to place the service data in ten volumes. A more customized volume configuration is possible, but the configurator script only supports single- or ten-volume volume sets. This is a limitation of the configurator script, not of Zimbra Collaboration Suite or of Red Hat Cluster Suite.*

Configure the SAN device and create the partitions for the volumes. Refer to the Red Hat Cluster Suite documentation for configuration requirements.

- If you select to configure the SAN in one volume with subdirectories, all service data goes under a single SAN volume.
- If you select to partition the SAN into 10 volumes, the SAN device is partitioned to provide the following volumes for each Zimbra server in the cluster.

- | | |
|------------------|---|
| • conf | Volume for the service-specific configuration files |
| • log | Volume for the local logs for Zimbra server |
| • redolog | Volume for the redo logs for the Zimbra server |
| • db/data | Volume for the MySQL data files for the data store |
| • store | Volume for the message files |
| • index | Volume for the search index files |
| • backup | Volume for the backup files |

- logger/db/data Volume for the MySQL data files for logger service's MySQL instance
- openldap-data Volume for OpenLDAP data
- postfix/spool Volume for Postfixspool

Installing the Zimbra Cluster Software

The Zimbra Cluster software consists of **install.pl**, **postinstall.pl**, and **configure-cluster.pl** scripts to automate the cluster configuration process and files that are used during the Zimbra cluster service operation.

Installing and configuring a single server for a cluster environment requires that you configure both servers in a specific sequence.

Flow of Installation:

1. On the Active node
 - Run cluster install.pl to install the necessary files, define users and groups, and create the mount points for the clustered service.
 - Install Zimbra Collaboration Suite. All packages are installed
2. On the Standby node
 - Run cluster install.pl to install the necessary files and define users and groups
 - Install Zimbra Collaboration Suite. All packages are installed (software only)
 - Set up syslog and MTA auth
3. On the Active node
 - Set up syslog and MTA auth
4. On the Standby node
 - Run the cluster postinstall.pl program
5. On the Active node
 - Mount the SAN volume(s)
 - Run the cluster postinstall.pl program
 - Run the cluster configurator script, configure-cluster.pl, to prepare the Red Hat Cluster Suite
 - Copy the cluster config. file to the standby node
 - Start Red Hat Cluster Suite daemons
6. On the standby node, start Red Hat Cluster Suite daemons

Installing and Configuring Single-Node Cluster Services

Column one displays the steps performed on the **Active Host**, column two, the steps performed on the **Standby**. The arrow identifies when you must continue the configuration on the other host.

IMPORTANT: These steps must be followed precisely because what you do on one node requires the other node to be in a specific state in order to be correctly configured.

Active

1. Bring up the service IP address on the active node.

```
[root@node1 ~]# ip addr add
xx.xx.xxx.xx dev eth0
```

2. Run the cluster install.pl:
 - **tar xzvf zcs-cluster.tgz**
to unpack the file
 - **cd zcs-cluster**
to change to the correct directory
 - **./install.pl**
to begin the installation

The necessary files are installed.

Each Zimbra cluster node requires Zimbra and Postfix users and groups. The same user and group IDs must be used on both nodes.

- a. Type the zimbra group ID (GID) to be used. The default is 500.
- b. Type the postfix group ID. The default is 501.
- c. Type the postdrop group ID. The default is 502.
- d. Type the zimbra user ID (UID) to be used. The default is 500.
- e. Type the postfix user ID. The default is 501.

The root directory for the mount points is created.

Standby

Active

- f. Mount point(s) are created for the cluster service. Type the service name when prompted.
 - g. Type **Done**, when finished.
3. Install the ZCS software.
All packages should be installed. SNMP is optional. See the Quick Start Installation Guide for detailed installation instruction.
To automatically install the Zimbra license during the install process, type `./install.pl -l /<directory>/ZCSLicense.xml`.
<directory> is the directory were you saved this file.
During the installation make the following changes
When the DNS error to resolve MX displays, enter yes to change the domain name. Modify the domain name to the cluster service hostname (not the active node hostname).


On the **Main Menu** make the following changes

- **Host name and LDAP master host name** must be changed from the active node hostname to the cluster service hostname.
- Note the LDAP password. You will need it later.
- Change the admin password.
- Install the license file. Enter the directory and file name for the Zimbra license.

When the ZCS installation is complete, there should be no reference to the active node hostname.

Standby

Active



Standby

4. On the Standby host run, run the cluster install.pl:

- **tar xzvf zcs-cluster.tgz** to unpack the file
- **cd zcs-cluster** to change to the correct directory
- **./install.pl** to begin the installation

The necessary files are installed.

Each Zimbra cluster node requires Zimbra and Postfix users and groups. The same user and group IDs must be used on both nodes.

- a. Type the zimbra group ID (GID) to be used. The default is 500.
- b. Type the postfix group ID. The default is 501.
- c. Type the postdrop group ID. The default is 502.
- d. Type the zimbra user ID (UID) to be used. The default is 500.
- e. Type the postfix user ID. The default is 501.

The root directory for the mount points is created.

- f. Mount point(s) are created for the cluster. Type the service names when prompted. These are the same service names as on the active host.
- g. Type **Done**, when finished.

Active

Standby

When you install ZCS on the standby node, you must configure the node as described below.

5. Install the ZCS software. Install the same Zimbra packages as installed on the active host. During the installation make the following changes
 - When the DNS error to resolve MX displays, enter yes to change the domain name. Modify the domain name to the cluster service name (not the server node name).
 - The DNS error appears again. This time when the installer asks "Re-Enter domain name?", type **No**.

Make these changes to the following Main Menu sections.

- **LDAP master host name** must be changed to point to the LDAP server on the active node (mail.example.com).
Note: this name is the service name, not the active node name.
- Change the **LDAP password** to the password set on the active node.
- **zimbra-ldap.**
Disable LDAP on the standby node.
- **zimbra-store - Admin user to create:**
Type No. An admin account should not be created on the standby node as it is already created on the active node.

Active

Standby


- **zimbra-store - SMTP host:** If SMTP is configured, change the SMTP host to the cluster service host. (mail.example.com)
- **zimbra-mta - MTA Auth host:** Change the MTA's auth host name to the cluster service host (mail.example.com)
- **zimbra-spell.** Enter the URL.
- **zimbra-logger**
Disable logger on the standby. It is enabled on the active node.

Complete the ZCS installation on the standby node.

In order for remote management and postfix queue management, the ssh keys must be manually populated on each server.

6. To set up syslog and MTA auth keys, as Zimbra user (**su - zimbra**). Type **zmupdateauthkeys** and press **Enter**. The keys are added to **/opt/zimbra/.ssh/authorized_keys**.

7. To set up syslog and MTA auth keys, as Zimbra user (**su - zimbra**). Type **zmupdateauthkeys** and press **Enter**. The key are added to **/opt/zimbra/.ssh/authorized_keys**



Run cluster postinstall.pl . Postinstall must be run on the standby node first because execution of postinstall requires that the LDAP server be running. Zimbra cluster post install script is used after Zimbra Collaboration Suite is installed on the servers to move the data files from the local disk to the volume(s) created on the SAN

Active

9. Mount the SAN volume (s). You can mount one volume for all services or you can mount ten separate volumes. The following command is to mount one volume for all services. To mount by label as root type:
[root@node1 zcs] mount LABEL=mysanvol /opt/zimbra-cluster/mountpoints/mail.example.com.
10. Run cluster postinstall.pl.
To start the Zimbra post install cluster configuration script, **cd** to the zcs-cluster directory created in step 2.
Type **./postinstall.pl** to begin post install.

Standby

8. To start the Zimbra post install cluster configuration script, **cd** to the zcs-cluster directory created in step 2.
Type **./postinstall.pl** to begin post install.

The Zimbra processes are stopped, various cluster-specific adjustments are made to the Zimbra Collaboration Suite installation, and unnecessary data files are deleted



Active

The Zimbra processes are stopped, various cluster-specific adjustments are made to the Zimbra Collaboration Suite installation, and the data files are moved to the SAN volume(s).

When the postinstall is complete use the Zimbra cluster configurator script to prepare Red Hat Cluster Suite to run the Zimbra Collaboration Suite. **The cluster configurator script is run on only the active mailbox node.**

The cluster configurator asks a series of questions to gather information about the cluster and generate the cluster configuration file, `/etc/cluster/cluster.conf`. This is the main configuration file of Red Hat Cluster Suite.

The cluster configurator installs the generated configuration file on each cluster node as `/etc/cluster/cluster.conf`.

11. To start the Zimbra configuration script, `cd` to the `zcs-cluster` directory created in step 2.

Type `./configure-cluster.pl`.

The configurator checks to verify that the server installation is correct.

12. When **Is installation finished on all cluster nodes?** displays, type `y` to continue.

13. Enter a name to identify this cluster. Press **Enter**.

Important: Each cluster on the same network must have a distinct name. Make sure you enter a name that is not in use! Each Red Hat Cluster Suite cluster on the same network must have a distinct name to avoid interfering with another Red Hat Cluster Suite cluster.

Standby

Active

14. Select the network power switch type that is used as the fence device. Configure the fence device host name/IP address, login, and password.

15. Enter the fully-qualified hostname for the nodes in the cluster and the plug number associated with the node's power cord. When the two nodes are identified, type **Done**.

For each service, you need to choose a preferred node to run on and enter the list of volumes to be mounted from the SAN.

16. Select the cluster service. In this cluster configuration, only one service is available. Select **1**.

17. Choose the preferred node on which to run service mail.example.com, node 1.

Standby

Active

18. A Zimbra cluster service must mount service-specific data volumes. All service data can be placed on a single volume or the different types of data can be distributed over multiple volumes. Choose the volume setup type, single volume or multiple volumes.
19. When “Choose a service...”, displays, select **2**. The configuration is complete.
20. Press **Enter** again to view a summary of the configuration.
21. After viewing the summary, save the configuration to a file. You can either accept the default name or rename the configuration file.
22. The configuration file must be copied to the standby node. If you want the script to copy the file to the standby node, enter **Yes**. (Enter the root password, if prompted.)

23. When asked, press **Enter**, to continue.
24. Bring down the cluster service IP address. At **root@node1 ZCS-cluster]#** type **ip addr del xx.xx.xx.xx dev eth0**. You can now proceed with starting the RHCS daemons, which will bring up ZCS on one of the nodes.
25. Start the cluster for the first time on the active node. See “Starting the Red Hat Cluster Suite Daemons” section below.
26. When clustat shows the cluster services running on the active node, the cluster configuration is complete.

Standby

Start the Red Hat Cluster Suite Daemons

After the cluster configuration file is copied, you can start the Red Hat Cluster Suite daemons.

Important: In order to start the cluster daemons correctly, you must be logged on to each node before proceeding, and to see any errors, you should have two sessions open for each node. You enter a command for one node, then enter the same command for the second. You must enter each command on both nodes, before proceeding to the next command.

- Log on to each node as root.
- Run `tail -f /var/log/messages`, on each node to watch for any errors.
- Open another session for each node.

To start the Red Hat Cluster Service on a member, type the following commands in this order. Remember to enter the command on each node before proceeding to the next command.

1. **service ccsd start.** This is the cluster configuration system daemon that synchronizes configuration between cluster nodes.
2. **service cman start.** This is the cluster heartbeat daemon. It returns when both nodes have established heartbeat with one another.
3. **service fenced start.** This is the cluster I/O fencing system that allows cluster nodes to reboot a failed node during failover.
4. **service rgmanager start.** This manages cluster services and resources.

The `service rgmanager start` command returns immediately, but initializing the cluster and bringing up the Zimbra Collaboration Suite application for the cluster services on the active node may take some time.

After all commands have been issued on both nodes, run `clustat` command on the active node, to verify that the cluster service has been started.

When `clustat` shows all services are running on the active node, the cluster configuration is complete.

What to do if cluster services does not relocate to preferred node

If the services does not relocate to the active node after several minutes, you can issue Red Hat Cluster Suite utility commands to manually correct the situation.

Note: Not starting correctly on the preferred node usually is an issue that happens only the first time the cluster is started.

For the cluster service that is not running on the active node, run `clusvcadm -d <cluster service name>`, as root on the active node.

```
[root@node1.example.com]#clusvcadm -d mail1.example.com
```

This disables the service by stopping all associated Zimbra processes, releasing the service IP address, and unmounting the service's SAN volumes.

To enable a disabled service, run `clusvcadm -e <service name> -m <node name>`. This command can be run on any cluster node. It instructs the specified node to mount the SAN volumes of the service, bring up the service IP address, and start the Zimbra processes.

```
[root@node1.example.com] #clusvcadm -e mail1.example.com -m
node1.example.com
```

Testing the Cluster Set up

To perform a quick test to see if failover works:

1. Log in to the remote power switch and turn off the active node.
2. Run `tail -f /var/log/messages` on the standby node. You will observe the cluster becomes aware of the failed node, I/O fence it, and bring up the failed service on the standby node.

View Zimbra Cluster Status

Go to the Zimbra administration console to check the status of the Zimbra cluster. The **Server Status** page shows the cluster server, the node, the services running on the cluster server, and the time the cluster was last checked. The standby node is displayed as standby. If a service is not running, it is shown as disabled.

Zimbra, Inc. Copyright © Zimbra, Inc. 2006-2007. All rights reserved.

The Zimbra logo and logo type are trademarks of Zimbra, Inc. All other trademarks are the property of their respective owners.

January 2007

rev 2