
Zimbra Cluster Installation - Multi-Node Configuration

For Red Hat Cluster Suite Integration

Zimbra Collaboration Suite can be integrated with Red Hat® Enterprise Linux® Cluster Suite version 4, update 3 to provide high availability. In a cluster implementation, all Zimbra mailbox servers are part of a cluster under the control of the Red Hat Cluster Manager.

This guide explains how to install ZCS on Zimbra servers in a cluster environment. Before you install your cluster environment, read the Multi-Server Installation Guide for information about how to prepare your server environment, plan for the install, and determine the system requirements. To get the latest copy of the documentation, go to <http://www.zimbra.com/products/documentation.html>.

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

Pre-configuration Requirements

All servers must meet the requirements described in the Systems Requirements section of the Multi-Server Installation Guide, in addition to the requirements described here.

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 work to provide high availability.

ZCS License Requirement

A Zimbra license is required in order to create accounts in the ZCS. See the Multi-Server Installation Guide for more details.

The license file should be saved to a directory on the first active node to be installed. 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 mailbox 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. Use either a APC or a WTI network power switch.

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 mailbox 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 on each mailbox server node.

Preparing the SAN

Configure the SAN device and create the partitions for the volumes. Refer to the Red Hat Cluster Suite documentation for configuration requirements. The SAN device must be partitioned to provide the following volumes for each Zimbra mailbox 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
- **mysql** 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** Volume for the MySQL data files for logger service's MYSQL instance

Overview of Cluster Installation

Red Hat Cluster Suite integration requires planning the cluster design and precisely executing the configuration. The Zimbra Cluster software automates the setup on the nodes. The scripts in the Zimbra Cluster software configure the Zimbra Collaboration Suite servers for Red Hat Cluster integration. In most

cases, you may not need to use Red Hat's graphical Cluster Configuration Tool to configure the Zimbra cluster. If you do, refer to the Red Hat Cluster Suite documentation for detailed configuration and management instructions.

The Zimbra Cluster software includes:

- Zimbra Cluster install script, used before the Zimbra Collaboration Suite installer to create the mount points for the SAN volumes.
- Zimbra Cluster post install script, used after Zimbra Collaboration Suite is installed on the servers to move the data files from the local disk to the volumes created on the SAN.
- Zimbra Cluster Configurator script that runs on one active node. The configurator script automates the Red Hat Cluster configuration process, taking you through the steps to create the `/etc/cluster/cluster.conf` file. In addition, the configurator script copies the `cluster.conf` file to each node.

Cluster Scenario

The screen-shots in this chapter describe configuring a cluster environment with two active nodes, one standby node, and two cluster services and separate LDAP and MTA servers that are not under the control of Red Hat Cluster Suite. The domain name is **example.com**.

The following Zimbra servers are configured:

- One Zimbra LDAP server, **ldap.example.com**
- One Zimbra MTA server, **mta.example.com** ()
- Three Zimbra mailbox nodes. Two mailbox nodes are active servers. One mailbox node is the standby server.
 - Active mailbox node 1, **node1.example.com**
 - Active mailbox node 2, **node2.example.com**
 - Standby mailbox node, **node3.example.com**
- Two cluster services, one for each of the active nodes
 - Cluster Service 1, **mail1.example.com**
 - Cluster Service 2, **mail2.example.com**

Sixteen volumes are configured on the SAN for this example cluster, eight for each of the two services.

Installing and Configuring the Software

You should install and configure ZCS servers in the following order:

1. Zimbra LDAP server
2. Active and standby mailbox nodes as the Zimbra mailbox servers in the cluster.

3. MTA servers. The MTA server is last because you need to configure one of the active cluster services' hostname as the MTA auth host.

Install the LDAP Server

See the Multi-Server Installation Guide, for instructions about how to install the Zimbra LDAP server.

Install the Active Mailbox Nodes

For each active mailbox node, install and configure the following software:

- Red Hat Cluster Suite software
- Zimbra Cluster software
- Zimbra Collaboration Suite software

Installing the Red Hat Cluster Suite Software

On each node, install the required RPMs and the **rgmanager** RPM for *Red Hat Cluster Suite with DLM*. See the Red Hat Cluster Suite documentation, *Determining RPMs To Install Determining* section for descriptions and the installation instructions.

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.

The software is a standard compressed tar file. Save the file to the computer from which you will install the software.

1. Log in as **root** to the Zimbra server and **cd** to the directory where the Zimbra **zcs-cluster.tgz** file is saved. Type the following commands:
 - **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 scripts, files, and Red Hat Cluster Suite patches are installed.

```
[root@node1 zcs-cluster]# ./install.pl

Each cluster node needs zimbra user and zimbra group. The same user ID
and group ID must be used on all cluster nodes to allow files on SAN
owned by zimbra user/group to be accessible on every node.

Enter zimbra group ID [500]:
.... groupadd -g 500 zimbra

Enter zimbra user ID [500]:
.... useradd -u 500 -g zimbra -G tty -d /opt/zimbra -s /bin/bash zimbra
.... chown root:root /opt/zimbra

Creating root directory for mount points
.... mkdir -p /opt/zimbra-cluster/mountpoints
```

2. Type the Zimbra group ID (GID) to be used. The same group ID number must be configured on every node. The default is 500. Change the default, if this group ID is not available on all the nodes in the cluster.
3. Type the Zimbra user ID (UID) to be used. The same user ID number must be configured on every node. The default is 500. Change the default, if this user ID is not available on all the nodes in the cluster.
4. Type the first cluster service name, press **Enter**. Type **mail1.example.com**. This is the public hostname. The eight volume mount points for the cluster service are created.
5. Type additional cluster service names until all services are configured.
6. Type **Done**, when finished.

On every mailbox server node you need to create mount points for all cluster services. Enter one service name per prompt.

```
Enter cluster service name ("done" to finish): mail1.example.com
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail1.example.com
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail1.example.com/conf
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail1.example.com/log
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail1.example.com/redolog
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail1.example.com/db/data
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail1.example.com/store
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail1.example.com/index
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail1.example.com/backup
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail1.example.com/logger/
db/data
```

```
Enter cluster service name ("done" to finish): mail2.example.com
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail2.example.com
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail2.example.com/conf
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail2.example.com/log
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail2.example.com/redolog
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail2.example.com/db/data
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail2.example.com/store
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail2.example.com/index
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail2.example.com/backup
.... mkdir -p /opt/zimbra-cluster/mountpoints/mail2.example.com/logger/
db/data
```

```
Enter cluster service name ("done" to finish): done
```

```
Mount points were created for the following cluster services:
```

```
mail1.example.com
mail2.example.com
```

Installing the Zimbra Collaboration Suite Software

Important: *If you install the Logger package, it must be installed on each mailbox node but only enabled on the first active node.*

For each active node in the cluster, install the Zimbra Collaboration Suite as follows. For a smooth installation, note these configuration points.

- *When the Zimbra software is installed, the installation detects the hostname configured for the server and automatically inserts this name as the default hostname for various values. The server hostname must be changed to the cluster service name configured in [Step 4](#) in *Installing the Zimbra Cluster Software* section.*
- *The LDAP server name and LDAP password are required. To find the LDAP password, after the LDAP server is installed, on the LDAP server, type `su - zimbra`, then type `zmlocalconfig -s ldap_root_password`.*

1. Log in as **root** to the server and **cd** to the directory where the Zimbra `zcs.tgz` file is saved. Type the following commands.

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

For the first active node only, to install your Zimbra license during the installation process, type **./install.pl -l /<directory>/ZCSLicense.xml**. If you do not install it now, you are asked to install it when you configure the Zimbra store.

The installation process checks to see if Sendmail, Postfix, and MySQL software are running. If any of these are, you are asked to disable them. The default is **Yes** to disable them. Disabling MySQL is optional, but highly recommended.

The `install.sh` script displays a reference to the Zimbra Public License with an address to view the license, and then reviews the installed software to verify that the prerequisite software is installed. If any is missing, the installation stops.

2. When asked to select the packages to install, type **N** for the Zimbra-LDAP, and Zimbra MTA packages. Zimbra Store, Zimbra SNMP, Zimbra Logger and Zimbra Spell should be marked **Y**. Press **Enter**. (Of these packages, only Zimbra Store is required.)

```
Installing:
  zimbra-core
  zimbra-logger
  zimbra-snmp
  zimbra-store
  zimbra-apache
  zimbra-spell
```

The selected packages are installed on the mailbox server.

At this point the **Main menu** displays the default entries for the mailbox server you are installing.

```

Main menu

1) Hostname:                mailhost.example.com
2) Ldap master host        mailhost.example.com
3) Ldap port:              389
4) Ldap password:         set
5) zimbra-ldap:           Enabled
   +Create Domain:         yes
   +Domain to create:      mailhost.example.com
6) zimbra-store:          Enabled
   +Create Admin User:     yes
   +Admin user to create:  admin@mailhost.example.com
***** +Admin Password    UNSET
   +Enable automated spam training: yes
   +Spam training user:    fdi0j@mailhost.example.com
   +Non-spam(Ham) training user: s3nnl@mailhost.example.com
   +Global Documents Account wiki@mailhost.example.com
   +SMTP host:             mailhost.example.com
   +Web server HTTP port:  80
   +Web server HTTPS port: 443
   +Web server mode:       http
   +Enable POP/IMAP proxy: no
   +IMAP server port:      143
   +IMAP server SSL port:  993
   +POP server port:       110
   +POP server SSL port:   995
   +Use spell check server: yes
   +Spell server URL:      http://
mailhost.example.com:7780/aspell.php
   ****+License file name: UNSET
7) zimbra-mta:             Enabled
8) zimbra-snmp:            Enabled
9) zimbra-logger:         Enabled
10) zimbra-spell:          Enabled
11) Enable default backup schedule: yes
r) Start servers after configuration yes
s) Save config to file
x) Expand menu
q) Quit

Address unconfigured (**) items or correct ldap configuration (? - help)

Checking ldap on :389...FAILED

```

3. Change the **Hostname** to one of the cluster service names entered in [Step 4](#) in *Installing the Zimbra Cluster Software* section (In our example, this is **mail1.example.com**). Type **1**, and then type the cluster service name, press **Enter**.
4. Set the LDAP master host and password.
 - Type **2** and then type the LDAP host name.
 - Type **4** and then type the LDAP password.

As you enter each of these values the server tries to contact the LDAP server. You can proceed when the LDAP server is successfully contacted.

5. Type **6** to configure the admin password, the SMTP host, and to set the web server mode, if your configuration is not http.

- On the first active node only, type **4** and set the password for the administrator account. The password is case sensitive and must be a minimum of six characters. The admin account is provisioned on the Zimbra server and allows you to log on to the administration console. The administrator name, mail address, and password are required to log in to the administration console.

On the other active nodes, type **2**, Create Admin User and type No. You only need to create one admin account.

- Type **9** to set the SMTP host.
- Type **10**, if you are changing the default. The communication protocol options are HTTP, HTTPS, or mixed. Mixed mode uses HTTPS for logging in and HTTP for normal session traffic. All modes use SSL encryption for back-end administrative traffic
- If you are setting up IMAP/POP proxy servers, type **13** to enable. When you enable these, IMAP and POP server port numbers and proxy port numbers are automatically changed. See the "Planning for the Installation" chapter, Configuring IMAP and POP Proxy Server.
- If **20, License filename** is displayed, your Zimbra license is not installed. Type **20** and enter the directory and file name for the Zimbra license. For example, if you saved the license file to the tmp directory, you would type **/tmp/ZCSLicense.xml**. If you do not have the license, you can proceed but will not be able to create new accounts until the license is installed.

6. Type **8** to configure the SMTP host and set the web server mode, if it is not http.

- Type **2** and then type the Zimbra MTA host name. The MTA does not need to be installed, but the name you enter must be the correct name for the MTA.
- Type **3**, if you are changing the default mode. The communication protocol options are HTTP, HTTPS, or mixed. Mixed mode uses HTTPS for logging in and HTTP for normal session traffic. All modes use SSL encryption for back-end administrative traffic.

Important: For clustering, the Web mode must be identical on all nodes.

```

Store configuration

1) Status:                               Enabled
2) SMTP host:                             smtp.example.com
3) Web server mode:                       https
4) IMAP server port:                      143
5) IMAP server SSL port:                 993
6) POP server port:                      110
7) POP server SSL port:                 995
8) Use spell check server:              yes
9) Spell server URL:                    http://
mail1.example.com:7780/aspell.php

Select, or 'r' for previous menu [r]
Checking ldap on ldap.example.com:389...Success

```

7. If you installed the SNMP package, you will need to modify the default notification addresses. Type 6 to modify the SNMP packages.

Configure whether to be notified by SNMP or SMTP. The default is **No**. If you enter yes, you must enter additional information.

- For SNMP, enter the SNMP Trap host name.
- For SMTP, enter the SMTP source email address and destination email address. Type the same host address as configured in the LDAP server.

```

Snmpp configuration

1) Status:                               Enabled
2) Enable SNMP notifications:            yes
3) SNMP Trap hostname:                  snmp.example.com
4) Enable SMTP notifications:           yes
5) SMTP Source email address:          admin@example.com
6) SMTP Destination email address:     admin@example.com

Select, or 'r' for previous menu [r]

```

8. When Logger is installed, it must be enabled on **the first** node. All other nodes must install but disable Logger. To disable logger, type the menu number for Logger and press **Enter**.
9. If you have no other changes, type **a** to apply the configuration changes. Press **Enter**, after **Save configuration data?** displays.
10. When **The system will be modified - continue?** appears, type **Y** and press **Enter**.
11. After the **Operations logged to /tmp/zmsetup.log.xxx**, press **Enter**. The server is modified. Installing all the components and configuring the server can take a few minutes.
12. When **Installation complete - press return to exit** displays, press **Enter**.

Mounting Volumes for Cluster Service

- Mount the eight-volume set for a cluster service entered in Step 4 of the Installing the Zimbra Cluster Software. The volumes must be mounted before proceeding.

Important: Verify that the mounted volumes are empty before proceeding.

Running Zimbra Cluster Post Install Script

1. To start the Zimbra post install cluster configuration script, `cd` to the `zcs-cluster` directory created in the Installing the Zimbra Cluster Software section. Type `./postinstall.pl` to begin post install.

```
[root@node1 zcs-cluster]# ./postinstall.pl
Disabling boot-time auto start of Zimbra applications.
.... chkconfig --del zimbra

Enabling remote syslogging
.
.
Installing RPMS
.
Installing RHCS patches
.
.
.
Modifying /etc/sudoers...

Backing up existing cluster.conf
.... mv /etc/cluster/cluster.conf /etc/cluster/cluster.conf.bak

Checking node type (active vs. standby)...
Enabling cluster administration in Admin Console

Setting cluster root in localconfig
.... su - zimbra -c 'zmlocalconfig -e zimbra_cluster_root=/opt/
zimbra-cluster'

Detecting service name installed on this node...
service name = mail1.example.com

Service-specific data files must be moved to SAN volumes. All SAN
volumes
for mail1.example.com service must now be mounted using mount points
created in /opt/zimbra-cluster/mountpoints/mail1.example.com
directory.

Are the SAN volumes mounted for mail1.example.com service? (Y/N) y
```

2. Type **Y** to confirm that the SAN volumes are mounted for the selected service.

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 service-specific volumes.

```
Stopping Zimbra processes...
.... su - zimbra -c 'zmcontrol stop'
Host maill.example.com
    Stopping antispam...Done
    Stopping antivirus...Done
    Stopping logger...Done
    Stopping mailbox...Done
    Stopping snmp...Done
    Stopping spell...Done

Moving data files to /opt/zimbra-cluster/mountpoints/maill.example.com
.... mv -f /opt/zimbra/conf/* /opt/zimbra-cluster/mountpoints/
maill.example.com/conf
.... rmdir /opt/zimbra/conf
.... chown zimbra:zimbra /opt/zimbra-cluster/mountpoints/
maill.example.com/conf
.... mv -f /opt/zimbra/log/* /opt/zimbra-cluster/mountpoints/
maill.example.com/log
.... rmdir /opt/zimbra/log
.... chown zimbra:zimbra /opt/zimbra-cluster/mountpoints/
maill.example.com/log
.... mv -f /opt/zimbra/redolog/* /opt/zimbra-cluster/mountpoints/
maill.example.com/redolog
.... rmdir /opt/zimbra/redolog
.... chown zimbra:zimbra /opt/zimbra-cluster/mountpoints/
maill.example.com/redolog
.... mv -f /opt/zimbra/db/data/* /opt/zimbra-cluster/mountpoints/
maill.example.com/db/data
.... rmdir /opt/zimbra/db/data
.... chown zimbra:zimbra /opt/zimbra-cluster/mountpoints/
maill.example.com/db/data
.... mv -f /opt/zimbra/store/* /opt/zimbra-cluster/mountpoints/
maill.example.com/store
mv: cannot stat `/opt/zimbra/store/*': No such file or directory
.... rmdir /opt/zimbra/store
.... chown zimbra:zimbra /opt/zimbra-cluster/mountpoints/
maill.example.com/store... mv -f /opt/zimbra/index/* /opt/zimbra-
cluster/mountpoints/maill.example.com/index
mv: cannot stat `/opt/zimbra/index/*': No such file or directory
.... rmdir /opt/zimbra/index
.... chown zimbra:zimbra /opt/zimbra-cluster/mountpoints/
maill.example.com/index... mv -f /opt/zimbra/backup/* /opt/zimbra-
cluster/mountpoints/maill.example.com/backup
mv: cannot stat `/opt/zimbra/backup/*': No such file or directory
.... rmdir /opt/zimbra/backup
.... chown zimbra:zimbra /opt/zimbra-cluster/mountpoints/
maill.example.com/backup
.... mv -f /opt/zimbra/logger/db/data/* /opt/zimbra-cluster/
mountpoints/maill.example.com/logger/db/data
.... rmdir /opt/zimbra/logger/db/data
.... chown zimbra:zimbra /opt/zimbra-cluster/mountpoints/
.
.
.
```

3. When the data has been moved, the eight volumes are unmounted. Press **Enter**, when asked.

```
About to unmount volumes for mail1.example.com service
Volumes to be unmounted:
 /opt/zimbra-cluster/mountpoints/mail1.example.com/conf
 /opt/zimbra-cluster/mountpoints/mail1.example.com/log
 /opt/zimbra-cluster/mountpoints/mail1.example.com/redolog
 /opt/zimbra-cluster/mountpoints/mail1.example.com/db/data
 /opt/zimbra-cluster/mountpoints/mail1.example.com/store
 /opt/zimbra-cluster/mountpoints/mail1.example.com/index
 /opt/zimbra-cluster/mountpoints/mail1.example.com/backup
 /opt/zimbra-cluster/mountpoints/mail1.example.com/logger/db/data
Press Enter to unmount volumes.
... umount /opt/zimbra-cluster/mountpoints/mail1.example.com/conf
... umount /opt/zimbra-cluster/mountpoints/mail1.example.com/log
... umount /opt/zimbra-cluster/mountpoints/mail1.example.com/redolog
... umount /opt/zimbra-cluster/mountpoints/mail1.example.com/db/data
... umount /opt/zimbra-cluster/mountpoints/mail1.example.com/store
... umount /opt/zimbra-cluster/mountpoints/mail1.example.com/index
... umount /opt/zimbra-cluster/mountpoints/mail1.example.com/backup
... umount /opt/zimbra-cluster/mountpoints/mail1.example.com/logger/
db/data

Done.
```

Repeat these steps for every active node in the cluster.

Configuring the Standby Mailbox Server Node

For the standby mailbox server node, install and configure the following software:

- Red Hat Cluster Suite software
- Zimbra Cluster software
- Zimbra Collaboration Suite software

Installing the Red Hat Cluster Suite Software

Install the required RPMs and the **rgmanager** RPM for *Red Hat Cluster Suite with DLM*. See the Red Hat Cluster Suite documentation, *Determining RPMs To Install Determining* section for descriptions and the installation instructions.

Installing the Zimbra Cluster Software

The Zimbra Cluster software is installed and run on each standby node. The software automates the cluster configuration process. The software is a standard compressed tar file. Save the file to the computer from which you will install the software.

The stand-by node is configured exactly the same as the active nodes. You define the same group ID and user ID and identify the cluster service names.

1. Log in as **root** to the Zimbra mailbox server and go to the directory where the Zimbra **zcs-cluster.tgz** file is saved. Untar and type **./install.pl** to begin.
2. Type the Zimbra group ID (GID) to be used. The same group ID number must be configured on every node. The default is 500. Change the default, if you changed it for the active nodes.
3. Type the Zimbra user ID (UID) to be used. The same user ID number must be configured on every node. The default is 500. Change the default, if you changed it for the active nodes.
4. Type the first cluster service name, press **Enter**. Type as **mail1.example.com**. Mount points are created.
5. Continue to add each standby cluster service. The same cluster service names must be entered as on the active nodes
6. Type **Done**, when finished.

Installing the Zimbra Collaboration Suite Software on the Standby Node

Install the Zimbra Collaboration Suite on the standby node. For detailed description of the installation process, review the Zimbra Collaboration Suite Multi-Server Installation Guide.

Important: For a smooth installation, note these configuration points.

- When the Zimbra software is installed, the installation detects the hostname configured for the server and automatically inserts this name as the default hostname for various values. **For the standby node, do not change from the default.**
 - The LDAP server name and LDAP password are required. To find the LDAP password, after the LDAP server is installed, on the LDAP server, type **su - zimbra**, then type **zmlocalconfig -s ldap_root_password**.
1. Log in as **root** to the Zimbra server and **cd** to the directory where the Zimbra **zcs.tgz** file is saved. Type the following commands.

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

The installation process checks to see if Sendmail, Postfix, and MySQL software are running. If any of these are, you are asked to disable them. The default is **Yes** to disable them. Disabling MySQL is optional, but highly recommended.

The **install.sh** script displays a reference to the Zimbra Public License with an address to view the license, and then reviews the installed software to verify that the prerequisite software is installed. If any is missing, the installation stops.

2. When asked to select the packages to install, install the same packages you installed on the active nodes. In our example, type **N** for the Zimbra-LDAP, and Zimbra MTA packages. Zimbra store, Zimbra SNMP, Zimbra Logger and Zimbra Spell should be marked **Y**. Press **Enter**. (Logger, Spell, and SNMP packages are optional, but if installed on the active nodes, must be installed on the standby node.)

```
Installing:
  zimbra-core
  zimbra-logger
  zimbra-snm
  zimbra-store
  zimbra-apache
  zimbra-spell
```

3. The selected packages are installed on the mailbox server. At this point, the **Main menu** displays the default entries for the mailbox server you are installing.
4. Set the LDAP host and password.
 - Type **2**, and then type the LDAP host name.
 - Type **4**, and then type the LDAP password.

As you enter each of these values, the server tries to contact the LDAP server. You can proceed when the LDAP server is successfully contacted.

5. Modify zimbra-store. Type **5** to configure the SMTP host and set the web server mode, if it is not http.
 - Type **2**, for SMTP host, and then type the Zimbra MTA host name.
 - Type **3**, if you are changing the default mode. The communication protocol options are HTTP, HTTPS, or mixed. Mixed mode uses HTTPS for logging in and HTTP for normal session traffic. All modes use SSL encryption for back-end administrative traffic.

Important: For clustering, the Web mode must be identical on all nodes.

```

Store configuration

1) Status:                               Enabled
2) SMTP host:                             smtp.example.com
3) Web server mode:                       https
4) IMAP server port:                       143
5) IMAP server SSL port:                   993
6) POP server port:                        110
7) POP server SSL port:                     995
8) Use spell check server:                 yes
9) Spell server URL:                       http://
mail1.example.com:7780/aspell.php

Select, or 'r' for previous menu [r]
Checking ldap on ldap.example.com:389...Success
    
```

6. If you installed the SNMP package, you will need to modify the default notification addresses. Type 6 to modify the SNMP packages.

Configure whether to be notified by SNMP or SMTP. The default is **No**. If you enter yes, you must enter additional information.

- For SNMP, enter the SNMP Trap host name.
- For SMTP, enter the SMTP source email address and destination email address. Type the same address as configured in the LDAP server.

```

Snmpp configuration

1) Status:                               Enabled
2) Enable SNMP notifications:              yes
3) SNMP Trap hostname:                     snmptrap.com
4) Enable SMTP notifications:              yes
5) SMTP Source email address:               admin@example.com
6) SMTP Destination email address:         admin@example.com
    
```

7. If Logger is installed, it must be disabled on all standby nodes. To disable logger, type the menu number for logger and press **Enter**.
8. When you have no other changes, type **a** to apply the configuration changes. Press **Enter** after **Save configuration data?** displays.
9. When **The system will be modified - continue?** appears, type **Y** and press **Enter**.
10. After the **Operations logged to /tmp/zmsetup.log.xxx**, press **Enter**. The server is modified. Installing all the components and configuring the server can take a few minutes.
11. When **Installation complete - press return to exit** displays, press **Enter**.

Running the Cluster Post Install Script

Now you prepare this server to be the standby server in the cluster. Start the Zimbra cluster post install script.

Note: *Unlike installation of active nodes, no SAN volumes are mounted on standby nodes prior to running the post install script.*

1. Log in as **root** to the Zimbra server and **cd** to the directory where the Zimbra **zcs-cluster.tgz** file is saved. Type the following commands:
 - **cd zcs-cluster** to change to the correct directory
 - **./postinstall.pl** to begin the post install. The Zimbra processes are stopped, various cluster-specific adjustments are made to the Zimbra Collaborating Suite installation, and unnecessary data files are deleted.

```
[root@node3 zcs-cluster]# ./postinstall.pl
Disabling boot-time auto start of Zimbra applications.
.... chkconfig --del zimbra

Creating cluster configuration directory
.... mkdir -p /etc/cluster

Checking node type (active vs. standby)...
This is a standby node.

.... su - zimbra -c 'zmtlctl https 2> /dev/null'

Deleting data files
.... rm -rf /opt/zimbra/tomcat/conf
.... ln -s /opt/zimbra/conf/tomcat /opt/zimbra/tomcat/conf
.... rm -rf /opt/zimbra/tomcat/logs
.... ln -s /opt/zimbra/log/tomcat /opt/zimbra/tomcat/logs

Done.
```

Install the MTA Server

See the Multi-Server Installation Guide, for instructions about how to install the Zimbra MTA server.

Modify Zimbra LDAP and MTA Servers for Logger Service

You must modify the syslog setup on the Zimbra LDAP server and Zimbra MTA servers.

1. On the LDAP server, as root, run **/opt/zimbra/bin/zmsyslogsetup**.
2. On the MTA server, as root, run **/opt/zimbra/bin/zmsyslogsetup**.

Configuring Red Hat Cluster for ZCS

When all the software is installed and the Zimbra installation on the servers configured, 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 one of the active mailbox nodes.**

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`.

Note: *The Zimbra cluster configurator should generate correct configuration file for most installations, but some cases are more complicated. For instance if you are using multiple fence devices or highly customized SAN setup, the configurator script will not work. In those cases, use the configurator to generate an initial cluster.conf. Then run the graphical Red Hat Cluster Configuration Tool, to make the necessary changes. Using the Zimbra Cluster configurator script first is recommended, because the script automates the steps for the basic configuration. After using the Red Hat Cluster Configuration Tool, you must manually copy the final cluster.conf file to each cluster host.*

The Zimbra configurator script guides you through creating the cluster configuration file. The following is configured:

- Fence Device - This is the network power switch. Each mailbox node in the cluster is plugged into the fence device. The cluster uses the fence device for I/O fencing during a failover.
- Cluster Nodes - This section is used to add members to the cluster and configure a fence device setting for each member.
- Managed Resources - The preferred node for each service and the list of volumes to be mounted from the SAN are configured

To use the configurator script

1. To start the Zimbra configuration script, `cd` to the `zcs-cluster` directory created in the Installing the Zimbra Cluster Software section. Type `./configure-cluster.pl`. The configurator checks to verify that the server installation is correct.
2. All servers in the cluster must be installed before you can proceed. When **Is installation finished on all cluster nodes?** displays, type `y` to continue.
3. Enter a name to identify this cluster. Press **Enter**. Each cluster on the same network must have a distinct name.

Important: 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.

```
[root@node1 zcs-cluster]# ./configure-cluster.pl

Zimbra Collaboration Suite Cluster Configurator

This script will guide you through creating an initial configuration
file for Red Hat Cluster Suite.  A series of questions will be asked
to collect the necessary information.  At the end, the configuration
data will be saved to a file and the file will be copied to all
cluster nodes, as /etc/cluster/cluster.conf on each node.

Press Enter to continue.
-----

Checking for Zimbra home... Found.
Checking for Zimbra cluster root... Found.
Checking for cluster mount points root... Found.
Checking for Red Hat Cluster Suite RPMs...
ccs-1.0.2-0
cman-1.0.2-0
dlm-1.0.0-5
fence-1.32.6-0
rgmanager-1.9.39-0
system-config-cluster-1.0.16-1.0

Installation looks good on this node.

You must finish installation on all cluster nodes before configuring
the cluster.  Is installation finished on all cluster nodes? (Y/N) y
-----

Each Zimbra cluster on the network must have a unique name.
Enter the cluster name: zimbra-cluster
```

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

```
A fence device is needed by the cluster for I/O fencing during a failover. The power cord of each cluster node must be plugged into an APC or WTI network power switch device, and the cluster will control the power switch to reboot the node being fenced. While Red Hat Cluster Suite supports a variety of fence devices, for the purpose of this configuration process assume you are using APC or WTI, and also assume all nodes are plugged into a single device. If you are using a different fence device or more than one device, you can correct the generated configuration file later with the system-config-cluster GUI tool.
```

```
Choose device vendor:
```

- 1) APC
- 2) WTI

```
Choose from above (1-2): 1
```

```
Enter fence device hostname/IP address: apc.example.com
```

```
Enter fence device login [apc]:
```

```
Enter fence device password: <password>
```

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

```
For each cluster node you must provide its fully-qualified hostname and the plug number on the fence device.
```

```
Enter node hostname ("done" if no more): node1.example.com
```

```
Enter fence device plug number for node1.example.com: 1
```

```
Enter node hostname ("done" if no more): node2.example.com
```

```
Enter fence device plug number for node2.example.com: 2
```

```
Enter node hostname ("done" if no more): node3.example.com
```

```
Enter fence device plug number for node3.example.com: 3
```

```
Enter node hostname ("done" if no more): done
```

6. Next, you select the cluster service, the preferred node for that service, and the volume set-up to be mounted from the SAN.

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

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.

Choose a service:

- 1) mail1.example.com
- 2) mail2.example.com
- 3) Done

Choose from above (1-3): **1**

Choose preferred node on which to run service mail1.example.com:

- 1) node1.example.com
- 2) node2.example.com
- 3) node3.example.com

Choose from above (1-3): **1**

A Zimbra cluster service must mount service-specific data volumes. Two choices are provided in this configuration process. All service data can be placed on a single volume, or multiple volumes can be used for different types of data files. In the multiple-volumes case eight volumes are used per service.

Choose volume setup type:

- 1) single volume
- 2) multiple volumes

Choose from above (1-2): **2**

7. A prompt is displayed for each volume in the service's volume set. Enter the SAN volume device name for the mount point in the prompt. (These names are the volumes defined when you created the 8 volumes on the SAN as described in ["Preparing the SAN" on page 2.](#))

```
1. Configuration directory:
Volume maill.example.com-conf:
  mount point = /opt/zimbra-cluster/mountpoints/maill.example.com/
  conf
  Enter device name: /dev/sdb5
2. Log directory:
Volume maill.example.com-log:
  mount point = /opt/zimbra-cluster/mountpoints/maill.example.com/log
  Enter device name: /dev/sdb6
3. Redolog directory:
Volume maill.example.com-redolog:
  mount point = /opt/zimbra-cluster/mountpoints/maill.example.com/
  redolog
  Enter device name: /dev/sdb7
4. MySQL data directory:
Volume maill.example.com-db-data:
  mount point = /opt/zimbra-cluster/mountpoints/maill.example.com/db/
  data
  Enter device name: /dev/sdb8
5. Message store directory:
Volume maill.example.com-store:
  mount point = /opt/zimbra-cluster/mountpoints/maill.example.com/
  store
  Enter device name: /dev/sdb9
6. Search index directory:
Volume maill.example.com-index:
  mount point = /opt/zimbra-cluster/mountpoints/maill.example.com/
  index
  Enter device name: /dev/sdb10
7. Backup directory:
Volume maill.example.com-backup:
  mount point = /opt/zimbra-cluster/mountpoints/maill.example.com/
  backup
  Enter device name: /dev/sdb11
8. Logger MySQL data directory:
Volume maill.example.com-logger-db-data:
  mount point = /opt/zimbra-cluster/mountpoints/maill.example.com/
  logger/db/data
  Enter device name: /dev/sdb12
```

8. Continue to configure the preferred nodes and the volume sets for the remaining cluster services.
9. When finished choosing the services, select **Done**. Press **Enter**, and then press **Enter** again to view a summary of the configuration.

```
Finished collecting information.
Press Enter to view summary of the configuration.

Configuration Summary
-----

Cluster Name: example-cluster

Fence Device:
  name:  fence-device
  agent: fence_apc
  ipaddr: apc.example.com
  login:  apc
  passwd: <password>

Nodes:
  node1.example.com - fence port 1
  node2.example.com - fence port 2
  node3.example.com - fence port 3

Services:
.
.mail1.example.com
  ipaddr: 000.000.000.000
  preferred node: node1.example.com
  volumes:
    mail1.example.com-conf
    mountpoint: /opt/zimbra-cluster/mountpoints/
.
.
.
-----

About to save configuration file.
Enter filename [/tmp/cluster.conf.19003]:
-----
```

10. After viewing the summary, save the configuration to a file. You can either accept the default or rename the configuration file.

Note: If you made a mistake, press **Ctrl-C** to abort the configurator script and start over.

Copying the files to all cluster nodes

The configuration file must now be copied to all cluster nodes. The Zimbra configurator script can copy the files, or you can do it manually. This is a continuation of the configurator script.

11. The script offers to do the copy via scp. To automatically copy the **cluster.conf** file to all nodes, type **y**. Enter the root password of each node when asked.

```
Cluster configuration saved in /tmp/cluster.conf.17815
This file must be copied to all cluster nodes now. This script can
do it for you using scp, or you can do it manually.
Copy to all cluster nodes using scp? (Y/N) y

Copying /tmp/cluster.conf.17815 to node1.example.com:/etc/cluster/
cluster.conf .... scp /tmp/cluster.conf.17815
root@node1.example.com:/etc/cluster/cluster.conf
root@node1.example.com's password:
cluster.conf.17815                               100% 5439      5.3KB/s
00:00

Copying /tmp/cluster.conf.17815 to node2.example.com:/etc/cluster/
cluster.conf.... scp /tmp/cluster.conf.17815 root@node2.example.com:/
etc/cluster/cluster.conf
root@node2.example.com's password:
cluster.conf.17815                               100% 5439      5.3KB/s
00:00

Copying /tmp/cluster.conf.17815 to node3.example.com:/etc/cluster/
cluster.conf.... scp /tmp/cluster.conf.17815 root@node3.example.com:/
etc/cluster/cluster.conf
root@node3.example.com's password:
cluster.conf.17815                               100% 5439      5.3KB/s
00:00

Configuration generated and pushed to all cluster nodes.

If necessary, use system-config-cluster GUI tool to further customize
the cluster configuration. You must manually copy the updated
cluster.conf to all nodes.

Press Enter to continue.
```

Important: Use the Red Hat Cluster Configuration Tool if you want to further customize the cluster configuration after the configuration file is generated and copied to all cluster nodes. If you customize the configuration file, you must then manually copy the updated cluster.conf to all nodes.

Start the Red Hat Cluster Suite Daemons

After the cluster configuration file is copied to every node, 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. In our example, you would have six screens opened. You enter a command for one node, then enter the same command for the second, and so forth. You must enter each command on all 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 all nodes 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. The command may not complete on all nodes immediately. It returns when all 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 defined cluster services may take some time.

After all commands have been issued on all nodes, run `clustat` command on one node, to verify all cluster services have been started.

Continue to enter the `clustat` command, until it reports all nodes have joined the cluster, and all services have been started.

Because nodes may not join the cluster in sequence, some of the services may start on nodes that are different from the configured preferred nodes. This is expected and eventually will be restarted on the configured preferred node.

When `clustat` shows all services are running on the preferred nodes, 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 preferred nodes after several minutes, you can issue Red Hat Cluster Suite utility commands to manually correct the situation.

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

For each cluster service that is not running on the correct preferred node, run `clusvcadm -d <cluster service name>`, as root on one of the cluster nodes.

```
[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 an active mailbox node.
2. To watch the standby node take over the failed service, run `clustat`, on one of the other nodes.
3. Run `tail -f /var/log/messages`. You will observe the cluster becomes aware of the failed node, I/O fence it, and bring up the failed service on a standby node.

View Zimbra Cluster Status

Go 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 nodes are displayed as standby. If a service is not running, it is shown as disabled. Managing and maintaining the Zimbra Cluster is through the Red Hat Cluster Manager.

Zimbra, Inc. Copyright © Zimbra, Inc. 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

