
Upgrading Zimbra Collaboration Suite in a Multi-Node Cluster Environment

In a cluster implementation, all Zimbra mailbox servers are part of a cluster under the control of the Red Hat Cluster Manager. The cluster upgrade commands automatically stop all nodes within the cluster, unmount the SAN volumes for each node, run the ZCS upgrade scripts, and restart the cluster.

A multi-node cluster environment includes two or more active nodes, at least one standby node, and separate LDAP and MTA servers that are not under the control of Red Hat Cluster Manager.

The cluster upgrade does not upgrade the LDAP server or the MTA servers. This document explains how to upgrade the Zimbra mailbox servers that are in a cluster environment. To upgrade the LDAP server and the MTA servers, see the latest ZCS release notes for instructions.

BACKUP BEFORE YOU BEGIN TO UPGRADE!

Before you begin the upgrade, run a full backup of the mailbox servers. If there are any problems during the upgrade, you can restore to your previous version. After the upgrade, run a full backup immediately. See the Backup and Restore chapter in the Administration Guide.

Zimbra License

As of ZCS 4.0, the Network Edition requires a Zimbra license to create accounts.

IMPORTANT: If you are updating your cluster from 3.x to 4.0 or later, you must install the ZCS license on every active node before you start the upgrade. The ZCSLicense.xml file will be copied to the corresponding active node's `/opt/zimbra/conf` directory on the SAN. If the license is not installed, the upgrade will fail.

Contact Zimbra sales to purchase a regular license. For more information about Zimbra license, see the ZCS Administrator's Guide.

Upgrading Multi-Node Cluster Environment

Save the ZCS upgrade files, **zcs<version>.tgz**, and the cluster upgrade files, **zcs-cluster.tgz**, together in the same directory on each of the active and standby nodes in the cluster.

The ZCSLicense.xml file should be saved in each active node's **/opt/zimbra/conf** directory.

Save the ZCS upgrade file, **zcs<version>.tgz** to a directory on the LDAP server and MTA server.

You must upgrade ZCS servers in the following order:

1. Zimbra LDAP server
2. MTA servers
3. Active mailbox nodes
4. Standby nodes

Unpacking the Upgrade Packages

For each Zimbra server being upgraded, log on as **root** and **cd** to the directory where the upgrade packages are saved.

Type the following commands:

- **tar xzvf zcs<version>.tgz** to unpack the Zimbra Collaboration Suite files.
- **tar xzvf zcs-cluster.tgz** to unpack the cluster upgrade files in the same directory as where the ZCS files have been unpacked. The cluster upgrade scripts are in the **zcs-cluster/upgrade** directory.

Stopping Cluster Nodes and Preparing LDAP and MTA Servers

1. After the files are unpacked, type **cd zcs-cluster/upgrade** to change to the cluster upgrade directory.
2. On one of the active nodes, run **./stop-all.pl** to stop the services on all nodes. You are asked if it is okay to proceed. Enter **Y**.

The upgrade disables the services on each host and removes the rgmanager cluster daemon from the boot sequence. This prevents rgmanager from automatically starting, if the server is accidentally rebooted before the upgrade is finished.

Note: *The commands are executed with ssh and you may be prompted for a root password for each host.*

3. You are given another opportunity to stop. If you are ready to continue, enter **Y** to continue the upgrade. The rgmanager is stopped on all nodes.

4. If any errors are displayed, log in to each server and make sure `rgmanager` is stopped and is removed from the boot sequence. See “To manually stop a node” on page 5. You should also verify that no cluster SAN volume is mounted on any cluster node. If any volume is mounted, proceeding with the upgrade can cause data corruption. Log in to each node to verify the status of the SAN.

Note: *Instead of logging in to each node to verify the status of the SAN, you can let the upgrade script run `df -h` on each node to make sure the SAN volume is not mounted. When you see **Do You want to see df output?**, enter **Y**.*

5. When asked if you are ready to continue, enter **Y**.
6. Once the SAN volumes are unmounted, you can proceed. When you see **Are all SAN volumes unmounted?**, enter **Y**.
7. Now upgrade the LDAP server before proceeding with the cluster server upgrade. Perform the upgrade according to the directions in the latest ZCS release notes.

Note: *After you upgrade the LDAP server, run `zmcontrol status` to verify that the LDAP server is running before your proceed with cluster upgrade.*

8. Next, upgrade the MTA servers according to the directions in the latest ZCS release notes.

Upgrading Active Nodes

You are now ready to upgrade the active nodes.

If you are upgrading a large number of accounts, the upgrade process can take some time. Use the `nohup` command when running `upgrade-node.pl`, so the remote login session is not disconnected before the ZCS upgrade is completed. The upgrade script requires user input and therefore should be run with the `tee` command.

```
$ nohup ./upgrade-node.pl 2>&1 | tee /tmp/upgrade-node.out
```

There is little interaction required to upgrade the ZCS. The ZCS software and Zimbra service data are upgraded on the active nodes.

1. Enter `./upgrade-node.pl`, to start the ZCS upgrade scripts.
2. After the ZCS `./install.sh` is found, you are asked if this the right location? Enter **Y**.
3. You are asked again to verify that all Zimbra processes and `rgmanager` cluster daemons are stopped on all cluster nodes. If they are stopped, enter **Y**. The Zimbra-cluster package can be upgraded.

If they are not stopped, you must manually stop the node. See “To manually stop a node” on page 5.

- Both software and Zimbra service data are upgraded on each active node. When **Choose a Zimbra service to upgrade** displays, enter the number for one of the active nodes to be upgraded.

```
Choose a Zimbra service to upgrade:
 1) mail1.domain.com
 2) mail2.domain.com
 3) Stand-by node
Choose from above (1-3): 1
You chose mail1.domain.com. Is this correct? (Y/N) y
```

Enter **Y** to verify your selection. The ZCS `./install.sh` script begins.

Note: *You can upgrade multiple nodes in parallel.*

- When **Finished upgrading this node** displays, the active node has been upgraded.
- When all active nodes have been upgraded, proceed to upgrade the standby nodes.

Upgrading Standby Nodes

The standby nodes were prepared for the upgrade when you completed the steps described in the Stopping the Cluster Nodes section. Now you are going to upgrade the cluster software and the ZCS software on the standby nodes.

Running `./upgrade-node.pl`, runs the ZCS upgrade scripts. You should have unpacked the ZCS upgrade files into the same directory as the cluster upgrade files. If this is done, there is little interaction required to upgrade ZCS.

- Run `./upgrade-node.pl` on a standby node.
- After the ZCS `install.sh` is found, you are asked to confirm that it is the right location? Enter **Y**.
- You are asked again if all Zimbra processes and cluster daemons are stopped. If you have verified that they are stopped, enter **Y**. The Zimbra-cluster package will now be upgraded.
- Enter the number for the standby node to be upgraded. Enter **Y** to verify your selection. The ZCS `./install.sh` script begins.

```
Choose a Zimbra service to upgrade:
 1) mail1.domain.com
 2) mail2.domain.com
 3) Stand-by node
Choose from above (1-3): 3
You chose mail1.domain.com. Is this correct? (Y/N) y
```

5. When **Finished upgrading this node** displays, the standby node has been upgraded. Proceed to update all standby nodes.

To manually stop a node

1. Log in to each server as root and type `su - zimbra`
2. To stop the services, type `zmcontrol stop`. Type `zmcontrol status` to verify that all zimbra processes are stopped.
3. Type `exit`.
4. To stop the rgmanager cluster daemon, type `service rgmanager stop`. Type `service rgmanager status` to verify that rgmanager is stopped.

Finishing the Multi-Node Upgrade

1. When all nodes have been upgraded, on an active node, run `./resume.all`. The services are started on all nodes.

Note: *Commands are executed with ssh and you may be prompted to enter the root password for each host.*

As rgmanager is started on each host, the active nodes start their assigned Zimbra services. You are asked if you are ready to continue. Enter **Y**.

2. When **Done** displays, all nodes have been upgraded. Run the `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 upgrade is complete.

Important: *After the upgrade, run a full backup immediately! Changes in this release invalidate all old backups.*

Testing the Cluster Setup

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 fences it, and brings up the failed service on a standby node.

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 1