
Upgrading Zimbra Collaboration Suite's One + One Cluster Environment

A one + one cluster upgrade is also referred to as a single node cluster. The cluster includes one active node with all services running on it, including LDAP and MTA, and one standby node. The cluster is under the control of the Red Hat Cluster Manager.

The one + one cluster upgrade process automatically stops both nodes in the cluster, unmounts the SAN, runs the ZCS upgrade scripts, and restarts both nodes in the cluster.

BACKUP BEFORE YOU BEGIN TO UPGRADE!

Before you begin the upgrade, run a full backup of the mailbox server. 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 have the license installed on the active node before you start the upgrade. The ZCSLicense.xml file should be copied to `/opt/zimbra/conf/`. 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 Single Node Cluster Environment

Save the ZCS upgrade files, `zcs<version>.tgz`, and the cluster upgrade files, `zcs-cluster.tgz` together in a directory on the active and standby nodes.

The ZCSLicense.xml file should be in the active node's `/opt/zimbra/conf` directory.

Stopping Cluster Nodes

1. Log in as **root** to the Zimbra server and **cd** to the directory where the Zimbra **zcs-cluster.tgz** and **zcs <version>.tgz** files 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 files in the same directory as where the ZCS files have been unpacked. The cluster upgrade scripts are in the **zcs-cluster/upgrade** directory.

2. After the files are unpacked, type **cd zcs-cluster/upgrade** to change to the cluster upgrade directory.
3. On the active node, run **./stop-all.pl** to stop the services. 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 the root password for each host.

4. At this point, enter **Y** to continue the upgrade. The rgmanager is stopped on both nodes.
5. If there are any errors, log in to each server and make sure rgmanager is stopped and is removed from the boot sequence. When you are ready to continue, enter **Y**.

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**.

6. Once the SAN volumes are unmounted, you can proceed. When you see **Are all SAN volumes unmounted?**, enter **Y**.

All the services on the active node are stopped.

Upgrading the Active Node

You must upgrade the active node before upgrading the standby node.

If you are upgrading a large number of accounts, the upgrade process can take some time. Use the **nohub** 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 node.

1. Enter **./upgrade-node.pl**, to start the upgrade scripts.
2. After ZCS **install.sh** is found, you are asked if this the right location? Enter **Y**.
3. You are asked again if all Zimbra processes and rgmanager cluster daemon is stopped. If it is stopped, enter **Y**. The Zimbra-cluster package is upgraded.
If they are not, go back to [Step 3](#) and start over.
4. Enter the number for the Zimbra sever to be upgraded as the active node. Enter **Y** to verify your selection. The ZCS **install.sh** script begins.
5. When **Finished upgrading this node** displays, the active node has been upgraded. Now proceed to upgrade the standby node.

Upgrading the Standby Node

The standby node was 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 node.

1. Run **./upgrade-node.pl** on the standby node.
2. After the ZCS **install.sh** is found, you are asked if this the right location? Enter **Y**.
3. You are asked again if all Zimbra processes and cluster daemons are stopped. If they are stopped, enter **Y**. The zimbra-cluster package will now be upgraded.
4. Enter the number for the Zimbra server to be upgraded as the standby node. Enter **Y** to verify your selection. The ZCS **install.sh** script is run.
5. When **Finished upgrading this node** displays, the standby node has been upgraded.

Finishing the Cluster Upgrade

1. When both nodes have been upgraded, on the active node, run **./resume.all**. The services are started on both 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 node starts its assigned Zimbra services. You are asked if you are ready to continue. Enter **Y**.

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

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

Testing the Cluster Set up

To perform a quick test to see if failover works:

1. Login to the remote power switch device and turn off the active node to simulate node failure.
2. To watch the standby node take over the failed service, run **clustat** on the standby node.
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 the 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