



## Deploying COS

---

To deploy the COS 2.1.2 software, perform the following tasks in order:

1. [Pre-Deployment Tasks](#)
2. [Setting Up the External DNS Server](#)
3. [Configuring COS](#)

### Pre-Deployment Tasks

Before deploying the COS, you must complete the following preliminary tasks:

- 
- Step 1** Prepare your COS VMware datacenter topologies and networks. The COS requires the following networks:
- Management — This is the primary COS network that connects all of the components.
  - Cache Fill — This is the data network. It connects the COS Appliances.
- Step 2** Download the COS Service Manager Open Virtual Appliance (OVA) file and load the image into a repository (HTTP server) that is accessible by vCenter.
- Step 3** Determine the blade and VM layout that you will use, and your management IP address allocation scheme for the network interfaces.
- Step 4** Install the external DNS and NTP servers.
- Determine the origin service FQDN and prepare the downstream client (client-facing) DNS servers to point to the COS Appliance dataplane (primary interface) IP addresses.

### COS Node Hardware Installation

COS software can currently be deployed on the Cisco Content Delivery Engine (CDE) platforms CDE460 and CDE470. For information about installing the CDE460 and CDE470 platforms, see the [Cisco Content Delivery Engine 205/220/250/420/460/470 Hardware Installation Guide](#).

## Upgrading to COS 2.1.2

If upgrading to COS 2.1.2 from an earlier COS release, perform the following steps on each COS node:


- 
- Step 1** Log in to the COS node as the root user.
- Step 2** Copy the COS 2.1.2 software ISO image to /root on each target COS node.
- Step 3** Mount the ISO image to /mnt/cdrom.
- ```
[root@cos-node ~]# mount -o loop /root/cos_repo-2.1.2.iso /mnt/cdrom
```
- Step 4** Prepare the YUM software repository using the following sequence of commands:
- ```
[root@cos-node ~]# rm -f /etc/yum.repos.d/*.repo
[root@cos-node ~]# /mnt/cdrom/local_repo_setup
[root@cos-node ~]# yum clean all
```
- Step 5** Upgrade the COS software using the following sequence of commands:
- ```
[root@cos-node ~]# yum -y groupinstall cos cos_support cos_test
[root@cos-node ~]# yum -y upgrade
[root@cos-node ~]# reboot
```

The system reboots and returns to service running the upgraded COS software version.

## PAM Installation and Configuration

To deploy the PAM from the vCenter, do the following:

- 
- Step 1** Log into the vCenter.
- Step 2** Click **File > Deploy OVF Template**.
- Step 3** Choose the cos-mgmt ova file and click **Next**.
- Step 4** Verify the OVF template details and click **Next**.
- Step 5** Accept the license agreement and click **Next**.
- Step 6** Enter a name and select a location for the deployed template. Then, click **Next**.
- Step 7** Choose the UCS on which the PAM is to be deployed and click **Next**.
- Step 8** Do not change the disk format. Click **Next**.
- Step 9** In the **Network Mapping** screen, choose **VM Network** from the **Destination Networks** drop-down list for source network **Network for adapter 1**. Click **Next**.
- Step 10** In the **Networking** section, enter the following information:
- **Hostname** — Local hostname of the PAM. Do not enter the fully qualified domain name (FQDN). For example, if the FQDN is pam.cisco.com, enter only pam.
  - **Domain** — The domain name used by the PAM, such as testdns.com (required).
  - **IP Address** — Primary management IP address to which the MOS can SSH (required). This is the IP address for the Management network for the PAM and for configuring Ethernet port eth0. The management IP address must be accessible from a management PC (for example, from a PC that is running vSphere).
  - **Subnet Mask** — Default subnet mask for the PAM (required).

- Gateway — Default gateway IP address for the PAM (required).
- Step 11** In the Additional Domain Name Servers section, enter the following information:
- DNS Server — Name or IP address of the DNS server used by the PAM to forward queries that it could not resolve (that is, the DNS forwarder). Enter names for all of the DNS servers used by the PAM.
- Step 12** In the Additional Network Time Protocol Servers section, enter the following information:
- NTP Server — Name or IP address of the external Network Time Protocol (NTP) server used by the PAM to synchronize its clock. Enter names for all of the NTP servers used by the PAM.
- You must enter at least one NTP server.
- Step 13** In the External DNS Credentials section, enter the following information:
- DNS Server — Name or IP address of the external DNS server used by the PAM to forward queries that it could not resolve.
  - Transaction Signature Key — Transaction Signature (TSIG) key that is configured on the external DNS server.
  - Transaction Signature Algorithm — Name of the TSIG algorithm that is configured on the external DNS server.
-  **Note** The administrator must ensure that the DNS zones are configured in the external DNS server.
- Step 14** Click **Next**, then click **Finish** to deploy the OVA and create the PAM VM. The progress bar shows the status of the deployment.
- Step 15** After the deployment is complete, map the adapter to the correct network, and power on the PAM VM.
- Step 16** Verify connectivity and open an SSH session into the PAM VM.

## Upgrading the PAM COS Controller

When upgrading to COS 2.1.2 from an earlier COS release, perform the following steps to upgrade the PAM COS controller:

- Step 1** Log on to the PAM COS controller as the root user.
- Step 2** Copy the `cos-aic rpm` to be upgraded to the PAM root directory.
- Step 3** Run the command `supervisorctl stop cos-controller`.
- Step 4** Get the current `cos-controller rpm` using the command `rpm -qalgrep cos-aic`.
- Step 5** Remove the current rpm using `rpm -e current cos-aic.rpm`.
- Step 6** Install the new rpm using the command `rpm -ivh new cos-aic.rpm`.
- Step 7** Start the `cos-controller` again using `supervisorctl start cos-controller`.

## Downgrading the PAM COS Controller

When downgrading from COS 2.1.2 to an earlier COS release, perform the following steps to downgrade the PAM COS controller:

- 
- Step 1** Log in to the PAM Console as the root user.
  - Step 2** Copy the `cos-aic` rpm to be downgraded to the PAM root directory.
  - Step 3** Run the command **`supervisorctl stop cos-controller`**.
  - Step 4** Get the current `cos-controller` rpm using the command **`rpm-qalgrep cos-aic`**.
  - Step 5** Remove the current rpm using `rpm -e current cos-aic.rpm`.
  - Step 6** Install the new rpm using the command **`rpm -ivh new cos-aic.rpm`**.
  - Step 7** Start the `cos-controller` again using **`supervisorctl start cos-controller`**.

**Note**

---

The RPM `current cos-aic.rpm` is the output of the previous command `rpm-qalgrep cos-aic`, which is currently running on the system. The RPM `cos-aic.rpm` is the new RPM to be installed on the PAM.

---

## Setting Up the External DNS Server

The administrator is responsible for creating the transaction signature (TSIG) key for the TSIG algorithm on the external DNS server. Valid TSIG algorithms are `hmac-md5`, `hmac-sha1`, `hmac-sha224`, `hmac-sha256`, `hmac-sha384`, and `hmac-sha512`. The following procedure shows how to create the TSIG key for the `hmac-md5` algorithm.

- 
- Step 1** Generate the TSIG key by entering the following command:

```
dnssec-keygen -a HMAC-MD5 -b 128 -n HOST testdns.com.
```

where:

- **HMAC-MD5** is the TSIG algorithm.
- **128** is the number of bits in the key.
- **testdns.com.** is the name of the key. The name of the key, the domain name of the PAM, and the DNS zone in the external DNS server should all be the same (in this example, **testdns.com**).
- The command must end with a period (`.`), which is required when generating the key.

Sample output:

```
dnssec-keygen -a HMAC-MD5 -b 128 -n HOST testdns.com.
```

```
Ktestdns.com.+157+05519
```

This command creates a key file and a private key file.

- Sample `.key` file:

```
Ktestdns.com.+157+05519.key
testdns.com. IN KEY 512 3 157 ujLdXfCZenQQZKZIFy42fw==
```
- Sample `.private` key file:

```

Ktestdns.com.+157+05519.private
\Private-key-format: v1.3
Algorithm: 157 (HMAC_MD5)
Key: ujLdXfCZenQZQKZIFy42fw==
Bits: AAA=
Created: 20140325141250
Publish: 20140325141250
Activate: 20140325141250

```

**Step 2** Create the key file (in this example, testdns.com.key) in the /etc/ directory.

Sample key file:

```

testdns.com.key
key testdns.com. {
    algorithm hmac-md5;
    secret "ujLdXfCZenQZQKZIFy42fw==";
};

```

where:

- **hmac-md5** is the TSIG algorithm.
- **ujLdXfCZenQZQKZIFy42fw==** is the TSIG key.

**Step 3** Add the key file path to the /etc/named.conf file by inserting the following line:

```
include "/etc/testdns.com.key";
```

**Step 4** Configure the DNS zones (used when deploying the PAM) on the external DNS server as shown in the following examples.

- a. Update the /etc/named.conf file with the details of the DNS zone (in this example, testdns.com) and reverse zone, for all interfaces.

In the following sample information:

- The DNS zone is related to the 172.20.216.xx subnet, which can be the Management interface.
- For the Data In and Data Out interfaces, similar information related to the reverse zone must be added.
- Data In is related to the 15.1.1.x subnet.
- Data Out is related to the 25.1.1.x subnet.
- **testdns.com.** is the key name.

#### DNS Zone Details

```

zone testdns.com IN {
type master;
file "slaves/db.testdns.com";
allow-update { key "testdns.com."; };
notify yes;
};

```

#### Reverse Zone Details (Management Interface)

```

zone 216.20.172.IN-ADDR.ARPA IN {

```

```

type master;
file "slaves/db.216.20.172";
allow-update { key "testdns.com."; };
notify yes;
};

```

#### Reverse Zone Details (Data In Interface)

```

zone 1.1.15.IN-ADDR.ARPA IN {
type master;
file "slaves/db.1.1.15";
allow-update { key "testdns.com."; };
notify yes;
};

```

#### Reverse Zone Details (Data Out Interface)

```

zone 1.1.25.IN-ADDR.ARPA IN {
type master;
file "slaves/db.1.1.25";
allow-update { key "testdns.com."; };
notify yes;
};

```

- b. Create db.\* files in the /var/named/slaves directory in the external DNS server for the DNS zone and reverse zone, for all interfaces.

#### Sample db File for DNS Zone (db.testdns.com)

```

@ 86400 IN SOA dns pam 2014031001 3600 1800 604800 86400
@ 86400 IN NS dns
dhcp 86400 IN CNAME dns
dns 86400 IN A pam_ip_address
ntp 86400 IN CNAME dns

```

#### Sample db File for Reverse Zone of Management Interface (db.216.20.172)

```

$ORIGIN .
$TTL 86400      ; 1 day
216.20.172.IN-ADDR.ARPA  IN SOA  dns.testdns.com.  pam.testdns.com. (
                                2012071021 ; serial
                                3600 ; refresh (1 hour)
                                1800 ; retry (30 minutes)
                                604800 ; expire (1 week)
                                86400 ; minimum (1 day)
                                )
    NS      dns.testdns.com.
$ORIGIN 216.20.172.IN-ADDR.ARPA.
$TTL 7200      ; 2 hours
xx              PTR      dns.testdns.com. //xx = last two digits of the PAM IP

```

**Sample db File for Reverse Zone for Data In Interface (db.1.1.15)**

```

$ORIGIN .
$TTL 86400      ; 1 day
1.1.15.IN-ADDR.ARPA  IN SOA  dns.testdns.com.  pam.testdns.com. (
                        2012071021 ; serial
                        3600 ; refresh (1 hour)
                        1800 ; retry (30 minutes)
                        604800 ; expire (1 week)
                        86400 ; minimum (1 day)
                        )
NS      dns.testdns.com.
$ORIGIN 1.1.15.IN-ADDR.ARPA.
$TTL 7200      ; 2 hours
1          PTR      dns.testdns.com.

```

**Sample db File for Reverse Zone for Data Out Interface (db.1.1.15)**

```

$ORIGIN .
$TTL 86400      ; 1 day
1.1.25.IN-ADDR.ARPA  IN SOA  dns.testdns.com.  pam.testdns.com. (
                        2012071021 ; serial
                        3600 ; refresh (1 hour)
                        1800 ; retry (30 minutes)
                        604800 ; expire (1 week)
                        86400 ; minimum (1 day)
                        )
NS      dns.testdns.com.
$ORIGIN 1.1.25.IN-ADDR.ARPA.
$TTL 7200      ; 2 hours
1          PTR      dns.testdns.com.

```

**Step 5** Make sure that the PAM and the external DNS server are in sync for the time and date.

## Editing NTP Server Information

**Note**

NTP servers are configured when the PAM is deployed. You cannot add or delete NTP servers using this procedure.

To edit NTP server information using the COS Service Manager GUI, follow these steps:

---

**Step 1** Log in to the COS Service Manager GUI as described in [Appendix A, “Reference Information”](#).

**Step 2** Go to **Infrastructure > Platform Services**.

In the **NTP Servers** table, the following information is displayed for each NTP server:

- Name (required)
  - The name of the NTP server is a string of up to 30 characters. Acceptable characters include uppercase and lowercase letters, numbers, periods (.), dashes (-), and underscores (\_).
  - The name must not begin with a period (.), and it is not case-sensitive.
- Description — The description is a string of up to 100 characters, and can include uppercase or lowercase letters, numbers, and any special characters.
- Region — The region in which the NTP server resides.
- Servers — The hostname or IPv4 address of the NTP server.

**Step 3** Edit the Description, Region, and Servers fields of an NTP server as required.

## Viewing DNS Server Information



### Note

---

DNS servers are configured when the PAM is deployed.

---

To view DNS server information using the COS Service Manager GUI, follow these steps:

---

**Step 1** Log in to the COS Service Manager GUI as described in [Appendix A, “Reference Information”](#).

**Step 2** Go to **Infrastructure > Platform Services**.

In the **DNS Servers** table, the following information is displayed for each DNS server:

- Name (required)
  - The name of the DNS server is a string of up to 35 characters. Acceptable characters include uppercase and lowercase letters, numbers, periods (.), dashes (-), and underscores (\_).
  - The name cannot begin with a period (.), and it is not case-sensitive.
- Description — The description is a string of up to 100 characters, and can include uppercase or lowercase letters, numbers, and any special characters.
- Region — Region in which the DNS server resides.
- IPv4 Address — IPv4 address of the DNS server.
- TSIG — Transaction Signature (TSIG) of the DNS server.
- TSIG Algorithm — Name of the TSIG algorithm used by the DNS server.

## Configuring COS

To configure COS using the COS Service Manager GUI, follow these steps:

---

**Step 1** Log in to the COS Service Manager GUI as described in [Appendix A, “Reference Information”](#).



**Step 2** Go to **Infrastructure > Regions & Zones**.

A *zone* is a set of Cloud platform components (compute, network, storage, and security) that are fate-shared. The zone can be mapped to the underlying Cloud platform provider, such as a datacenter in vCenter, an Availability Zone, or any other combination of fate-shared Cloud resource topologies. Each zone is associated with one Cloud Controller.

A *region* is made up of one or more zones. It is an abstract representation of the underlying Cloud platform. A region can be associated with a geographical region, one or more datacenters, or a service area. Release 2.1.2 of the COS software supports only one region.

**Step 3** Expand a region to display the zones associated with that region.**Step 4** Click **Add Row** to add zones to the region, as needed.**Step 5** Go to **Infrastructure > Networks** to define IP pools for COS Node cache fill interfaces.**Step 6** On the COS node, execute `/opt/cisco/cos/config/cosinit` and configure the IP pool name, DocServer IP and Port (which is the COS Service Manager management IP). This will register the COS node with the COS Service Manager.**Step 7** Go to **Infrastructure > COS Nodes** and configure **Cluster Name** by choosing **cluster.ca.01** from the drop-down list.**Step 8** Ensure that the **Admin State** for the node is set to **Inservice**. Click **Save**.**Step 9** Go to **Services > Overview** and confirm that you can see the registered COS Node.

## Configuring the Service Instance Template

To configure the COS service instance template using the COS Service Manager GUI, do the following:

**Step 1** Log in to the COS Service Manager GUI as described in [Appendix A, “Reference Information”](#).**Step 2** Go to **Services > Cloud Object Storage**. The **Service Summary** page opens, displaying a list of COS service instances.**Step 3** Click the COS service instance to be configured. The **Service Definition** tab for that instance opens, containing the **General** and **Service Endpoints** areas.**Step 4** The fields of the **General** area are listed below. You can edit them appropriately.

- Title — The title of the COS service. The default title is **An unused COS Service**.
- Description — A brief description of the COS service. The default description is **COS Service**.
- Admin State — The administrative state of the COS service. From the drop-down list, you can choose to **Enable** or **Disable** the service.
- Service Template — Currently, this read-only field identifies the default service template, **Cisco Object Store (COS) Service**.

**Step 5** The **Service Endpoints** area currently lists only one service endpoint. The following parameters are displayed for the service endpoint and can be edited appropriately:

- Name — The name of the service endpoint is set to the default value **ce1** and cannot be changed.
- Description — A brief description of the service endpoint.
- Region — The region to which the service endpoint belongs.
- Min Nodes — The minimum number of nodes that must be associated the service endpoint.
- Desired Nodes — The desired number of nodes for the service endpoint.

- Max Nodes — The maximum number of nodes that can be associated with the service endpoint.
- Max Storage — This field is to be ignored for Release 2.0.1.
- Cluster Name — From the drop-down list, you can currently choose only the default COS cluster **cluster.ca.01**.
- Auth Profile — The default value is **auth-1**.
- State — The state of the service endpoint. From the drop-down list, choose **Disabled** or **Enabled**.

**Step 6** For your changes to take effect, click **Save**. To discard your changes, click **Cancel**.

## Configuring IP Pools

Before installing COS nodes, you must configure IP pools using the COS Service Manager GUI.

To configure an IP pool, follow these steps:

- 
- Step 1** Log in to the COS Service Manager GUI as described in [Appendix A, “Reference Information”](#).
- Step 2** Go to **Infrastructure > IP Pools and Networks**.
- Step 3** To add a new IP pool, click **IP Address Pools > Add Row**.
- Step 4** In the **Name** field, enter a unique name for the IP pool. In the **Description** field, enter a brief description for the IP pool.
- Step 5** To associate an IP address range with the IP pool, click **IP Address Ranges > Add Row**.
- Step 6** Enter appropriate values in the fields described below:

**Table 2-1** IP Address Ranges Fields

| Field       | Description                        |
|-------------|------------------------------------|
| Range Start | The first IP address of the range. |
| Range End   | The last IP address of the range.  |
| Netmask     | The netmask for the range.         |
| Gateway     | The Gateway for the range.         |



**Note** If you attempt to initialize a COS node using an IP pool with an insufficient number of available IP addresses, the initialization fails and an Event is generated. For more information on this class of events, see [“COS AIC Server Events”](#).

- Step 7** For your changes to take effect, click **Save**. To discard your changes, click **Cancel**.



**Note** After you have associated an IP pool with a COS node, do not edit or delete the pool before dissociating the pool from the node.

## Editing or Deleting IP Pools

To edit or delete configured IP pools using the COS Service Manager GUI, follow these steps:

- 
- Step 1** Log in to the COS Service Manager GUI as described in [Appendix A, “Reference Information”](#).
  - Step 2** Go to **Infrastructure > IP Pools and Networks**.
  - Step 3** Before editing or deleting an IP pool, ensure that it is not serving any COS nodes.
    - a. Dissociate any COS nodes associated with the IP pool to be edited or deleted, and link these nodes to another IP pool.
  - Step 4** Check the box against the name of the IP pool to be edited or deleted.
  - Step 5** Click **Edit** or **Delete**.

## Configuring a COS Cluster

As part of the COS installation, a single default node cluster is created. The cluster can be configured using the COS Service Manager GUI.

To view or edit the settings of a COS node cluster, follow these steps:

- 
- Step 1** Log in to the COS Service Manager GUI as described in [Appendix A, “Reference Information”](#).
  - Step 2** Go to **Infrastructure > COS Clusters**.  
The default COS node cluster is displayed.
  - Step 3** Edit the fields described below based on your deployment.

**Table 2-2** COS Node Cluster Fields

| Field               | Description                                                  |
|---------------------|--------------------------------------------------------------|
| Authentication FQDN | The FQDN for COS authentication requests.                    |
| Storage FQDN        | Currently, this must be the same as the Authentication FQDN. |

- Step 4** For your changes to take effect, click **Save**. To discard your changes, click **Cancel**.

## COS Node Initial Configuration

COS nodes ship with a pre-installed image. While provisioning a COS node on a network, you can upgrade the node to a different image.

After the image is installed and the node rebooted, you will be prompted to start configuration by running the `cosinit` script, which must be run to perform the initial setup and add the node to the COS management system.

- Enter **yes** to run `cosinit` immediately.
- If you wish to configure the node later, enter **no**. You can run the `cosinit` later using the following command: `run /opt/cisco/cos/config/cosinit`

The `cosinit` script can be run with the options shown below:

```
cosinit [-skipnw] [-input <configFile>] [-help]
```

- [-skipnw] — Skip network configurations
- [-input <configFile>] — Specify an input file for configuration
- [-help] — Displays the usage

#### Examples:

```
/opt/cisco/cos/config/cosinit -skipnw
```

This is used to configure IP pool and DocServer, and register the cos node into COS Service Manager.

```
/opt/cisco/cos/config/cosinit -input configFile
```

This enables cosinit to run with a config file. All the parameters will be read from the config file.

Sample configFile

```
InterfaceConfig      : skip
```

(Optional. If the above line is added, management network configuration will be skipped.)

```
IpAddress            : 172.22.99.218
NetMask              : 255.255.254.0
Broadcast            : 172.22.99.255
DefaultGateway       : 172.22.98.1
HostName             : CDE460-218
PartNum              : CDE460-4R1
DocServerHost        : 172.22.116.35
DocServerPort        : 5087
BaudRate             : 9600
eth4                 : 4
eth5                 : 4
eth6                 : 4
eth7                 : 4
```

Executing the cosinit script completes the following tasks:

- Configuring the management interface — ethX, IP, netmask, broadcast.
- Configuring the default gateway.
- Configuring the hostname.
- Mapping the network interfaces to IP addresses in the network.
- Configuring Doc Server IP address and Port.
- Recording appliance model name — CDE 470-4R2 or CDE 460-4R1.
- Configuring the baud rate by reading the value in install\_baud\_rate. If the install\_baud\_rate file is not found under the /root directory, the operator is prompted to enter a baud rate during cosinit execution or specify the value in the configFile.

When the cosinit script is successfully executed,

- /etc/sysconfig/network-scripts/ifcfg-eth0 is generated and the management interface configuration is saved in it.
- /etc/sysconfig/network is generated, and the hostname and default gateway info are saved in it.
- /etc/hosts is generated and the host info is saved in it.
- Network is restarted.
- /tmp/.cosnodeinit file is created with the configuration specified in cosinit.

- The baud rate is updated in the `/boot/grub/grub.conf` file.
- The COS AIC client and other applications are started.
- The COS AIC client reads the `cosnodeinit` generated by `cosinit`.
- Using the Doc Server IP address and port mentioned in `cosnodeinit`, the AIC client connects to the Doc Server.
- The AIC client sends the `cosannounce` to the Doc Server. `cosannounce` is a partially populated `smcosnode` document.

The COS node is now registered with the COS management system. You can configure the node using the COS Manager GUI, add it to a COS cluster, and make it operational.

## Adding a COS Node to a COS Cluster

After a COS node is installed and initialized, it must be added to a COS cluster to begin servicing COS requests.

To add a node to the COS cluster using the COS Service Manager GUI, do the following:

- 
- Step 1** Log in to the COS Service Manager GUI as described in [Appendix A, “Reference Information”](#).
  - Step 2** Go to **Infrastructure > COS Nodes**.
  - Step 3** Check the box against the name of the node you want to edit and click **Edit**.
  - Step 4** From the **Cluster Name** drop-down list, choose the default COS cluster, **cluster.ca.01**.
  - Step 5** For your changes to take effect, click **Save**. To discard your changes, click **Cancel**.

When a COS node is added to a COS cluster, the following files are written on to the node:

- `setupfile` — This file holds the primary configuration data for the COS node and is required by CServer.
- `RemoteServers` — This file holds the service interface IP addresses for all of the COS nodes in the cluster.
- `SubnetTable` — This file holds the IP, Netmask, Gateway and Network data for each service interface of a COS node.
- `cosd.conf` — This file holds the configuration for the `cosd` service.
- `cassandra.yaml` — This file, located in `/etc/cassandra/conf/`, is responsible for the configuration of the Cassandra service.

When the node **Admin State** is set to **Inservice**, all enabled service interfaces of that node are written to the DNS (internal or external).



### Note

To avoid possible bootstrapping issues with the Cassandra service, be sure to provide (or use a script to provide) a delay of at least two minutes between adding two nodes in sequence to the cluster.

If the replication factor is too high, use the following script on any one node in the cluster to adjust it:

```
sh /opt/cisco/cos-aic-client/cassandra/cassandra-adjust-replication.sh { 1 | 2 | 3 }
```

For the final argument, use 1 for one node, 2 for two nodes, and 3 for three or more nodes.

## Configuring a COS Node

After a COS Node is installed and initialized, you can modify its configuration parameters using the COS Service Manager.

To edit the settings of a functioning COS node using the COS Service Manager GUI, follow these steps:

**Step 1** Log in to the COS Service Manager GUI as described in [Appendix A, “Reference Information”](#).

**Step 2** Go to **Infrastructure > COS Nodes**.

**Step 3** Check the box against the name of the node you want to edit and click **Edit**.

**Step 4** Edit the fields described below based on your requirements.

**Table 2-3** *COS Node Fields*

| Field        | Description                                                                  |
|--------------|------------------------------------------------------------------------------|
| Description  | A brief description of the node.                                             |
| Model        | The part number of the node.                                                 |
| Zone         | The zone to which the node belongs.                                          |
| Cluster Name | The COS cluster to which the node belongs.                                   |
| Admin State  | Indicates whether the node is <b>Inservice</b> or under <b>Maintenance</b> . |



**Note** The Admin State of a COS node must be set to **Maintenance** before removing the node from a COS cluster.

**Step 5** To modify the service interfaces of the node, follow the procedure [“Configuring the Service Interface of a COS Node”](#).

**Step 6** For your changes to take effect, click **Save**. To discard your changes, click **Cancel**.

## Configuring the Service Interface of a COS Node

To modify the settings of a service interface of a functioning COS node using the COS Service Manager GUI, follow these steps:

**Step 1** Log in to the COS Service Manager GUI as described in [Appendix A, “Reference Information”](#).

**Step 2** Go to **Infrastructure > COS Nodes**.

**Step 3** Check the box against the name of the node you want to edit and click **Edit**.

**Step 4** Check the box against the name of service interface you wish to edit and click **Edit**.

**Step 5** Edit the fields described below based on your requirements.

**Table 2-4 Service Interface Fields**

| Field   | Description                                                                                                                                                                                          |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IP Pool | The IP pool from which this interface will be assigned an IP address. You may associate a different IP pool with this interface only when the node <b>Admin State</b> is set to <b>Maintenance</b> . |
| Enabled | From the drop-down list, choose <b>True</b> or <b>False</b> to enable or disable the service instance, respectively.                                                                                 |

**Step 6** For your changes to take effect, click **Save**. To discard your changes, click **Cancel**.

## Configuring Cassandra Database Maintenance

The Cassandra database requires periodic maintenance with an anti-entropy repair that must be manually configured to execute every two days on each COS node. For more information on the Cassandra repair process, see the Cassandra 2.1 documentation available at:

[www.datastax.com/documentation/cassandra/2.1/](http://www.datastax.com/documentation/cassandra/2.1/)

To execute the periodic repair, we recommend configuring a CRON job on each node, and scheduling the CRON jobs on the nodes to begin at different times so as to avoid the repair running concurrently on more than one node at a time. As a reference, we recommend providing 15 minutes between scheduling of the repair operation on each node.

The command syntax for the repair is:

```
/usr/bin/nodetool repair -par -inc
```

