Deploying COS

To deploy the COS 3.5.2 software, perform the following tasks in order:

1. Pre-Deployment Tasks, page 2-1
2. Setting Up the External DNS Server, page 2-3
3. PAM Installation, page 2-6
4. COS Installation on the C3x60, page 2-7
5. COS Installation on the CDE465, page 2-15
6. Viewing NTP Server Information, page 2-20
7. Viewing DNS Server Information, page 2-21
8. Configuring COS from the GUI, page 2-21
9. Automated COS Node Configuration (Optional), page 2-28
10. Configuring Cassandra Database Maintenance, page 2-30

Note COS Release 3.5.2 and later support remote network installation of the COS client using the Intel Preboot Execution Environment (PXE) in combination with the Red Hat Enterprises network installation feature using NFS, FTP, or HTTP, PXE. For details, see PXE Network Installation, page F-1.

Pre-Deployment Tasks

Before deploying the COS, complete the following preliminary tasks:

Step 1 Prepare your COS VMware datacenter topologies and networks. 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.

Hardware Options

COS 3.5.2 software currently can be deployed on the following hardware:

- Cisco UCS C3260-4U3 Dual Node Rack Server with 56 x 6 TB hard drives (336 TB total storage, 28 hard drives or 168 TB per server node)
- Cisco UCS C3160-4U2 Rack Server with 54 x 6 TB hard drives (324 TB total storage)
- Cisco UCS C3160-4U1 Rack Server with 54 x 4 TB hard drives (216 TB total storage)
- Cisco Content Delivery Engine CDE465-4R4 with 36 x 6 TB hard drives (216 TB total storage)

Note You can convert a C3160 to a C3260 in the field. For details, see **Migrating a Cisco UCS C3160 Server to a Cisco UCS C3260 Server** in the *Cisco UCS C3260 Rack Server Installation and Service Guide*.

For information about installing the hardware, see the following:

- *Cisco UCS C3260 Rack Server Installation and Service Guide*
- *Cisco UCS C3160 Rack Server Installation and Service Guide*
- *Cisco Content Delivery Engine 465 Hardware Installation Guide*

Before you begin, be sure that you have the following:

- Server hardware installed per manufacturer instructions
- Cisco Integrated Management Controller (CIMC) Connection to the server
- ISO image of the COS Software

COS Network Architecture

Figure 2-1 provides a view of the network architecture of a COS cluster based on the UCS C3x60.

Figure 2-1 C3x60 COS Cluster Architecture

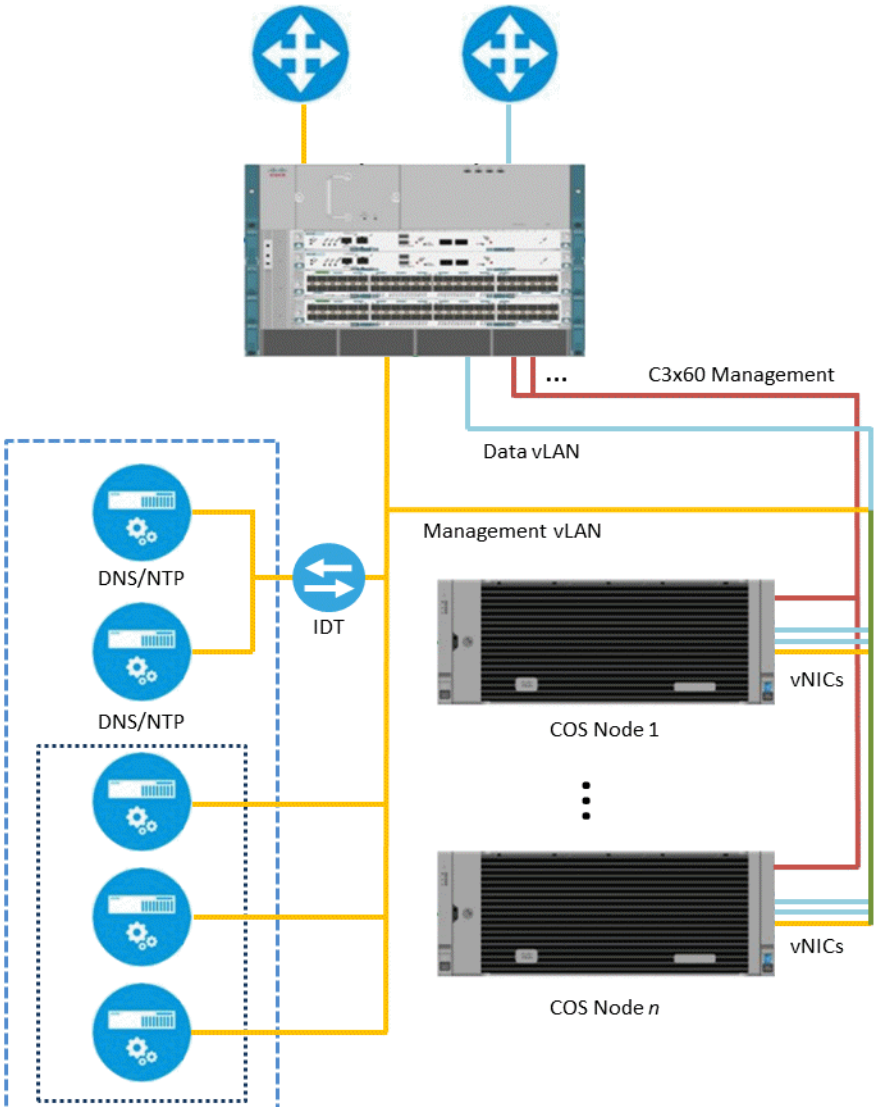

Setting Up the External DNS Server

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
- 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 ujLdXfCZenQZQKZlFy42fw==
```

- Sample .private key file:

```
Ktestdns.com.+157+05519.private
\Private-key-format: v1.3
Algorithm: 157 (HMAC_MD5)
Key: ujLdXfCZenQZQKZlFy42fw==
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 "ujLdXfCZenQZQKZlFy42fw==";
};
```

where:

- **hmac-md5** is the TSIG algorithm.
- **ujLdXfCZenQZQKZlFy42fw==** 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 ns.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.

PAM Installation

To deploy the PAM from the vCenter:

Step 1 Log into the vCenter.

Step 2 Click **File > Deploy OVF Template**. The Source dialog appears.

Step 3 Select the location of the OVF template (for example, **mendocino-shenandoah-301.16144**) and click **Next**. The OVA Template Details dialog appears.

Step 4 Verify the OVF template details and click **Next**. The End User License Agreement dialog appears.

Step 5 Accept the End User License Agreement and click **Next**. The Name and Location dialog appears.

Step 6 Specify a name and location for the deployed template and click **Next**. The Deployment Configuration dialog appears.

Step 7 Select the **MOS PAM** deployment configuration and click **Next**. The Storage dialog appears.

Step 8 Select a destination storage file for the VM files and click **Next**. The Disk Format dialog appears.

Step 9 Confirm that **Thin Provision** disk format for virtual disk storage is selected and click **Next**. The Network Mapping dialog appears.

Step 10 For each Source Network listed, choose the appropriate **Destination Network** and click **Next**. The Properties dialog appears.

Step 11 Enter or choose the following information to customize the software solution for the deployment:

- IP address(es) of the NTP server(s)
- IP address of the DNS server
- External DNS transaction signature key (enter twice to confirm)
- Transaction signature algorithm
- IP address or hostname of two HA peers, if enabled
- (Optional) machine hostname
- Machine domain name
- Network adapter IP address
- Network adapter subnet mask
- Network adapter gateway IP address

Step 12 Click **Next**, and then click **Finish** to deploy the OVA and create the PAM VM. The progress bar shows the status of the deployment.

Step 13 Verify connectivity and open an SSH session into the PAM VM.

COS Installation on the C3x60

Note on Crash Partition Location

When installed on a C3160, the previous (3.5.1) COS release created a crash partition on one of the SSDs at the rear of the chassis. With COS Release 3.5.2, the location of the crash partition depends on the node hardware, as follows:

- When installed on a C3260, COS 3.5.2 creates a crash partition along with other system partitions on the software RAID SSDs at the rear of the chassis.
- When installed on a C3160, COS 3.5.2 creates a crash partition along with other system partitions on the RAID system drives, which are the SSDs in chassis slots 55 and 56.

These locations assume a fresh installation and not an upgrade (not supported in COS 3.5.2 in any case).

C3260 Installation

Note COS 3.5.2 is fully qualified for the UCS C3260 dual-node configuration. Qualification and production support for C3260 single-node 28- and 56-disk configurations is planned for a future COS release. The single-node options described below, while logically valid, are not currently supported.

Network Requirements

The C3260 provides four 40 Gbps ports for video and management traffic. To separate the management traffic from the content data, each port on the system I/O controller (SIOC) is divided into two four virtual NICs (vNICs) and assigned to one of two user-defined virtual LANs (vLANs). The first vNIC on each port is limited to 1 Gbps and is assigned to the management vLAN. The remaining vNICs are assigned to the data vLAN.

Each vNIC appears in Linux as a physical network interface. A pre-installation script, described below, is provided to establish this vNIC and vLAN configuration.

Note You must configure the first downstream switch(es) to receive the traffic from the two defined vLAN IDs, and route the traffic through the switch appropriately.

Pre-Installation Script for C3260 Servers

Before installing COS on a Cisco UCS C3260 Rack Server, you must configure the C3260 for either single-node service (not fully supported in COS 3.5.2) or dual-node service, and configure the data and management vLAN IDs for one or both nodes, as needed. For a C3260 in dual-node configuration, you will then install COS on each of the two server nodes separately.

A Bash script named **preinst_setup_UCSC-C3260.sh** is provided to help configure the C3260. This script is located at the root of the full COS product ISO image, which you can download from the COS product support page.

Note You must execute the script from a remote Linux node with network access to the target C3260 CIMC. Execution occurs after the C3260 CIMC IP address is defined, but before COS is installed.

The help for this script appears as follows:

```
# ./preinst_setup_UCSC-C3260.sh -h
Setup UCS C3260 for VDS product installation
   options:
      -a <speed>     Admin Speed ('4x10', '1x40')
      -c <ip>        CIMC IP address
      -s <ip1[,ip2]> SIOC IP address
      -b <ip1[,ip2]> BMC IP address
      -d <id>        Data vlan ID
      -m <id>        Management vlan ID
      -p <pw>        CIMC password
      -t             Test: Write, but do not execute, the commands to preinst_setup.log
      -u <user>      CIMC username
      -N             Nondestructive; will not modify drive configuration
```

Note

- Option **-c** is the IP address defined for the CIMC.
- Option **-a** defines the cable type plugged in the two SIOC network ports.
- The C3260 requires three IP addresses for single-node configuration and five IP addresses for dual-node configuration. If the chassis has two compute notes, define the additional IP addresses for the second compute note.
- Use of the **-N** option is not normally required.

Example - Dual-Node Configuration

```
./preinst_setup_UCSC-C3260.sh -a 1x40 -c 99.99.193.132 -s 99.99.193.133,10.99.193.134 -b
99.99.193.135,99.99.193.136 -d 100 -m 2 -p <password> -u admin
```

Script Behavior

This script removes all currently defined virtual RAID devices.

For C3260s in single-node configuration, all drives are assigned to the one node. For dual-node systems, the storage devices are divided equally between nodes. The drives in slots 1-28 are assigned to compute node 1, and the drives in slots 29-56 are assigned to compute node 2.

As described above, the vNICs are configured so that each 40 Gbps port on the SIOC has one 1 Gbps virtual management interface and three virtual data interfaces assigned to their provided vLAN.

Note In the CIMC, the vNICs are listed in their order of appearance in Linux. Thus, eth0 and eth1 are the management interfaces, and all other eth ports are the data interfaces.

C3160 Installation

Note You can convert a C3160 to a C3260 in the field. For details, see **Migrating a Cisco UCS C3160 Server to a Cisco UCS C3260 Server** in the *Cisco UCS C3260 Rack Server Installation and Service Guide*.

Network Requirements

The C3160 provides a total of four 10 GbE ports for video data and management traffic. To separate the management traffic from content data, port 1 on each SIOC is divided into two vNICs, with the management vNIC limited at 1 Gbps and the data vNIC limited at 9 Gbps. The second port on each SIOC is configured as a 10 Gbps data vNIC. The management vNICs are assigned to a user-defined management vLAN ID. The remaining data vNICs are assigned to a user-defined vLAN ID.

Each vNIC appears in Linux as a physical network interface. A pre-installation script, described below, is provided to establish this vNIC and vLAN configuration.

Note You must configure the first downstream switch to receive the traffic from the two vLAN IDs and route the traffic through the switch appropriately.

Pre-Installation Script for C3160 Servers

A Bash script named **preinst_setup_UCSC-C3160.sh** is provided to help configure the C3160. This script is located at the root of the full COS product ISO image, which you can download from the COS product support page.

Note You must execute the script from a remote Linux node with network access to the target C3160 CIMC. Execution occurs after the C3160 CIMC IP address is defined, but before COS is installed.

The help for this script appears as follows:

```
# ./preinst_setup_UCSC-C3160.sh -h
Setup UCS C3160 for VDS product installation
   options:
      -c <ip>        CIMC IP address
      -d <id>        Data vlan ID
      -m <id>        Management vlan ID
      -p <pw>        CIMC password
      -u <user>      CIMC username
```

When it runs, the script prompts for any parameters missing from the command line. Password values entered in response to such a prompt are not echoed to the screen.

Example

```
./preinst_setup_UCSC-C3160.sh -c 10.10.10.10 -u admin -m 42 -d 2001
Password:
```

C3x60 Installation and Configuration

Installing and configuring COS on the C3x60 involves the following procedures:

1. Configure the PAM using the COS Configuration Wizard
2. Prepare the C3x60 for COS installation
3. Install the COS ISO image on the C3x60
4. Configure the COS node network on the C3x60
5. Register COS nodes to the PAM
6. Create user accounts and verify node and cluster health

This section describes each of these steps.

Configure the PAM Using the COS Configuration Wizard

Step 1 Log into the COS Service Manager GUI as described in Using the COS Service Manager GUI, page B-1.

Step 2 Follow the steps described in Using the COS Configuration Wizard, page B-12 to configure the PAM.

Note For this procedure, you must include the final step, Create and Configure Initialization Profiles (Optional), page B-15. This step creates URLs that you use later to add COS nodes to the PAM.

Prepare the C3x60 for COS Installation

Step 1 For the C3160 only, remove any hard drives from the rear horizontal tray (up to four possible).

Step 2 Connect the CIMC Ethernet port on SIOC1 (the rightmost RJ45 port on each SIOC) to your management network.

Step 3 Connect the 40 Gbps (for C3260) or 10 Gbps (for C3160) cables to your downstream switch.

Note Your video network switch should already be configured to support the data and management vLAN IDs.

Step 4 Connect the power cords to the server.

Step 5 Temporarily attach a monitor and keyboard to the server.

Step 6 Apply power to the server, and immediately after the function key banner appears (early in the power-up sequence), press **F8** to enter Cisco CIMC Configuration.

Step 7 Under NIC Properties, set the CIMC IP, mask, and gateway as appropriate for your installation.

Step 8 Press **F10** to save your changes, then press **F5** to refresh the screen and confirm that the IP address is set properly.

Step 9 Press **ESC** to exit and reboot (the system will not yet fully reboot).

Step 10 Configure the server as follows:

- For the C3260, configure for single-node or dual-node service and configure vLANs using the C3260 pre-installation script as described in Pre-Installation Script for C3260 Servers, page 2-8.
- For the C3160, configure the vLANs using the C3160 pre-installation script as described in Pre-Installation Script for C3160 Servers, page 2-10.

Install the COS ISO Image on the C3x60

Step 1 Attach the COS installation ISO file using *one* of the following methods:

- Connect an external DVD drive at any USB port on the C3x60.

Note When using an external DVD drive, installation may fail with missing files reported. If this occurs, the likely cause is that the DVD drive is drawing excessive current from the USB port of the C3x60. To work around this issue, use a DVD drive that has its own AC power adapter.

- Open a web browser and use the IP address provided to CIMC in step 6 of Prepare the C3x60 for COS Installation, page 2-11 to access the CIMC web page. Then navigate to the Virtual Media page using the appropriate path below and create a virtual media device.
  - For C3260: **Navigation Button (page upper left) > Compute: Server *X* > Remote Management > Virtual Media > Add New Mapping**.
  - For C3160: **Server > Remote Presence > Virtual Media > Add New Mapping**.

Note Experience has shown that using the HTTP protocol produces the best results.

After the virtual media mapping is defined, select the CIMC **BIOS > Config Boot Order** option, define a **CIMC Mapped DVD** subtype mapping, and place it in boot position **1**.

- Use the CIMC Java KVM Console **Virtual Media** menu option to create a virtual media device associated with a COS ISO found on your desktop computer. After the virtual media mapping is defined, select the CIMC **BIOS > Config Boot Order** option, define a **KVM Mapped DVD** mapping, and place it in boot position **1** (see the **Advanced** tab for the C3260).

Note Optionally, you can set the boot device by pressing **F6** early in the C3x60 power-up sequence and selecting the desired boot device.

Step 2 Reboot the C3x60 and confirm that it boots from the COS ISO image.

Installation is automatic, and typically requires 10-15 minutes to complete from DVD (installation from virtual media is typically somewhat slower).

Note The default system baud rate is 9600. To define a different baud rate, enter **auto ks_baud_rate=<rate>** at the installation boot: prompt before the 10 second auto-installation timeout occurs.

Step 3 When prompted to reboot, disconnect the installation DVD drive or disable the virtual installation media, and then press **Enter** to reboot the C3x60.

Configure the COS Network on the C3x60

To configure the COS network on the C3x60:

Step 1 When the C3x60 finishes rebooting, log in to the console using the following credentials:

- Username – **root**
- Password – **rootroot**

The COS initialization (cosinit) script launches automatically to set up the management network that will be located on bond0.

```
Executing cosinit for platform configurations only

ATTENTION!!!

cosinit script should be run only to configure the device after an image installation.
This script modifies the network and other critical configurations based on the deployment
type. Improper use of this script may result in mis-configuring the device or making it
inaccessible.

If a new image is installed on this server, a reboot is required before running cosinit.
If a reboot is already performed, please continue. Otherwise, please exit and execute
cosinit after rebooting the server

Do you want to continue ? (yes/no) [y]: y
```

Note If you exit the cosinit script now, you can relaunch it manually from the /opt/cisco/cos/config folder.

Step 2 When asked if you want to continue, enter **y** and configure the management interface as follows (IP addresses shown below are examples only):

```
Enter management interface [bond0]:
Please ensure an IP address and netmask are configured for
management interface bond0:

Select option 1 to configure Management Interface
Select the option for 'Done' to exit this menu:
1. Configure management interface (bond0)
2. Done
Choice [2]: 1
Do you want to disable interface bond0? (yes/no) [n]: n
Enter an IP address for bond0: 10.10.10.20
Enter a netmask for bond0: 255.255.255.0
Enter a broadcast address for bond0 [10.10.10.255]:

Select option 1 to configure Management Interface
Select the option for 'Done' to exit this menu:
1. bond0    ip:10.10.10.20 mask:255.255.255.0 bcast:10.10.10.255
2. Done
Choice [2]: 2
Backing up old scripts in /etc/sysconfig/network-scripts
Writing new ifcfg-bondX scripts
```

Note If you exit the cosinit script now, you can relaunch it manually from the /opt/cisco/cos/config folder.

Step 3 Continue the script by entering the hostname of the node and the details of the default gateway as follows (IP addresses shown below are examples only):

```
Enter a hostname [localhost.localdomain]: COS-Node-1
Enter the number of the bond interface that connects to the gateway: 0
Enter the default gateway IP address [10.10.10.1]: 10.10.10.1
Backing up /etc/sysconfig/network
Writing new /etc/sysconfig/network
Backing up /etc/hosts
Writing new /etc/hosts
Restarting network services, this may take a minute:
Shutting down loopback interface:  [  OK  ]
Bringing up loopback interface:  lo: Disabled Privacy Extensions
[  OK  ]
Bringing up interface bond0:  ADDRCONF(NETDEV_UP): bond0: link is not ready
igb: bond0 NIC Link is Up 1000 Mbps Full Duplex, Flow Control: RX
ADDRCONF(NETDEV_CHANGE): bond0: link becomes ready
[  OK  ]
Network services restarted; may take a few seconds to establish connectivity
Reboot for hostname changes to take effect
Network configuration complete
```

Step 4 Ping the gateway to validate the configuration:

```
[root@localhost ~]# ping 10.10.10.1
64 bytes from 10.10.10.1: icmp_seq=1 ttl=255 time=2.58 ms
64 bytes from 10.10.10.1: icmp_seq=1 ttl=255 time=2.58 ms
64 bytes from 10.10.10.1: icmp_seq=1 ttl=255 time=2.58 ms
64 bytes from 10.10.10.1: icmp_seq=1 ttl=255 time=2.58 ms
```

Step 5 Reboot the COS node.

```
[root@localhost -]# reboot
```

Register COS Nodes to the PAM

After installing a COS image and rebooting the node, you are prompted to start configuration by running the cosinit script. This script performs initial setup and adds the node to the COS management system.

The cosinit script can be run with the following options:

- *–skipnw* – This option skips the management network configuration and registers the COS node to the COS Service Manager. You must manually input the IP pool name and select the C/F interfaces that use this IP pool.
- *–input <configFile>* – This option enables cosinit to run with a specified config file or a URL of COS nodes initialization profile.
- *–help* – This option displays the full help for the cosinit usage.

When the cosinit script executes, the COS node is registered with the COS management system. You can configure the node using the COS Service Manager web GUI, add it to a COS cluster, and make it operational. For more information, see COS Node Initial Configuration, page 2-24.

Step 1 Use the URLs generated in Configure the PAM Using the COS Configuration Wizard, page 2-10 to run cosinit on each COS node, thus adding each node to the PAM.

Example

/opt/cisco/cos/config/cosinit -nameserver 97.98.99.1 -input
https://service-mgr.cisco.com:8043/v2/cosnodeinstances/c3260

Step 2 In the COS Service Manager GUI, navigate to **PAM > Services > Cloud Object Store**.

Step 3 Click **An unused COS Service**, and then click the **Overview** tab to view the service details.

Step 4 Verify that the COS node services are running.

Create User Accounts and Verify Node and Cluster Health

To create users and accounts for a COS node, install the Swauth package on any Linux server and use the commands described below.

Note For additional details on the CLIs, see COS Command Line Utilities, page E-1 and the *Cisco Cloud Object Storage Release 3.5.1 API Guide*.

COS Swift CLI

The cos-swift command line utility is based on the Swift API, and can be used to manage storage accounts, storage containers, and storage objects in a COS cluster.

Syntax

cos-swift [–t <auth-token>] [–a <storage-url>] [–v/– –verbose][–h/– –help <subcommand>] [subcommand <options>]

COS Swauth CLI

The cos-swauth command line utility is based on the Swauth API, and can be used to manage authentication accounts and users in a COS cluster.

COS Installation on the CDE465

Note For information about installing the CDE465 platform, see the *Cisco Content Delivery Engine 465 Hardware Installation Guide.*

When installing the COS software on the CDE645 platform for the first time, recommended practice is to use the Supermicro Intelligent Platform Management Interface (IPMI) as described below.

Installation Using IPMI

Note This procedure requires that IPMI first be enabled on the CDE465. If necessary, enable IPMI on the CDE465 by enter the BIOS, selecting IPMI management, and then rebooting the CDE465.

To install COS using IPMI:

Step 1 Log in to the CDE IPMI as follows:

a. Enter the IPMI management IP address in the browser.

b. Enter the default Username **admin** and Password **admin**.

c. Click **login**.

Step 2 Launch the KVM console as follows:

a. Choose **Remote Control > Console Redirection**.

b. Click **Launch Console**.

Step 3 In the KVM console:

a. Select **Virtual Media > Virtual Storage**.

b. Insert the COS ISO file into **Device1**.

c. Click **Plug in**.

Step 4 Cycle power to the CDE465, enter the BIOS, set the machine to boot from **Device1: USB Drive**, and then reboot the CDE465.

After rebooting, the CDE465 starts COS installation automatically without manual intervention.

Step 5 When prompted at the end of installation, reboot the CDE465 and remove the ISO file from Device1.

Configuring the COS Network on the CDE465

After the CDE465 has rebooted, configure the COS network on the CDE465 as follows:

Step 1 Log in to the console using the following credentials:

- Username – **root**
- Password **rootroot**

The COS initialization script (cosinit) launches automatically to set up the management network that will be located on bond0.

```
Executing cosinit for platform configurations only

ATTENTION!!!

cosinit script should be run only to configure the device after an image installation.
This script modifies the network and other critical configurations based on the deployment
type. Improper use of this script may result in mis-configuring the device or making it
inaccessible.

If a new image is installed on this server, a reboot is required before running cosinit.
If a reboot is already performed, please continue. Otherwise, please exit and execute
cosinit after rebooting the server

Do you want to continue ? (yes/no) [y]: y
```

Note If you exit the cosinit script now, you can relaunch it manually from the /opt/cisco/cos/config folder.

Step 2 When asked if you want to continue, enter **y** and configure the management interface as follows (IP addresses shown below are examples only):

```
Enter management interface [bond0]:
Please ensure an IP address and netmask are configured for
management interface bond0:

Select option 1 to configure Management Interface
Select the option for 'Done' to exit this menu:
1. Configure management interface (bond0)
2. Done
Choice [2]: 1
Do you want to disable interface bond0? (yes/no) [n]: n
Enter an IP address for bond0: 10.10.10.20
Enter a netmask for bond0: 255.255.255.0
Enter a broadcast address for bond0 [10.10.10.255]:

Select option 1 to configure Management Interface
Select the option for 'Done' to exit this menu:
1. bond0 ip:10.10.10.20 mask:255.255.255.0 bcast:10.10.10.255
2. Done
Choice [2]: 2
Backing up old scripts in /etc/sysconfig/network-scripts
Writing new ifcfg-bondX scripts
```

Note If you exit the cosinit script now, you can relaunch it manually from the /opt/cisco/cos/config folder.

Step 3 Continue the script by entering the hostname of the node and the details of the default gateway as follows (IP addresses shown below are examples only):

```
Enter a hostname [localhost.localdomain]: CDE465-1
Enter the number of the bond interface that connects to the gateway: 0
Enter the default gateway IP address [10.10.10.1]: 10.10.10.1
Backing up /etc/sysconfig/network
Writing new /etc/sysconfig/network
Backing up /etc/hosts
Writing new /etc/hosts
Restarting network services, this may take a minute:
Shutting down loopback interface: [ OK ]
Bringing up loopback interface: lo: Disabled Privacy Extensions
[ OK ]
```

```
Bringing up interface bond0: ADDRCONF(NETDEV_UP): bond0: link is not ready
igb: bond0 NIC Link is Up 1000 Mbps Full Duplex, Flow Control: RX
ADDRCONF(NETDEV_CHANGE): bond0: link becomes ready
[ OK ]
Network services restarted; may take a few seconds to establish connectivity
Reboot for hostname changes to take effect
Network configuration complete
```

Step 4 Ping the gateway to validate the configuration:

```
[root@localhost ~]# ping 10.10.10.1
64 bytes from 10.10.10.1: icmp_seq=1 ttl=255 time=2.58 ms
64 bytes from 10.10.10.1: icmp_seq=1 ttl=255 time=2.58 ms
64 bytes from 10.10.10.1: icmp_seq=1 ttl=255 time=2.58 ms
64 bytes from 10.10.10.1: icmp_seq=1 ttl=255 time=2.58 ms
```

Step 5 Verify that the file /etc/cosd.conf, has been updated by the cos-aic-client to uncomment the db host field (about two pages down), and to specify the IP address of the local management interface as the DB Host.

```
db host <mgmt ip>
```

Step 6 Reboot the COS node.

```
[root@localhost -]# reboot
```

Step 7 Configure IP Pools as described in Configuring IP Pools, page 2-23.

Initialize the COS Node - PAM Configuration

Step 1 Provided that cosinit has run in an earlier procedure, you can now execute cosinit to skip the network configuration and begin PAM configuration, as shown in the following example:

```
[root@CDE465-1 config]#./opt/cisco/cos/config/cosinit -skipnw

ATTENTION!!!

cosinit script should be run only to configure the device after an image installation.
This script modifies the network and other critical configurations based on the deployment
type. Improper use of this script may result in mis-configuring the device or making it
inaccessible.

If a new image is installed on this server, a reboot is required before running cosinit.
If a reboot is already performed, please continue. Otherwise, please exit and execute
cosinit after rebooting the server

Do you want to continue ? (yes/no) [y]:
```

Step 2 Specify the configuration of your hardware platform as follows:

```
Please choose your platform from the following list of valid platforms:
1. UCSC-CDE465-4R4
Choice: 1
Configuring the BAUD rate to 9600 ...
```

Note The choices cosinit offers here are only those that COS supports for the hardware platform cosinit has detected. In this example, cosinit detects CDE465 hardware and so lists the sole CDE465 configuration option. If cosinit detects C3x60 hardware, it lists different configuration options accordingly.

```
=====================================================
IP Pool Configuration
=====================================================
Enter IP Pool name (Hit 'Enter' to skip IP Pool configuration): cde465-14

Please select the interfaces you want to add to colusa-69:
[ ] 1. eth1
[ ] 2. eth2
[ ] 3. eth4
[ ] 4. eth5
5. Done

Choice [5]: 1
```

Note The IP pool name should be the same as the one created in the COS Service Manager GUI.

Step 3 Enter the appropriate IP address of the COS Service Manager PAM, as follows:

- For a single PAM, enter the IP address as the docServer address.
- For multiple PAMs in an HA environment, enter the fully qualified domain name (FQDN) of the docServer that the use defined during the PAM deployment; for example, cos88.cisco.com.

Note The docserver port is 5087.

```
Please select the interfaces you want to add to colusa-69:
    [*] 1. eth1
    [ ] 2. eth2
    [ ] 3. eth4
    [ ] 4. eth5
    5. Done

    Choice [5]: 2

    Please select the interfaces you want to add to colusa-69:
    [*] 1. eth1
    [*] 2. eth2
    [ ] 3. eth4
    [ ] 4. eth5
    5. Done

    Choice [5]: 3

    Please select the interfaces you want to add to colusa-69:
    [*] 1. eth1
    [*] 2. eth2
    [*] 3. eth4
    [ ] 4. eth5
    5. Done

    Choice [5]: 4

    Please select the interfaces you want to add to colusa-69:
    [*] 1. eth1
    [*] 2. eth2
    [*] 3. eth4
    [*] 4. eth5
    5. Done

    Choice [5]: 5
```

```
Enter IP Pool name (Hit 'Enter' to complete IP Pool configuration):
Enter Cluster Name:
```

Note Enter a cluster name at this point *only* if you plan to use automated node configuration.

- If you enter a cluster name, cosinit automatically enables the data interfaces associated with the IP pool, adds the node to the specified cluster, and updates the COS node status to InService. COS then launches automated node configuration, which starts the ntpd, cosd, cassandra, and cserver services.
- If you do not enter a cluster name, cosinit registers the COS node to the PAM with the data interfaces that were added to the IP pool, and places them in Maintenance state. You must then enable all interfaces manually by assigning the COS node to a cluster and updating the node status to InService from the PAM GUI.

Caution Use automated node configuration *only* for a fresh installation. Automated node configuration starts CServer by executing **service cserver start -C**. If used on a preconfigured COS server, this command results in the loss of all existing stored content.

```
==================================================
DocServer Configurations
==================================================
Enter DocServer IP or FQDN: 97.98.99.1
Enter DocServer Port [5087]:
Enter expected Cluster Size:
Backing up exisitng /tmp/.cosnodeinit...
/tmp/.cosnodeinit generated successfully
Shutting down monit:                                       [ OK ]
Starting monit: Monit start delay set -- pause for 120s [ OK ]
Shutting down cos-aic-client:                              [ OK ]
monit: Cannot connect to the monit daemon. Did you start it with http support?
Starting cos-aic-client:                                   [ OK ]
monit: Cannot connect to the monit daemon. Did you start it with http support?
Stopping syslog-ng:                                        [ OK ]
Starting syslog-ng:                                        [ OK ]

cosinit finished successfully. Please reboot the box!!!
```

Note For **Enter expected Cluster Size** above, enter the expected number of nodes in the cluster that this node will join. This value is used later to calculate resiliency values for distributed erasure coding (DEC).

Step 4 Reboot the node if prompted.

Note If you did not enter a cluster name, cosinit does not prompt you to reboot the CDE465 as shown in this example, as reboot is not needed.

When cosinit executes, the following services are started automatically:

- syslog-ng
- ntpd
- cosd
- cassandra

- cserver

If the cluster name is provided, the command **service cserver start -C** is executed automatically when cserver starts. In addition, execution of cosinit results in the following:

- The file /etc/sysconfig/network-scripts/ifcfg-bond0 is generated, containing the management interface configuration.
- The file /etc/sysconfig/network is generated, containing the hostname and default gateway.
- The file /etc/hosts is generated, containing the host information.
- The network is restarted.
- The file /tmp/.cosnodeinit is created, reflecting the configuration specified in cosinit.
- The file /boot/grub/grub.conf is updated, reflecting the current baud rate.
- The COS AIC client service (cos-aic-client) is started.

Step 5 Confirm that the COS AIC client service is started as follows:

```
[root@Colusa-69 config]# service cos_aicc status
cos-aic-client is running
```

The COS node is now registered with COS Service Manager.

Viewing NTP Server Information

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

To view NTP server information using the COS Service Manager GUI:

Step 1 Log in to the GUI as described in Using the COS Service Manager GUI, page B-1.

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.

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:

Step 1 Log in to the GUI as described in Using the COS Service Manager GUI, page B-1.

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.
- Domain – Domain of the DNS server.
- TSIG Algorithm – Name of the TSIG algorithm used by the DNS server.

Configuring COS from the GUI

To configure COS using the COS Service Manager GUI:

Step 1 Log in to the GUI as described in Using the COS Service Manager GUI, page B-1.

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 3.5.2 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 > IP Address Pools and Networks** to define IP pools for the COS node.

Step 6 Go to **Infrastructure > COS Cluster** and create a COS cluster.

Step 7 Go to **Services > Cloud Object Stores**, click the COS service instance, and create one COS endpoint.

Step 8 On the COS node, execute **/opt/cisco/cos/config/cosinit** and configure the IP pool name, DocServer IP or domain name (depending on PAM configuration; see note below) and Port (which is the COS Service Manager management IP). This will register the COS node with the COS Service Manager.

Note For a single PAM, the DocServer IP address is the same as the COS Service Manager IP address, so you can enter it here directly. For HA PAMs, however, you must enter the DocServer domain name here (for example, mgmt-docserver.cos.cisco.com) and then enter the NameServer IP address, to help locate the mgmt-docserver leader.

Step 9 Go to **Infrastructure > COS Nodes** and configure **Cluster Name** by choosing the created cluster name from the drop-down list.

Step 10 Ensure that the **Admin State** for the node is set to **Inservice**. Click **Save**.

Step 11 Go to **Services > Cloud Object Stores**, select the COS service instance, and then click the **Overview** tab to 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:

Step 1 Log in to the GUI as described in Using the COS Service Manager GUI, page B-1.

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. Edit the field values as appropriate.

- 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 There is no default COS endpoint when the PAM is brought up, so you must create one manually.

The following parameters are displayed for the service endpoint. Edit the fields as appropriate:

- 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 – No default cluster is loaded when the PAM is brought up, so you must create one manually.
- Auth Profile – The default value is **auth-1**.
- Asset Redundancy Policy – The resiliency method for the service endpoint. From the drop-down list, choose **redundancy-pol-cos-erasurecoding**, **redundancy-pol-cos-mirroring**, or **redundancy-pol-cos-mixed**.
- State – The state of the service endpoint. From the drop-down list, choose **Disabled** or **Enabled**.

Step 6 Click **Save** to save your changes, or click **Cancel** to discard them and start over.

Configuring IP Pools

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

Step 1 Log in to the GUI as described in Using the COS Service Manager GUI, page B-1.

Step 2 Go to **Infrastructure > IP Address Pools and Networks** and create a suitable network for the IP pools.

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 try to initialize a COS node with an IP pool having fewer IP addresses available than are needed, the initialization fails and an Event is generated. For details, see COS AIC Server Events, page 3-8.

Step 7 Click **Save** to save your changes, or click **Cancel** to discard them and start over.

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

Editing or Deleting IP Pools

To edit or delete configured IP pools using the COS Service Manager GUI:

Step 1 Log in to the GUI as described in Using the COS Service Manager GUI, page B-1.

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 as follows:

a. Dissociate any COS node(s) associated with the IP pool to be edited or deleted.

b. Link the dissociated node(s) 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** as appropriate.

Configuring a COS Cluster

When COS is installed, it creates a single default node cluster. You can configure the cluster using the COS Service Manager GUI.

To view or edit the settings of a COS node cluster from the COS Service Manager GUI:

Step 1 Log in to the GUI as described in Using the COS Service Manager GUI, page B-1.

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	Must be the same as the Authentication FQDN

Step 4 Click **Save** to save your changes, or click **Cancel** to discard them and start over.

COS Node Initial Configuration

A COS node may ship with a preinstalled image that can be upgraded when provisioning the node on a network. After installing the image and rebooting the node, you are prompted to start configuration by running the cosinit script.

When prompted:

- Enter **yes** to run cosinit immediately.
- Enter **no** to configure the node later using the command:

run /opt/cisco/cos/config/cosinit

When it runs, the cosinit script performs the initial setup and adds the node to the COS management system.

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

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

- *[–skipnw]* – Skip network configurations
- *[–input <configFile>]* – Specifies a local input file for configuration (see **Example 2** below) or the URL of a COS node initialization profile
- *[–help]* – Displays the usage

Example 1: -skipnw option

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

Running cosinit with the *-skipnw* option skips the management network configuration and registers the COS node to the COS Service Manager. You must manually input the IP pool name and select the C/F interfaces that use this IP pool.

When cosinit prompts for a cluster name, you can either skip this step or enter a cluster name that has been created on the PAM. If you enter a cluster name, cosinit helps you to enable all the C/F interfaces that are added to the IP pool, and to update COS node to **cluster** and **inservice**.

Example 2: -input option

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

Running cosinit with the *-input* option enables cosinit to run with a config file. All parameters are read from the config file.

```
InterfaceConfig : bond0 --if put "skip" here ,will skip the mgmt interface config
IPADDRESS : 97.98.99.227
NETMASK : 255.255.255.128
BROADCAST : 97.98.99.255
DefaultGateway : 97.98.99.129
HOSTNAME : Colusa-227
PartNum : UCSC-C3160-4U2 --Is important, pls make sure this field is correct type
ClusterName : cluster.ca.01 --When put to null, user need to manually add node to cluster
DocServerHost : mgmt-docserver.mos.cisco.com -- can be ip address. If ip, then no need the
“NameServer" field.
DocServerPort : 5087
NameServer : 10.74.124.202
eth1 : colusa-227
eth2 : colusa-227
eth4 : colusa-227
eth5 : colusa-227
```

Executing the cosinit script completes the following tasks:

- Configures the management interface – bondX, IP, netmask, broadcast.
- Configures the default gateway.
- Configures the hostname.
- Maps the network interfaces to IP addresses in the network.
- Configures Doc Server IP address and Port.
- Records the appliance model name.

- Configures 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-bond0 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 Service Manager GUI, add it to a COS cluster, and make it operational.

Adding a COS Node to a COS Cluster

After installing and initializing a COS node, you must add the node to a COS cluster to begin servicing COS requests.

To add a node to the COS cluster using the COS Service Manager web GUI:

Step 1 Log in to the GUI as described in Using the COS Service Manager GUI, page B-1.

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

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

Step 4 In the **Service Interfaces** list, check the box beside the interfaces to add to the IP pool, and then enable them.

Step 5 From the **Cluster Name** list, choose the created COS cluster.

Step 6 Click **Save** to save your changes, or click **Cancel** to discard them and start over.

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 installing and initializing a COS node, you can modify its configuration parameters using the COS Service Manager GUI as follows:

Step 1 Log in to the GUI as described in Using the COS Service Manager GUI, page B-1.

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 **Inservice** or under **Maintenance**

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, page 2-28.

Step 6 Click **Save** to save your changes, or click **Cancel** to discard them and start over.

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:

Step 1 Log in to the GUI as described in Using the COS Service Manager GUI, page B-1.

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 **Admin State** is set to **Maintenance.**
Enabled	From the drop-down list, choose **True** or **False** to enable or disable the service instance, respectively.

Step 6 Click **Save** to save your changes, or click **Cancel** to discard them and start over.

Automated COS Node Configuration (Optional)

To use automated configuration, you provide a file to cosinit that includes a ClusterName and IP Pool reference for at least one service interface. This enables the system to configure the COS node without manual intervention through either the COS Service Manager GUI or the API.

Caution Using automated configuration results in the deletion of all existing content from the disks in the node being configured. Do not use this option unless you intend to wipe all content from these disks.

Example

Run cosinit with the input file:

/opt/cisco/cos/config/cosinit -input cosinit_auto

```
[root@cos470-1 ~]# cat cosinit_auto
InterfaceConfig : skip
PartNum : C3160-R1
ClusterName :cluster-1
DocServerHost : 97.98.99.1
DocServerPort : 5087
eth1 :      colusa-57
eth2 :      colusa-57
eth4 :      colusa-57
eth5 :      colusa-57
```

This results in one of the following:

- If ClusterName has a value, automated configuration is triggered. Then, if at least one service interface has an IP Pool configured, the AIC Client sets the adminstate to inService in the smcosnode document before sending cosannounce to DocServer. Otherwise, adminstate is set to Maintenance.

  The AIC Server handles the rest of the configuration automatically by assigning an IP address from the specified IP Pool and adding the node to the cluster. The AIC Client then writes the necessary configuration files to the COS node as usual. When configuration is complete, the AIC Client automatically starts the cassandra, cosd, and cserver cos node services.
- If ClusterName has no value, manual configuration is needed. You must manually set the COS node adminstate from Maintenance to inService, assign the node to a cluster, and then enable the service interfaces from the GUI.

You can verify the input by viewing the /tmp/.cosnodeinit file:

```
[root@perf-4t-cos01] # vi /tmp/.cosnodeinit
Name : 171491989
Model : C3160-R1
DocServerHost : 10.56.194.152
DocServerPort : 5087
bond0 : 10.56.194.149
bond1 : v1860-pool
bond2 : v1860-pool
bond4 : v1860-pool
bond5 : v1860-pool
ClusterName :cluster1
```

Note When using automated configuration of multiple COS nodes, configure the first node, and then wait before configuring the second node until the Cassandra database service to appear as Running in the GUI of the first node configured. Otherwise, there may be unexpected behavior in the seed list configuration for the Cassandra database of the nodes added after the first node.

Verify That All Services are Running

Step 1 Verify that cassandra is running:

```
[cos-node@ root] nodetool status
Datacenter: datacenter1
=======================
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
-- Address      Load      Tokens Owns (effective)  Host ID                              Rack
UN 10.168.5.2   117.35 KB 256    100.0%         bf862009-40eb-474a-a86b-1700c724755a rack1
UN 10.168.5.18  768.4 KB  256    100.0%         39e62fbf-d0fb-4a92-9eb9-869b488ed5af rack1
```

Step 2 Verify that cosd is functional:

```
[cos-node@ root] curl –v http://fqdn/info

{"cluster":{"config_ver":2,"fqdn":"cos-utah50.lindon.lab.cisco.com","name":"local","enable
_wos":true},"swauth":{"reseller_prefix":"AUTH_","path_prefix":"auth/","max_key_len":256,"m
ax_user_len":256,"token_life":86400,"max_token_life":86400,"max_account_len":256},"swift":
{"version":"2.2.0","max_account_len":256,"max_container_len":256,"max_object_len":1024,"ma
x_container_list":10000,"max_object_list":10000},"log":{"default":"notice"},"rio":{"path_p
refix":"rio/"}}
```

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/

To execute periodic repair, we recommend configuring a CRON job on each node, and then 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