Cisco Smart PHY Application Install Guide, Release 22.1

First Published: 2022-03-23

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

CHAPTER 1

Cisco Smart PHY Deployment Overview

This guide provides information about the deployment of Cisco Smart PHY application in offline environments without the Internet connectivity.

Offline Deployment

Cisco Smart PHY software image is a compressed tarball file that contains all the scripts, helm charts, and container images required for installing the Cisco Operations Hub cluster and the Cisco Smart PHY application. It also contains a copy of these instructions and configuration examples.

Use the `deployer` script available in the software image to set up the Deployer virtual machine (VM) and clusters.

The installation process creates the following components:

- Deployer: The controller used to configure and deploy the cluster.
- Cisco Operations Hub Cluster: The Cisco Smart PHY application runs on this cluster.

The Deployer VM supports two types of Cisco Smart PHY deployments:

- All-in-one (AIO) cluster: Runs as a single VM.
- Multinode cluster: Consists of 12 VMs deployed across three VMware ESXi hosts. Each VMware ESXi host server hosts a control-plane VM, etcd VM, Infra VM, and Operations VM.

  The multinode cluster provides complete support for high availability (HA) and supports two sizes of clusters: small and normal for production environments. The default size is small, when the `size` is not specified.

  For the current number of RPDs supported on the Cisco Smart PHY application, we recommend the small multinode cluster for production deployments due to their increased resiliency.

The following table shows the minimum resources for deployer and AIO.

VM Type	CPU Cores	RAM Size (GB)	SSD Storage Size (GB)
Deployer	8	16	320

VM Type	CPU Cores	RAM Size (GB)	SSD Storage Size (GB)
All-in-one	18	96	1541

The following table shows the minimum resources for each type of VM when the cluster size is set to multinode `small`:

VM Type	CPU Cores	RAM Size (GB)	SSD Storage Size (GB)
Control Plane	2	16	125
etcd	2	16	125
infra	8	64	1000
ops	8	64	320

The following table shows the minimum resources for each type of VM when the cluster size is set to multinode `normal`:

VM Type	CPU Cores	RAM Size (GB)	SSD Storage Size (GB)
Control Plane	2	16	125
etcd	2	16	125
infra	14	96	1500
ops	16	176	620

Prerequisites for Cluster Deployment

The following prerequisite components are required to install, operate, and manage the Cisco Smart PHY application running on a Cisco Operations Hub cluster.

- Staging server: Physical or virtual machine to run the installation script
- An ESXi host to run the Deployer VM.
- VMware vSphere virtualization platform

Prerequisites for ESXi Hosts

Three ESXi hosts are required to run a Cisco Operations Hub multinode cluster.

The minimum compute, storage, and networking requirements for the VMware ESXi host are listed in the following table:

Component	Small Size Cluster	Normal Size Cluster
Processor	20 vCPUs	34 vCPUs
Memory	160 GB	304 GB

Component	Small Size Cluster	Normal Size Cluster
Storage	1640 GB Minimum 50000 IOPS (Input/output operations per second) Latency of < 5 ms	2440 GB Minimum 50000 IOPS (Input/output operations per second) Latency of < 5 ms
NIC	2x 10G vNIC	2x 10G vNIC

Prerequisites for VMware vSphere

VMware ESXi and VMware Center Server are mandatory components of a Cisco Smart PHY cluster deployment.

- Hypervisor: VMware ESXi 7.0
- Host Management: VMware vCenter Server 7.0

Deploy Cisco Smart PHY

Deploying Cisco Smart PHY in an offline environment involves the following process.

1. (Optional) Configure UCS server: Not required if you are deploying using third party servers. For details, see Configure UCS Servers, on page 25.
2. Prepare a staging server.
3. Prepare a cluster configuration file.
4. Deploy the cluster.

If required, repeat the step 3 and 4 to deploy another cluster.

CHAPTER 2

Prepare Staging Server

The staging server may be a physical server, a virtual machine, or even a laptop. The staging server must be connected to the target VMware vSphere Infrastructure, vCenter Server, and cluster nodes with correct credentials.

Prerequisites

The staging server requires the following software:

- docker 18.09.7 or later
- python 3.6 or later

Unpack Cisco Smart PHY Application Package

The Cisco Smart PHY software image is a compressed tarball file that is self-sufficient for installing the Deployer, Cisco Operations Hub cluster, and Cisco Smart PHY application. It contains the following files:

- Installation script
- All relevant product images
- Sample configuration files
- README file

Before you begin

Make sure that you have a minimum of 50 G disk space to extract the image.

Step 1 Unpack the signed TAR software image of the Cisco Smart PHY application:

```
smartphy-installer-<version>.SSA.tgz
```

The file is approximately 10 G.

After downloading the image, extract all individual files, and verify the signature of the files using the following steps.

Step 2 Run the following command to extract the TAR file: `tar -zxovf smartphy-installer-<version>.SSA.tgz`

This command extracts the following files:

- `cs-verify.sh`
- `SMART_PHY_REL_KEY-CCO_RELEASE.cer`
- `image.tgz`
- `image.tgz.signature`
- `signed_files`

Step 3 Run the following command to extract all individual files of the cluster, Operations Hub, and Cisco Smart PHY:

```
tar -zxovf smartphy-installer-<version>.tgz
```

Example:

The `smartphy-installer-<version>.SSA.tgz` file is extracted to the `smartphy-installer-<version>` directory.

Step 4 Change the directory to `smartphy-installer-<version>` directory.

```
cd smartphy-installer-<version>
```

The new staging directory `smartphy-installer-<version>` has the following content:

```
$ tree -a
.
├── README.md
├── cluster-deployer-<version>.tar
├── cluster-deployer-<version>.tar.signature
├── deploy
├── deploy.signature
├── docker-images
│   ├── ccmts-customization_<version>.tar
│   └── ccmts-customization_<version>.tar.signature
├── examples
│   ├── aio-smartphy-config.yaml
│   ├── aio-smartphy-standby-config.yaml
│   ├── deployer-sample-config.yaml
│   ├── multinode-smartphy-config.yaml
│   └── multinode-smartphy-standby-config.yaml
├── offline-products
│   ├── cee-<versioin>.tar
│   ├── cee-<versioin>.tar.signature
│   ├── opshub.tar
│   ├── opshub.tar.signature
│   ├── smartphy-<version>.tar.signature
│   └── smartphy-<version>.tar
├── smi-install-disk.iso
├── smi-install-disk.iso.signature
├── upgrade-prep
├── upgrade-prep.signature
└── utility-images
    ├── autodeploy_<version>.tar
    ├── autodeploy_<version>.tar.signature
    ├── cluster-manager-docker-deployer_<version>.tar
    └── cluster-manager-docker-deployer_<version>.tar.signature
```

This directory is referred to as the staging directory in this document.

Step 5 Run the `cs-verify.sh` script.

Example:

```
./cs-verify.sh SMART_PHY_REL_KEY-CCO_RELEASE.cer smartphy-installer-<version>.tgz
```

The following messages appear:

```
Verifying signature

Signature verification succeeded
```

If the signature verification fails, error messages appear on the screen. If error messages appear, download the software package once again.

CHAPTER 3

Prepare Cluster Configuration File

Prepare the cluster configuration file by completing the following tasks:

1. Get details of VMware vCenter.
2. Get IP Addresses for deployer and cluster.
3. Add VMware vCenter environment configuration.
4. Add deployer configuration.
5. Add cluster configuration.
6. Add Cisco Smart PHY CIN configuration.

VMware vCenter Details

To contact the VMware vCenter server, the `deployer` script and the deployer VM require the following details:

- Server name or IP address
- Username and password
- Datacenter and cluster name
- Host server and datastore names

  The deployer and the single-node cluster require one host server and the multinode cluster requires three host servers.

IP Addresses for Deployer and Cluster

Deploying the Cisco Smart PHY software offline requires the following IP addresses:

- One management IP address for the deployer
- Management IP addresses for cluster (1 for single-node, 12 for multi-node cluster)
- CIN network IP addresses for Cisco Smart PHY (1 per CIN interfaces per Operations VM)
- One virtual IP address for management network and one for each CIN network (multi-node cluster)

Sample Configuration Files

The `examples` directory contains sample configuration files for automatic deployment:

- `deploy-sample-config.yaml`: Configuration file with only the deployer.
- `aio-smartphy-config.yaml`: Configuration file with the deployer and the single-node Smart PHY cluster.
- `multinode-smartphy-config.yaml`: Configuration file with deployer and multinode Cisco Smart PHY cluster.
- `aio-smartphy-standby-config.yaml`: Configuration file with deployer and single-node Cisco Smart PHY cluster that is used for standby (without CIN config).
- `multinode-smartphy-standby-config.yaml`: Configuration file with deployer and multinode Cisco Smart PHY cluster that is used as a standby (without CIN config).

Cluster Configuration File

Place the configuration file in the staging directory. This configuration file is in the standard YAML language format, with the following three sections:

- Environments
- Deployers
- Clusters (Smart PHY multi-node/single-node)

Each section can contain multiple items. Replace <...> with actual values.

Environment Configuration

This section provides details of the VMware vCenter access and network access for creating and provisioning the deployers and cluster virtual machines.

```
environments:
  <environment name>:
      server: <vCenter name or IP address>
      username: <vCenter user name>
      datacenter: <vCenter datacenter name>
      cluster: <vCenter cluster name>
      nics: [ <list of vCenter networks> ]
      nameservers: [ <list of DNS servers> ]
      search-domains: [ <list of search domains> ]
      ntp: [ <list of ntp server names or IP addresses> ]
      https-proxy: <HTTP proxy server>
      no-proxy: <list of domains not using proxy>
```

Guidelines for configuring the VMware vCenter environment:

- The environment name can have only lowercase letters, digits, and hyphens (-).
- The NIC's list must have only one network, although the NIC configuration allows multiple networks. This network is used as the management network in the deployer or cluster that refers to this environment.
- Configure multiple environments for this vCenter if your vCenter has more than one network that serves as a management network. One for each network. In addition, refer to the corresponding environment in the deployer or cluster based on the management network it uses.
- Make sure the NIC's name-servers and search-domains fields are configured as lists.

Deployer Configuration

Before creating and deploying a deployer VM, define a minimum of one environment.

```
deployers:
  <deployer name>:
      environment: <environment of vCenter hosting the deployer>
      address: <deployer VM IP address in CIDR format>
      gateway: <gateway IP address>
      username: <user name for deployer>
      # SSH private-key-file with path relative to the staging directory
      # If the line is missing, ssh private key will be auto-generated and saved inside
.sec/
      private-key-file: <path and filename for ssh private key>
      host: <ESXi host IP address>
      # datastore-folder parameter is optional, it accepts subfolder structure too. If you
 don't specify the folder, cluster and deployer are created under root folder of VM datastore.

      datastore-folder: "my-folder/cluster"
      datastore: <vCenter datastore name for host>
      # ingress-hostname only supports '.' and alphanumeric characters
      ingress-hostname: "deployer.example.com"
      # Use docker-subnet-override to override default docker-subnet address in deployer
VM. This is an optional parameter. If you do not specify it, 172.17.0.0/16 subnet range
will be assigned by default. Use this option if the default docker IP subnet (172.17.0.0/16)
 of the deployer VM may overlap with IP networks in your environment.
      docker-subnet-override:
        # pool name of docker bridge address pool, like pool1, pool2 etc..
        - pool-name: pool1
          # docker bridge subnet range in CIDR format.
          base: 172.XX.X.X/16
          # subnet range is 8-24.
          size: 24
```

Guidelines for configuring the deployer VM:

- The name of the deployer VM can have only lowercase letters, digits, and hyphens (-).

- The deployer supports FQDN. Use the field `ingress-hostname` to configure FQDN.
- If you specify `ingress-hostname` (for example `deployer.example.com`), add the FQDN entries to DNS. Use the following FQDNs:
    - `charts.deployer.example.com`
    - `docker.deployer.example.com`
    - `deployer.example.com`
    - `files-offline.smi-deployer.deployer.example.com`
    - `cli.smi-deployer.deployer.example.com`
    - `restconf.smi-deployer.example.com`
- The `private-key-file` field, when present, must refer to the SSH private key file. This file must be in the staging directory and must not be accessible (read/write/execute) to other users.

  If the `private-key-file` line is missing, the `deployer` script generates an SSH private key for the deployer and places it in the `.sec` subdirectory under the staging directory. The filename is `<deployer-name>_auto.pem`.
- To avoid a resource contention, do not run the deployer VM on the same ESXi hosts running any of the Cisco Smart PHY cluster's 12 VMs.

Cluster Configuration

Before creating and deploying a cluster, configure one environment and one deployer. A cluster has an environment field to reference to its corresponding environment.

```
clusters:
 <SMI cluster name>
   type: "opshub"
   # optional cluster `size` field. Support `small` or `normal`.Default value is `small`
if not specified.
   size: small
   environment: <environment of vCenter hosting the SMI cluster>
   gateway: <gateway IP address>
   username: <user name for the SMI cluster>
   # SSH private-key-file with path relative to the staging directory
   # If the line is missing, ssh private key will be auto-generated and saved inside .sec/
   private-key-file: <path and filename for ssh private key>
   # The following two fields are for multi-node cluster only
   primary-vip: <virtual IP address for the management network in CIDR format>
   vrouter-id: <VRRP ID for the management network>
   # ingress-hostname only supports '.' and alphanumeric characters
   ingress-hostname: "smartphy.example.com"
   pod-subnet: <IP address range for kubernetes pod in CIDR format "192.168.120.0/24">
   #pod-subnet is an optional field. If you do not specify the IP address, "192.168.0.0/16"
 will be assigned by default.
  service-subnet: <IP address range for kubernetes service in CIDR format "10.96.120.0/24">

   # service-subnet is an optional field. If you do not specify the IP address, "10.96.0.0/12"
 will be assigned by default.
   docker-bridge-subnet: ["IP address range for docker bridge in CIDR format
'10.96.120.0/24'"]
```

```
   # docker-bridge-subnet is an optional field. If you do not specify the IP address,
"172.17.0.0/16" will be assigned by default.
   # For Multi-Node cluster only
   nodes:
    - host: <ESXi host 1 IP address>
     addresses: [ <CONTROL-PLANE 1 IP>, <ETCD 1 IP>, <INTRA 1 IP>, <OPS 1 IP> ]
     datastore: <vCenter datastore for host 1>
     # datastore-folder parameter is optional, it accepts subfolder structure too. If you
don't specify the folder, cluster and deployer are created under root folder of VM datastore.

     datastore-folder: "my-folder/cluster"
    - host: <ESXi host 2 IP address>
      addresses: [ <CONTROL-PLANE 2 IP>, <ETCD 2 IP>, <INTRA 2 IP>, <OPS 2 IP> ]
      datastore: <vCenter datastore for host 2>
      # datastore-folder parameter is optional, it accepts subfolder structure too. If you
 don't specify the folder, cluster and deployer are created under root folder of VM datastore.

      datastore-folder: "my-folder/cluster"
    - host: <ESXi host 3 IP address>
      addresses: [ <CONTROL-PLANE 3 IP>, <ETCD 3 IP>, <INTRA 3 IP>, <OPS 3 IP> ]
      datastore: <vCenter datastore for host 3>
      # datastore-folder parameter is optional, it accepts subfolder structure too. If you
 don't specify the folder, cluster and deployer are created under root folder of VM datastore.

      datastore-folder: "my-folder/cluster"
   apps:
      - smartphy:
          nodes:
            - host:  <ESXi host 1 IP address>
              nics: <vCenter network for CIN>
              ops:
                interfaces:
                  - addresses: [ <OPS 1 IP> ]
                    vip: [ <LIST of virtual IP for CIN network in CIDR format> ]
                    vrouter-id: <VRRP ID for CIN network>
                    routes:
                      - { dest: [ <LIST of destination subnets> ], nhop: <next hop IP> }
                      - { dest: [ <LIST of destination subnets> ], nhop: <next hop IP> }
            - host:  <ESXi host 2 IP address>
              nics: <vCenter network for CIN>
              ops:
                interfaces:
                  - addresses: [ <OPS 2 IP> ]
                    vip: [ <LIST of virtual IP for CIN network in CIDR format> ]
                    vrouter-id: <VRRP ID for CIN network>
                    routes:
                      - { dest: [ <LIST of destination subnets> ], nhop: <next hop IP> }
                      - { dest: [ <LIST of destination subnets> ], nhop: <next hop IP> }
            - host:  <ESXi host 3 IP address>
              nics: <vCenter network for CIN>
              ops:
                interfaces:
                  - addresses: [ <OPS 3 IP> ]
                    vip: [ <LIST of virtual IP for CIN network in CIDR format> ]
                    vrouter-id: <VRRP ID for CIN network>
                    routes:
                      - { dest: [ <LIST of destination subnets> ], nhop: <next hop IP> }
                      - { dest: [ <LIST of destination subnets> ], nhop: <next hop IP> }

   # For Single-Node cluster only

clusters:
 "cicd-smi-aio":
   type: "opshub"
   environment: "cicd-vcenter"
   username: "build"
   # private-key-file must exist in the path of staging/install directory.
   # file path is relative to the staging/install directory.
   private-key-file: "cmts.pem"
   # pod-subnet is an optional field. If you do not specify the IP address, "192.168.0.0/16"
 will be assigned by default.
   pod-subnet: "192.168.120.0/24"
   # service-subnet is an optional field. If you do not specify the IP address, "10.96.0.0/12"
 will be assigned by default.
   service-subnet: "10.96.120.0/24"
   # docker-bridge-subnet is an optional field. If you do not specify the IP address,
"172.17.0.0/16" will be assigned by default.
   docker-bridge-subnet: ["172.17.120.0/24"]
   gateway: "172.22.80.1"
   ingress-hostname: smartphy.example.com
   nodes:
   - host: <ESXi host IP address>
     addresses: [ <AIO VM IP address> ]
     datastore: <vCenter datastore for host>
     # datastore-folder parameter is optional, it accepts subfolder structure too. If you
don't specify the folder, cluster and deployer are created under root folder of VM datastore.

     datastore-folder: "my-folder/cluster"
   apps:
    - smartphy:
        nodes:
          - host: <ESXi host IP address>
            nics: [<vCenter network for CIN>]
            control-plane:
               interfaces:
                 - addresses: [ <LIST of IP addresses for CIN network in CIDR format> ]
                   routes:
                   - { dest: [ <LIST of destination subnets> ], nhop: <next hop IP> }
                   - { dest: [ <LIST of destination subnets> ], nhop: <next hop IP> }
```

Command	Description
<cluster name>	Cluster name.
type	Use `opshub` for Cisco Smart PHY cluster.
size	Small or normal. When the size is not specified, the default value is small.
environment	Reference to vCenter environment.
gateway	Gateway for the cluster nodes.
username	Username of the cluster.
private-key-file	SSH private-key-file with the path relative to the staging directory. If the line is missing, the SSH private key will be auto-generated and saved inside `.sec/`.
primary-vip	Primary virtual IP address in CIDR format (multinode only).
vrouter-id	VRRP ID for management network (multinode only).

Command	Description
ingress-hostname	Fully Qualified Domain Name (FQDN) assigned to the cluster. Only alphanumeric characters and period (.) are allowed. Your authoritative DNS server must be configured to resolve the specified FQDN and the following subdomain: • `opscenter.<fqdn>` Alternatively, if your authoritative DNS server supports wildcards, you must configure the DNS to resolve the specified FQDN and a wildcard record covering the subdomains listed here. If you do not specify an FQDN: • The cluster IP address is used to generate an FQDN leveraging `nip.io` as the domain and top-level domain (TLD). For example, if the IP address of the cluster is 10.0.0.2, the generated FQDN is `10.0.0.2.nip.io`. The subdomains listed here are also leveraged. • Your DNS servers must allow the resolution of the `nip.io` domain. If resolution of `nip.io` is blocked, you cannot access the cluster.
host	ESXi IP address where VMs are hosted.
service-subnet	Service subnet range to configure k8s and calico in CIDR format. The default value is 10.96.0.0/12.
pod-subnet	Pod subnet to configure k8s and calico in CIDR format. The default value is 192.168.0.0/16.
docker-bridge-subnet	IP address range for docker bridge in CIDR format. It is an optional field. If you do not specify the IP address, 172.17.0.0/16 will be assigned by default.
apps	Application to be installed on top of the platform, in this case `smartphy`.
addresses	IP addresses assigned to control-plane, etcd, infra and docsis or operations nodes respectively.
CIN Configuration	
vip	Virtual IP address in CIDR format.
vrouter-id	VRRP ID for CIN.
addresses	CIN IP addresses in CIDR format.
nics	vCenter NICs for CIN.
For Single-Node cluster	
host	ESXi IP address where VM is hosted.
control-plane	Cisco Smart PHY CIN configuration.

Guidelines for configuring a cluster:

- The name of the cluster can have only lowercase letters, digits, and hyphens (-).
- The `private-key-file` field, when present, should refer to the SSH private key file. This file must be in the staging directory and must not be accessible (read/write/execute) to other users.

  If the private-key-file line is missing, the `deployer` script generates an SSH private key for the cluster and places it in the `.sec` subdirectory under the staging directory. The filename is `<cluster-name>_auto.pem`.
- Configure the virtual IP address of the Smart PHY cluster and VRRP ID (`vrouter-id` at cluster level) for the management network for multinode clusters. The management network supports only IPv4. The `vrouter-id` can take values 1–254.
- If multiple clusters share the same management subnet, the VRRP ID for each cluster must be unique in the management subnet.

Note If the cluster ingress hostname is `smartphy.example.com`, you can access opscenter CLI and RESTCONF endpoints as follows:

- Use `opscenter.smartphy.example.com/cee-data/restconf` to access cee-data opscenter RESTCONF endpoint.
- Use `opscenter.smartphy.example.com/cee-data/cli` to access cee-data opscenter CLI.
- Use `opscenter.smartphy.example.com/opshub-data/restconf` to access opshub-data opscenter RESTCONF endpoint.
- Use `opscenter.smartphy.example.com/opshub-data/cli` to access opshub-data opscenter CLI.
- Use `opscenter.smartphy.example.com/smartphy-data/restconf` to access smartphy-data opscenter RESTCONF endpoint.
- Use `opscenter.smartphy.example.com/smartphy-data/cli` to access smartphy-data opscenter CLI.

Configure TLS Certificate

When Cisco Operations Hub cluster is deployed, a self-signed certificate is configured by default. You can replace the self-signed certificate with a CA signed certificate through the Deployer CLI. Use the following commands as example to configure a CA signed TLS certificate.

```
product opshub# config terminal
Entering configuration mode terminal
product example deployer(config)# clusters {k8s-cluster-name}
product example deployer(config-clusters-******)# secrets tls opshub-data cert-api-ingress
 ?
Possible completions:
  certificate   Path to PEM encoded public key certificate.
  private-key   Private key associated with given certificate.
  <cr>
product example deployer(config-clusters-******)# secrets tls nginx-ingress
default-ssl-certificate ?
Possible completions:
  certificate   Path to PEM encoded public key certificate.
  private-key   Private key associated with given certificate.
  <cr>
product example deployer(config-clusters-******)#commit
product example deployer(config-clusters-******)#exit

product example deployer(config-clusters-******)# cluster <cluster-name> actions sync run
force-vm-redeploy false
```

Cisco Smart PHY CIN Configuration

Configure Converged Interconnect Network (CIN) for the Cisco Smart PHY cluster. One or more CIN networks can be present. Configure CIN under each node.

Guidelines for configuring CIN:

- CIN should contain the network names (NICs) and the IP addresses (addresses).
- The routing table (routes) is optional.
- The virtual IP addresses (`vip`) and the VRRP ID (`vrouter-id`) fields are used only in multinode clusters. They are configured on the first node.
- The virtual IP addresses are mandatory. You can configure up to one IPv4 and one IPv6 addresses per CIN network.
- If multiple Smart PHY clusters share a CIN subnet, the VRRP ID should be unique for each cluster.
- For multinode cluster, all nodes must have the same number of CIN interfaces. If the NICs or route fields are missing for the second or third nodes, the corresponding value from the first node is used.
- You can also set up a Smart PHY cluster as backup cluster. For backup clusters, do not include any CIN configuration. The configuration should not have operations and interfaces under the nodes.

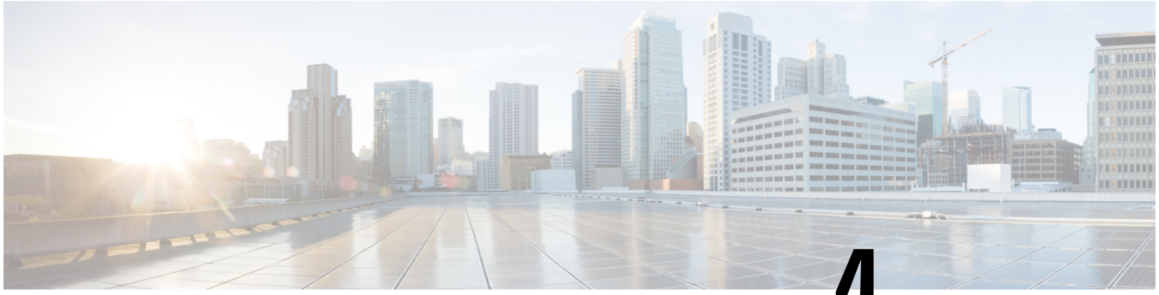

CHAPTER 4

Deploy the Deployer VM, Cisco Operations Hub, and Cisco Smart PHY Application

This section explains how to use the deployer script to deploy the deployer virtual machine (VM), Cisco Operations Hub, and the Cisco Smart PHY application.

Deploy the Deployer VM, Cisco Operations Hub, and Cisco Smart PHY Application

Deploy the Deployer

From the staging server, run the `deployer` script to deploy the clusters using the following command:

```
$ ./deploy
Usage ./deploy -c <config_file> [-v]
  -c <config_file> : Configuration File, <Mandatory Argument>
  -v               : Config Validation Flag, [Optional]
  -f               : Day0: Force VM Redeploy Flag [Optional]
                   : Day1: Force RPD Update Flag [Optional]
  -u               : Cluster chart Upgrade Flag [Optional]
  -s               : Skip Compare Flag [Optional]
  -sc              : Skip Compatibility check during upgrade Flag [Optional]
  -D               : Enable Debug Logs [Optional]
```

The following options are available in the `deployer` script:

- `-c <config_file>`: Configuration file (Mandatory Argument). This option is the first option in the command.
- `-u`: Cluster chart Update Flag [Optional]
- `-v`: Config Validation Flag, [Optional]
- `-f`: Redeploy the cluster. If you redeploy the cluster, cluster VM's will be rebooted and the data persisted on disk will be retained. You can use this option to modify some of the cluster parameters.

The `-u` flag is for updating CNF/charts in cluster.

The `deployer` script triggers the docker command that requires root permission to run. Depending on your setting, you can use the **sudo** to the deploy command.

The `deployer` script does the following operations:

- If you are running the `deployer` script for the first time, it prompts you to enter all passwords required for installation.

  - For vCenter environment: vCenter password for the user specified in the environment configuration.
  - For deployer: SSH password of the user admin for the deployer's Operation Center.
  - For Cisco Smart PHY cluster: SSH password for all VMs in the cluster (or user-specified in the cluster's configuration file). Also, the SSH passwords for the three Operation Centers (Cisco Smart PHY, Operations Hub, and CEE); for user admin.

  You are prompted twice to enter each password. The password is saved inside the staging directory in encrypted form for future use.

- Passwords for the deployer, the cluster, and the Operation Centers must be eight characters long, and must have a lowercase letter, uppercase letter, a digit, and a special character.
- The `deployer` script generates an SSH key pair when the `private-key-file` line is missing for the deployer or the cluster in the configuration file. The generated private key files are in the `.sec` sub directory under the staging directory, with `<cluster-name>_auto.pem` filename.
- The root user owns the generated private keys. When logging in using SSH and these private key files, make sure that you run it with `sudo`.
- If the deployer VM is not running, the `deployer` script installs the deployer VM.
- The `deployer` script checks if the deployer VM is missing any of the product packages that are found in the `offline-images` directory, and if it finds any missing, it uploads them to the deployer VM.
- The script also generates the configuration for each cluster and pushes them to the deployer VM.
- The `deployer` script triggers the deployer VM to perform the sync operation for the cluster. The sync operation applies the configuration to the cluster. If you have not set up the cluster, it installs the cluster. Or the sync operation updates the cluster with the configuration.
- If the sync operation times out, the `deployer` script triggers the sync operation again. The script waits for the sync operation to complete, and then continues to monitor the cluster to make sure that all helm charts are deployed and all pods are created.

You can repeat the `deployer` script to deploy more than one cluster by providing the corresponding configuration files. Alternatively, you can run this command appending a `-v` flag. The `-v` flag forces the `deployer` script to skip the synchronizing operation. Use this option to push the configuration of a cluster to the deployer without deploying or updating the cluster.

Wait for the installation process to complete. Following is a sample output after the process is complete:

```
Friday 22 October 2021  07:53:52 +0000 (0:00:00.123)       0:12:22.518 ********
install-cm-offline : Extract cluster manager file into /data ---------- 545.16s
vm-vsphere-iso : Wait for ssh ------------------------------------------ 88.51s
install-cm-offline : Deploy cluster manager ---------------------------- 85.14s
install-ntp-iso : force_time_sync --------------------------------------- 7.34s
vm-vsphere-iso : Create VM ---------------------------------------------- 3.85s
vm-vsphere-iso : Get VM Update needed ----------------------------------- 1.65s
install-ntp-iso : Cleaning cache ---------------------------------------- 1.53s
```

```
Gathering Facts --------------------------------------------------------- 1.34s
vm-vsphere-iso : Check if ISO file exists ------------------------------- 0.79s
vm-vsphere-iso : Test vCenter credentials are valid --------------------- 0.60s
install-ntp-iso : apt_update -------------------------------------------- 0.55s
vm-vsphere-iso : Create user data ISO ----------------------------------- 0.52s
install-ntp-iso : Remove "ntp" package ---------------------------------- 0.47s
install-cm-offline : Ensure /data/cm-install folder NOT exists ---------- 0.36s
install-ntp-iso : Install offline APT repo GPG key ---------------------- 0.34s
install-cm-offline : Ensure /data folder exists ------------------------- 0.33s
install-ntp-iso : restart_chrony ---------------------------------------- 0.28s
install-ntp-iso : enable chrony ntp ------------------------------------- 0.28s
download-iso : download base image ISO file ----------------------------- 0.28s
vm-vsphere-iso : Create netplan Template -------------------------------- 0.18s

Create deployers completed
```

Deploy the Cluster with CA signed certificate using deploy command

When you deploy the Cisco SmartPHY cluster, the cluster is configured with a self-signed certificate by default. You can deploy the cluster with a CA signed certificate by performing the following steps before running deploy script.

1. Generate a CA signed certificate with a common name as `ingress-hostname` used in the day 0 configuration YAML file.
2. On the stanging server, create a directory with the cluster name as the directory name under `<staging directory>/certs/client_certificates`. For example, if you use cluster name `testcluster`, the created directory will be `<staging directory>/certs/client_certificates/testcluster`. This directory is called **cluster ingress certificates directory**.
3. Create `cert-api-ingress` and `default-ssl-certificate` directories under **cluster ingress certificates directory**.
4. Place the CA Signed certificate and keys under `cert-api-ingress` directory. The CA signed certificate file has `.crt` extension and key file has `.key` extension.

Deploy the Cluster

Run the following `sync` command to deploy a new cluster or to update an existing cluster.

```
clusters <cluster> actions sync run
```

Enter `yes` at the prompt to start the deployment as a background synchronization job.

Important The sync command does not support updating network or node configurations. For such changes, redeploy the cluster.

Redeploy Cisco Operations Hub Cluster

To remove and redeploy a cluster, run the following command:

```
clusters <cluster> actions sync run force-vm-redeploy true purge-data-disks true
```

This command removes the VMs of the cluster and its data disks, before deploying the cluster.

Note Back up the configuration data before redeploying the cluster. Configuration data of the Cisco Operations Hub cluster is deleted after the process.

Verify Installation

After successfully deploying the Cisco Smart PHY application using the deployer script, the console shows a success message.

Log in to one of the control-plan nodes and make sure that all the pods are in the `Running` state.

```
kubectl get pod --all-namespaces
```

A few internal services and pods may need more time to complete the startup tasks and successfully establish communication with other services within the cluster. After a few minutes, you can initiate all operations from the Cisco Smart PHY web UI page.

CHAPTER 5

Troubleshoot Cisco Smart PHY Installation

This section provides tips that would help troubleshoot issues with the installation.

Access the Deployer

You can access the Deployer using a web browser or a terminal.

Step 1 To access the Deployer using a web browser, use the following URL: `https://cli.smi-deployer.deployer-fqdn/`

Step 2 Log in with the following user credentials:

- Username: `admin`
- Password: Available in the `init-k3` section of your deployer configuration file

You must change the password after your first login.

After you log in, you can access the operations center of the Deployer. The operations center provides a CLI environment, where, for example, you can run the `show run` command to show the running configuration.

Troubleshooting

Make sure that the IP addresses in the configuration file and the virtual machine (VM) names are not currently used, when deploying a new deployer or a new Cisco Smart PHY cluster.

Troubleshoot Deploying a New Deployer

- For deployers, the VM name is the same as the deployer name.
- For single-node clusters, the VM name is the cluster name with `-ops` appended.
- For multi-node clusters, there are 12 VMs. The names of these VMs are the cluster names with a comma (,) and `-ops-n` appended, where `n` is 1, 2, or 3. Check if the VM is created on a vCenter.

- Log into the deployer VM using SSH with the correct username and public key file.

  ```
  ssh -i <private-key-file> <deployer-user>@<deployer-address>
  ```

- Use **kubectl** command to find the internal IP address of the Operation Center service:

  ```
  kubectl get svc ops-center-smi-cluster-deployer -n smi
  ```

- Look for the `CLUSTER-IP` field in the output. Log into the deployer through SSH using this cluster IP address and the password for the deployer Operation Center:

  ```
  ssh admin@<cluster-ip> -p 2024
  ```

- Check whether the product tar files available in the `offline-products` directory are downloaded to the deployer:

  ```
  software-package list
  ```

Troubleshoot Deploying a New Cisco Smart PHY Cluster

- Check if the configuration for Cisco Smart PHY clusters is pushed to the deployer:

  ```
  show running-config
  ```

- Monitor the deployment status from the deployer VM:

  ```
  monitor sync-logs <cluster>
  ```

  (Press control-C to quit monitoring)

- Check whether the VMs of the cluster are created on the VMware vCenter.

- Log into the cluster VMs using SSH to see if they are accessible.

- For a single-node cluster, log into the `-ops` VM. For multinode clusters, log into one of the control plane VMs using SSH with the correct username and the SSH private key file.

  ```
  ssh -i <private-key-file> <cluster-user>@<vm-ip-address>
  ```

- Check the Kubernetes cluster using the **kubectl** command.

  For example, to check the status of all pods, use the following command:

  ```
  kubectl get pod --all-namespaces
  ```

  When all pods are in the `Running` state, you can log in to the Cisco Smart PHY UI page.

APPENDIX A

Configure UCS Servers

This section explains how to configure and prepare UCS servers for Cisco Smart PHY software installation.

- Configure the servers using the Cisco Integrated Management Controller (CIMC).
- Install the ESXi Hypervisor.
- Add the ESXi hosts to a vSphere cluster using VMware vCenter.

For more details, see the following sections.

Install and Configure Cisco Smart PHY Server

To install and configure the Cisco Smart PHY server, do the following:

1. Install UCS Server
2. Update Firmware
3. Configure Boot Drives
4. Configure Data Drives
5. Install VMWare ESXi Hypervisor
6. Reboot ESXi Host and Set Boot Device

Install UCS Server

Step 1 Rack mount the servers.

For more details, refer the Cisco UCS C220 M5 Server Installation and Service Guide.

Step 2 Ensure both power supplies are connected on each server and power on the servers.

Step 3 Connect the network cables.

- For CIMC, use the 1 Gb Ethernet dedicated management port.

- For ESXi Host Management, use the Ethernet Port 1 of the Dual 1Gb/10Gb Intel X550T onboard NIC.
- For Cisco Smart PHY application connectivity, connect port 1 on one of the PICe NICs to the management network and connect port 1 on the other PCIe NIC to the CIN Network.

Step 4 Connect the UCS KVM console adapter or connect a keyboard and a monitor directly to the server.

Step 5 Configure the Cisco IMC through the KVM console and update the Network Settings

Update Firmware

Download the latest Hardware Update Utility for the UCS C220 M5 server from Cisco's Software Download site. The Utility helps you to update the CIMC, BIOS, and device firmware for storage controllers, network adapters, SSDs, and other components.

Configure Boot Drives

Step 1 Enable the Cisco MSTOR Boot Optimized M.2 RAID Controller.

Step 2 Create RAID 1 virtual drive from 2 x M.2 SSD drives.

Step 3 Set Stripe Size to 64 KB.

Configure Data Drives

Step 1 Enable Cisco 12G SAS Modular RAID Controller.

Step 2 Create RAID 5 enabled virtual drive using 4 x SSDs.

Step 3 Set Stripe Size to 64 KB.

Step 4 Set the Write Cache Policy to Write Back with Good BBU.

Install the VMware ESXi Hypervisor

Step 1 Download the Cisco custom image for ESXi 6.5 U3 GA Install CD ISO from VMware.

Step 2 Install VMware ESXi 6.5 Update 3 on the M.2 RAID 1 Virtual Drive (Boot Drive).

Use the Cisco Custom ISO: `VMware_ESXi_6.5.0_13932383_Custom_Cisco_6.5.3.1.iso`

Step 3 Set a password for the root user per the installation process.

Step 4 Reboot the VMware ESXi host according to the installation process.

Reboot ESXi Host and Set Boot Device

Step 1 Interrupt the boot process with F2 after the host resets and boot into the BIOS.

Step 2 Under the **Boot Options** tab set the **Boot Option #1** to the UEFI target: `VMWARE ESXi`

Step 3 Disable all other boot options.

Step 4 Save changes and exit.

Step 5 Confirm whether the host boots directly into VMware ESXi.

Add Smart PHY ESXi Hosts to vSphere Virtual Infrastructure

Step 1 Configure ESXi host management networking.

a) Log in to the ESXi host through the DCUI with the root account.

b) Configure the Management Network: Update IP configuration, DNS configuration, custom DNS suffixes, and VLAN ID (if required.)

Step 2 Add ESXi hosts to VMWare vCenter server.

a) In vCenter, create a new, dedicated cluster for Smart PHY.

Do not enable DRS or any HA features.

b) Add each new Smart PHY ESXi host to the new Smart PHY cluster.

Step 3 Configure and enable required ESXi host features.

a) Configure time on the host: Enable NTP.

b) Apply ESXi host licenses.

c) Create a new datastore on the data drive storage device.

Step 4 Configure VM networking.

a) Ensure the VMWare vSwitch connectivty to the physical switch.

b) Create a PortGroup and vSwitch for K8s cluster node VM management network.