



## **Cisco Smart PHY Application Install Guide, Release 3.1.3**

**First Published:** 2020-12-17

### **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

#### **Deploy Cisco Smart PHY 1**

Overview 1

Prerequisites for Cluster Deployment 2

Prerequisites for Cisco Smart PHY Servers 2

Prerequisites for VMware Deployment 2

Deploy Cisco Smart PHY 3

---

### CHAPTER 2

#### **Configure UCS Servers 5**

Install and Configure Cisco Smart PHY Server 5

UCS Server Installation 5

Update Firmware 6

Configure Boot Drives 6

Configure Data Drives 6

Install the VMware ESXi Hypervisor 6

Reboot ESXi Host and Set Boot Device 7

Add Smart PHY ESXi Hosts to vSphere Virtual Infrastructure 7

---

### CHAPTER 3

#### **Prepare Staging Server 9**

Prerequisites 9

Unpack Cisco Smart PHY Application Package 9

---

### CHAPTER 4

#### **Prepare Cluster Configuration File 13**

Cluster Configuration File 13

VMware vCenter Environment Configuration 13

Deployer Configuration 14

Cluster Configuration 15

VMware vCenter Details 17  
IP Addresses for Deployer and Cluster 18  
Cisco Smart PHY CIN Configuration 18  
Sample Configuration Files 18

---

**CHAPTER 5**      **Deploy Cluster and Deployer 19**  
                    Deploy the Deployer or Cluster 19

---

**CHAPTER 6**      **Troubleshoot Cisco Smart PHY Installation 21**  
                    Troubleshooting 21



# CHAPTER 1

## Deploy Cisco Smart PHY

This guide provides information about the deployment of Cisco Smart PHY 3.1.3 in offline environments without Internet connectivity.

- [Overview, on page 1](#)
- [Prerequisites for Cluster Deployment, on page 2](#)
- [Deploy Cisco Smart PHY, on page 3](#)

### Overview

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

Use the `deployer` script available in the software image to set up the Deployer and clusters.

The installation process creates the following components:

- **Deployer:** A controller used to configure and deploy the cluster.
- **Cluster:** The cluster on which the target product application runs.

The Deployer supports two types of cluster deployments:

- **Multinode cluster:** Runs on three UCS servers, each with a control-plane, etcd, Infra, and Operations VMs, which add up to 12 VMs.
- **All-in-one (AIO) cluster:** Runs on a single virtual machine (VM).



**Note** We recommend the multinode cluster for production deployments due to their increased resiliency.

The following table contains the minimum resources for each type of node:

Node Type	CPU Cores	RAM Size (GB)	SSD Storage Size (GB)
Deployer	4	16	110
All-in-one	18	96	700

Node Type	CPU Cores	RAM Size (GB)	SSD Storage Size (GB)
Multinode	34	160	1950

## Prerequisites for Cluster Deployment

The following prerequisite components are required to install, operate, and manage a Cisco Smart PHY cluster.

- Staging server: A physical or virtual machine to run the installation script
- Cisco Smart PHY servers: Three UCS servers are required to run a multinode cluster
- VMware deployment

## Prerequisites for Cisco Smart PHY Servers

Three Cisco UCS C220 M5 servers are required to run a Cisco Smart PHY multinode cluster. The supported Cisco UCS server is UCSC-C220-M5SX.

The minimum compute, storage, and networking requirements for the Cisco UCS server are listed in the following table.

Component	Specification
Chassis	UCSC-C220-M5SX
Processor	Intel 6248 2.5GHz/150W 20C/27.5MB DCP DDR4 2933 MHz
Memory	384 GB DDR4-2933-MHz RDIMM
Storage	4 x 800 GB SSD
NIC	Use one of the following: <ul style="list-style-type: none"> <li>• 2 x Intel XL710-QDA2 (40G)</li> <li>• 2x Intel X550-T2 (10G RJ-45 Copper)</li> <li>• 2x Intel X710-DA2 (10G SFP)</li> </ul>

## Prerequisites for VMware Deployment

VMware ESXi and VMware vCenter Server are mandatory components of the Cisco Smart PHY servers, and is necessary for cluster deployment.

- Hypervisor: VMware ESXi 6.5 Update 3 or VMware ESXi 6.7
- Host Management: VMware vCenter Server 6.5 or VMware vCenter Server 6.7

If the VMware ESXi 6.7 is installed on the host, ensure that the VMware vCenter Server version is 6.7.

# Deploy Cisco Smart PHY

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

1. Configure UCS server.
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

# 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, on page 5](#)
- [Add Smart PHY ESXi Hosts to vSphere Virtual Infrastructure, on page 7](#)

## 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

## UCS Server Installation

---

**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 data, connect port 1 of the Intel XL710 40G NIC in PCIe Slot 1 to the SP Router/Leaf Switch using Cisco QSFP-40G-SR4.

**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** 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 2** Set a password for the root user per the installation process.

**Step 3** 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 connectivity to the physical switch.
    - b) Create a PortGroup and vSwitch for K8s cluster node VM management network.
-





## CHAPTER 3

# Prepare Staging Server

---

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

- [Prerequisites, on page 9](#)
- [Unpack Cisco Smart PHY Application Package, on page 9](#)

## 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 and the Smart PHY cluster. 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
.
├── cluster-deployer-airgap.vmdk
├── cluster-deployer-airgap.vmdk.signature
├── deploy
├── deploy.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-<version>.tar
│   ├── cee-<version>.tar.signature
│   ├── opshub-master-<version>.tar
│   ├── opshub-master-<version>.tar.signature
│   ├── smartphy-master-<version>.tar
│   └── smartphy-master-<version>.tar.signature
├── README.md
├── upgrade-prep
└── utility-images
    ├── autodeploy_<version>.tar.gz
    ├── autodeploy_<version>.tar.gz.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 4

# Prepare Cluster Configuration File

---

To set the cluster configuration file, do 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.
  - [Cluster Configuration File, on page 13](#)
  - [VMware vCenter Details, on page 17](#)
  - [IP Addresses for Deployer and Cluster, on page 18](#)
  - [Cisco Smart PHY CIN Configuration, on page 18](#)
  - [Sample Configuration Files, on page 18](#)

## 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.

## VMware vCenter 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: <ntp server name or IP address>
  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, 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: <vCenter datastore name for host>

```

Guidelines for configuring the deployer:

- The name of the deployer 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 deployer and places it in the `.sec` subdirectory under the staging directory. The filename is `<deployer-name>_auto.pem`.

- To avoid resource contention, do not run the deployer in an ESXi server that serves any Cisco Smart PHY cluster.

## 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: "smartphy"
  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
  master-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"

  # 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>
    ops:
      interfaces:
      - vip: [ <LIST of virtual IP for CIN network in CIDR format> ]
        vrouter-id: <VRRP ID for CIN network>
        addresses: [ <LIST of IP addresses for CIN network in CIDR format> ]
        nics: <vCenter network for CIN>
        # CIN routing table (optional)
        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>
    addresses: [ <CONTROL-PLANE 2 IP>, <ETCD 2 IP>, <INTRA 2 IP>, <OPS 2 IP> ]
    datastore: <vCenter datastore for host 2>
    ops:
      interfaces:
      - addresses: [ <LIST of IP addresses for CIN network in CIDR format> ]
  - 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>
    ops:
      interfaces:
      - addresses: [ <LIST of IP addresses for CIN network in CIDR format> ]

  # For Single-Node cluster only
  nodes:
  - host: <ESXi host IP address>
    addresses: [ <AIO VM IP address> ]
    datastore: <vCenter datastore for host>
    ops:
      interfaces:
      - addresses: [ <LIST of IP addresses for CIN network in CIDR format> ]
        nics: <vCenter network for CIN>
        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 <code>smartphy</code> for Cisco Smart PHY cluster.
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 <code>.sec/</code> .
master-vip	Promary virtual IP address in CIDR format (multinode only).
vrouter-id	VRRP ID for management network (multinode only).
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 subdomains: <ul style="list-style-type: none"> <li>• <code>cli.smartphy-data-ops-center.&lt;fqdn&gt;</code></li> <li>• <code>restconf.smartphy-data-ops-center.&lt;fqdn&gt;</code></li> <li>• <code>cli.opshub-data-ops-center.&lt;fqdn&gt;</code></li> <li>• <code>restconf.opshub-data-ops-center.&lt;fqdn&gt;</code></li> <li>• <code>cli.cee-data-ops-center.&lt;fqdn&gt;</code></li> <li>• <code>restconf.cee-data-ops-center.&lt;fqdn&gt;</code></li> <li>• <code>grafana.&lt;fqdn&gt;</code></li> </ul> 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: <ul style="list-style-type: none"> <li>• The cluster IP address is used to generate an FQDN leveraging <code>nip.io</code> as the domain and top-level domain (TLD). For example, if the IP address of the cluster is <code>10.0.0.2</code>, the generated FQDN is <code>10.0.0.2.nip.io</code>. The subdomains listed here are also leveraged.</li> <li>• Your DNS servers must allow the resolution of the <code>nip.io</code> domain. If resolution of <code>nip.io</code> is blocked, you cannot access the cluster.</li> </ul>
- host	ESXi IP address where VMs are hosted

Command	Description
addresses	IP addresses assigned to control-plane, etcd, infra and docsis or operations nodes respectively
<b>CIN Configuration</b>	
- 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.
<b>For Single-Node cluster</b>	
- host	ESXi IP address where VM is hosted.
ops	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.

## VMware vCenter Details

To contact the VMware vCenter server, the `deployer` script and the `deployer` 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 node)
- One virtual IP address for management network and one for each CIN network (multi-node cluster)

## 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.

## 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).



## CHAPTER 5

# Deploy Cluster and Deployer

This section provides details of the `deployer` script and how to run it. You can run the `deployer` script to deploy both the deployer and the cluster.

- [Deploy the Deployer or Cluster, on page 19](#)

## Deploy the Deployer or Cluster

Run the `deployer` script to deploy the clusters using the following command:

```
$ ./deploy
Usage: ./deploy -c <config_file> <options>
with the following optional parameters or flags:
  -i <install_opt> : Cluster installation option: deploy, redeploy, or upgrade
  -u                : Deployer upgrade flag
  -v                : Config validation flag
```

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.
- `-i <install_opt>`: Cluster installation options: deploy, redeploy, or upgrade [Optional]
- `-u`: Cluster Upgrade Flag [Optional]
- `-v`: Config Validation Flag, [Optional]

The `-u` flag is for updating the deployer. When this flag is present, the script processes all deployers in the `deployers` section in the configuration file. Also, ignores clusters in the `clusters` section.

For cluster installation, use one of the following options in the `-i` flag section:

- `deploy`: The default installation mode when you do not provide the `-i` option in the command. In this mode, the cluster is pinged. If it is not pingable, the cluster is deployed. Otherwise, no operation is performed on the cluster.
- `redeploy`: In this mode, the cluster is uninstalled, if it is already available. Then the new cluster is redeployed.  
With the `redeploy` option, all data in the original cluster is lost.
- `upgrade`: In this mode, the cluster is upgraded with the software in the package.

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 Smart PHY cluster: SSH password for all VMs in the cluster; for user specified in the cluster's config. 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 should be eight characters long, and should 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 is not running, the `deployer` script installs the deployer.
- The `deployer` script checks if the deployer 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.
- The script also generates the configuration for each cluster and pushes them to the deployer.
- The `deployer` script triggers the deployer 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 `sync` operation and the rest. Use this option to push the configuration of a cluster to the deployer without deploying or updating the cluster.





## CHAPTER 6

# Troubleshoot Cisco Smart PHY Installation

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

- [Troubleshooting, on page 21](#)

## Troubleshooting

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

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`.

For a multinode cluster, there are 12 VMs. Their names are cluster names with `-master-n`, `-etcd-n`, `-infra-n`, and `-ops-n`, where `n` is 1, 2, and 3.

### Troubleshoot Deploying a New Deployer

- 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. Use it to SSH into the deployer:  

```
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:

```
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.