



Cisco Smart PHY Application Install Guide, Release 23.2

First Published: 2023-07-31

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2023 Cisco Systems, Inc. All rights reserved.



CONTENTS

Full Cisco Trademarks with Software License ?

CHAPTER 1

Deploying Cisco Smart PHY 1

- Offline Deployment Overview 1
 - Deployment Components 1
 - Deployment Overview 2
 - Deployment Types 2
- Prerequisites for Deployment 3
 - Staging Environment 3
 - Domain Name System 3
 - Cluster Configuration 4
 - VMware vSphere 6
 - VMware ESXi Host Running the Deployer VM 7
 - VMware ESXi Hosts Running Smart PHY Cluster VMs 7
 - Connectivity 8
 - NTP Server 8
- Preparing the Staging Environment 9
 - Transferring the Smart PHY release package to your Staging Environment 9
 - Extracting Installation files and Performinng Signature Verification 9
- Preparing Cluster Configuration File 11
 - Sample Configuration Files 11
 - Cluster Configuration File 11
 - Environment Configuration 11
 - Deployer Configuration 12
 - Cluster Configuration 13
 - Cisco Smart PHY CIN Configuration 17

Deploying the Deployer VM and Cisco Smart PHY Cluster 17

 Verifying Installation 20

Configuring Dual Stack on External Cluster Interfaces 20

Deployment Limitations 21

Smart PHY Client Requirements 22

CHAPTER 2 **Performing Cisco Smart PHY In-Place Software Upgrade 23**

 Prerequisites for In-Place Upgrade 23

 Limitations for In-Place Upgrade 25

 Upgrading Smart PHY 25

 Troubleshooting Common Error Messages 26

CHAPTER 3 **Troubleshooting Cisco Smart PHY Installation 27**

 Access the Deployer 27

 Troubleshooting 27

APPENDIX A **Configuring UCS Servers for Hosting Operations Hub 29**

 Installing VMware ESXi 29

 Rebooting the VMware ESXi Host and Setting the Boot Device 29

 Adding ESXi Hosts to vSphere Virtual Infrastructure 30

 Configuring VMware ESXi Host Management Networking 30

 Adding ESXi Hosts to VMware vCenter Server 30

 Configuring and Enabling ESXi Host Features 31

 Configuring Virtual Machine Networking 31

 Preparing Supporting Software Components 31



CHAPTER 1

Deploying Cisco Smart PHY

This chapter provides information about deploying the Cisco Smart PHY software product package in an offline environment (without Internet connectivity).

- [Offline Deployment Overview, on page 1](#)
- [Prerequisites for Deployment, on page 3](#)
- [Preparing the Staging Environment, on page 9](#)
- [Preparing Cluster Configuration File, on page 11](#)
- [Deploying the Deployer VM and Cisco Smart PHY Cluster, on page 17](#)
- [Configuring Dual Stack on External Cluster Interfaces, on page 20](#)
- [Deployment Limitations, on page 21](#)
- [Smart PHY Client Requirements, on page 22](#)

Offline Deployment Overview

Feature History

Feature Name	Release information	Description
Support for SMI path based URL routing in the Deployer VM.	Cisco Smart PHY 22.3	Deployer VM supports path based URL. Smart PHY only needs two DNS entries as a prerequisite for installation. <ul style="list-style-type: none">• cluster.example.com (DNS entry for Cluster)• deployer.example.com (DNS entry for Deployer)

Cisco Smart PHY supports deployment in an offline operator-managed vSphere virtualization environment.

You can download the Smart PHY software package as a compressed file from the Cisco.com website. The software package contains instructions, sample cluster configuration files, a cluster deployment tool, and the Smart PHY software.

Deployment Components

- **Deploy tool**—A deployment automation tool that controls the deployment of an Smart PHY cluster.

- [Staging Environment, on page 3](#)—A desktop operating system or virtual machine that meets the requirements to run the deploy tool.
- [Cluster Configuration, on page 4](#) file—An YAML-formatted text file that contains the Smart PHY cluster configuration, including the vSphere configuration, and “Deployer” VM configuration. An 'admin' user creates the configuration file.
- **Deployer VM**—The deploy tool instantiates the Deployer VM. This virtual machine hosts the software, container image, and VM image repositories needed to complete an offline Smart PHY cluster deployment.

Deployment Overview

At a high level, deploying the Cisco Smart PHY cluster consists of the following steps:

1. (Optional) Configuring UCS Servers for Hosting Smart PHY. This step is only required when deploying on to Cisco UCS servers dedicated to the cluster.
2. Preparing the Staging Environment
3. Creating the cluster Configuration File
4. Executing the deploy tool. For more information, see [Deploying the Deployer VM and Cisco Smart PHY Cluster, on page 17](#). To deploy another cluster, repeat creating the configuration file and deploying the cluster procedures.

The deploy tool, which is run from your staging environment, reads the Smart PHY cluster configuration from the specified cluster configuration file. After validating the values from the cluster configuration file, the deploy tool instantiates a “Deployer” VM in your designated vSphere environment.

Once the “Deployer” VM boots completely, the deploy tool syncs the Smart PHY cluster configuration to a software agent running on the “Deployer” VM. The agent executes the following Smart PHY cluster operations:

- Copying cluster VM images to the vSphere datastore
- Instantiating cluster VM
- Configuring Guest OS
- Installing and configuring container orchestrations software
- Finally, launching Smart PHY’s containerized micro-services

Deployment Types

The deploy tool can create two types of Cisco Smart PHY clusters:

- **All-in-one (AIO) cluster**—Runs as a single VM on an ESXi host.
 - AIO clusters are best suited for labs and small production environments where high availability is not required.
- **Multinode cluster**—Consists of 12 VMs deployed across three ESXi hosts.
 - Each ESXi host runs one instance of these four VMs: control-plane, etcd, Infra, and Operations.

- Multinode clusters provide high availability and continues to operate even after a failure of one ESXi host.
- Two multimode cluster sizes are available:
 - Small—Best suited for labs and small production environment. This is the default deployment size when no value is specified in the cluster configuration file.
 - Normal—Best suited for large production environment.

Prerequisites for Deployment

This section provides details on prerequisites that must be met before deploying a Cisco Smart PHY cluster. The following resources are required to deploy, operate, and manage the Cisco Smart PHY cluster.

Related Topics

[Staging Environment](#), on page 3

[Domain Name System](#), on page 3

[Cluster Configuration](#), on page 4

[VMware vSphere](#), on page 6

[VMware ESXi Host Running the Deployer VM](#), on page 7

[VMware ESXi Hosts Running Smart PHY Cluster VMs](#), on page 7

[Connectivity](#), on page 8

[NTP Server](#), on page 8

Staging Environment

The staging environment is any Operating System or virtual machine with:

- High-speed, low latency connectivity to the vSphere environment
- At least 50GB of free disk

Prerequisites for Staging Environment

The following software must be installed:

- UNIX compatible shell
- Docker 18.09.7 or later
- Python 3.6 or later

Domain Name System

You can assign a Fully Qualified Domain Names (FQDN) to both the Smart PHY cluster and Deployer VM so that these resources function properly.

Two types of FQDN are available:

1. User-Specified FQDNs (Recommended)
2. Autogenerate FQDNs

User-Specified FQDN

User-Specified FQDN enables you to specify the hostname for both the Smart PHY cluster and deployer VM. You can specify your preferred cluster and deployer VM FQDNs using the “ingress-hostname” key-value pair (optional parameter) in the cluster configuration file.

Prerequisites for User-Specified FQDN

- Supports only alphanumeric characters in the FQDN.
- Unique FQDNs must be assigned to the cluster and deployer VM.

Ensure that you configure the corresponding DNS records (listed below) correctly in your authoritative DNS server before conducting a deployment with your specified FQDNs. If the records are not resolved correctly, then your cluster deployment fails.

Required DNS Records by entity:

- Operations Hub cluster:
 <cluster-fqdn>
- Deployer VM:
 <deployer-vm-fqdn>

For example, if the `ingress-hostname` value in the **clusters** section of your configuration file is `opshub.example.com`, then the required DNS records is `opshub.example.com`.

Autogenerate FQDN

You can trigger autogeneration of an FQDN by omitting the optional parameter “ingress-hostname” key-value pair from the relevant sections of the cluster configuration file.

The deploy tool autogenerates FQDNs by combining the entities’ management IP address, specified in the cluster configuration file, with the “nip.io” domain name.

For example, if 10.0.22.22 is assigned to the cluster management VIP and the `ingress-hostname` key-value pair is omitted, the autogenerated cluster FQDN will be `10.0.22.22.nip.io`.

Prerequisites for Autogenerate FQDN

- Ensure that your DNS servers resolve the nip.io domain properly. If nip.io resolution is blocked, then your cluster deployment fails.

Cluster Configuration

Feature History

Feature Name	Release Information	Description
vSphere Data Center Folder Path support	Cisco Operations Hub, Release 22.4	Operations Hub’s Autodeployer supports the ability to deploy Ops Hub clusters into a defined vSphere datacenter folder path.

You need the following information to prepare the cluster configuration file.

vSphere Environment

Collect or prepare the following information:

- vCenter server hostname or IPv4 address
- vCenter credentials (username and password)
- vCenter Datacenter
- vCenter Datacenter folder path (optional)
- vCenter cluster name
- vCenter networks
- Datastore name
- Datastore folder path (optional)
- ESXi hostnames
- DNS hostname or IPv4 addresses
- Search domains
- NTP hostnames or IPv4 addresses
- HTTPS Proxy Server IPv4 address (if required)
- Environment name (This name is referenced in the cluster configuration file)

Deployer VM

Collect or prepare the following information:

- Guest OS management network:
 - One IPv4 Address
 - Subnet mask in CIDR notation
 - Gateway address
- Ingress Hostname (Optional, but recommended. For more information, see [Domain Name System, on page 3](#).)
- Username (you'll need to choose a username)
- Deployer name (This name is referenced in the cluster configuration file)

Smart PHY Cluster

Collect or prepare the following information:

- Cluster & Guest OS Management network:
 - Cluster
 - Virtual Router Redundancy Protocol (VRRP) ID

- One IPv4 Virtual IP Address (VIP)
- Subnet mask in CIDR notation
- Gateway address
- Guest OS (Must be in the same IPv4 subnet as the cluster management network):
 - AIO: One IPv4 Address
 - Multinode: 12 IPv4 Addresses. (1 address for the Guest OS on each of the 12 cluster VMs)
- Ingress Hostname (Optional, but recommended. For more information, see [Domain Name System, on page 3](#).)
- Deployment size (small or normal)
- Username (you must choose an username)

To prepare the cluster configuration file, see [Cluster Configuration File, on page 11](#).

VMware vSphere

VMware vSphere is the only virtualization environment where the Smart PHY clusters are supported.

Supported Hypervisors

- VMware ESXi 7.0

Supported ESXi Host Management

- VMware vCenter Server 7.0.3

If VMware ESXi 7.0 is installed on the host, then ensure that the VMware vCenter Server version is also 7.0.



Note We recommend that you use VMware vCenter Server 7.0 with VMFS 6 Datastore type.

Datastore Cleanup in ESXi Hosts

The following tasks are applicable for a new deployment of Smart PHY.

- Before deploying SmartPHY, if there are any existing instances of inception server and cluster VM instances, then power off those VMs and use the **Delete from disk** option to delete them. When it is deleted from the host, go to the datastore folder of the host or hosts, and delete the data that is related to the inception and cluster.
- Starting with Cisco Smart PHY 23.1 Release, delete the **smi-base-image** folders in the datastore before deployment, as the Kubernetes version is upgraded.

VMware ESXi Host Running the Deployer VM

One ESXi host is required to run the “Deployer” VM. The “Deployer” VM must not be co-located on the ESXi hosts running cluster VMs.

Prerequisites for VMware ESXi Host Running the Deployer VM

Ensure that the VMware ESXi host has the recommended capacity for compute, storage, and networking which are listed in the following table:

Table 1: Minimum System Requirements - ESXi Hosts

Parameter	Value
Processor	8 vCPUs
Memory	16 GB
Storage	320 GB Minimum 50,000 IOPS (Input/output operations per second) Latency of < 5 ms
NIC	10G vNIC

VMware ESXi Hosts Running Smart PHY Cluster VMs

Three ESXi hosts are required to run a Cisco Smart PHY multinode cluster. Cluster VMs must not be co-located on the ESXi host running the “Deployer” VM.

Prerequisites for VMware ESXi Hosts Running Smart PHY Cluster VMs

Ensure that the VMware ESXi host has the recommended capacity for compute, storage, and networking which are listed in the following table:

Table 2: Minimum System Requirements - ESXi Hosts

Cluster Size	Small	Normal
Processor	20 vCPUs	34 vCPUs
Memory	160 GB	304 GB
Storage	1640 GB Minimum 50,000 IOPS (Input/output operations per second) Latency of < 5 ms	2440 GB Minimum 50,000 IOPS (Input/output operations per second) Latency of < 5 ms
NIC	2x 10G vNIC	2x 10G vNIC

The following tables show the minimum requirements for AIO cluster:

AIO

VM Type	CPU Cores	RAM Size (GB)	SSD Storage Size (GB)
AIO	18	96	1541

The following tables show the minimum requirements for each of the VM types deployed in a multimode cluster:

Small Multimode Cluster

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

Normal Multimode Cluster

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

Connectivity

From the Staging environment, the deploy tool must have connectivity to the following resources:

- Local DNS server
- vCenter server
- IPv4 subnet assigned to the Guest OS management network on the “Deployer” VM
- IPv4 subnet assigned to the Cluster VIP and Guest OS management network on the cluster VMs

The “Deployer” VM, once created, must have connectivity to the following resources:

- Local DNS server
- vCenter server
- IPv4 subnet assigned to the Cluster VIP and Guest OS management network on the cluster VMs

NTP Server

Ensure that the clocks for the staging server, Deployer VM, and Cluster VM are in sync, preferably pointing to the same NTP server.

Preparing the Staging Environment

This section provides details on how to prepare your staging environment for a Smart PHY cluster deployment.

Preparing the staging environment involves the execution of the following procedures

- [Transferring the Smart PHY release package to your Staging Environment, on page 9](#)
- [Extracting Installation files and Performing Signature Verification, on page 9](#)

Transferring the Smart PHY release package to your Staging Environment

-
- Step 1** Download the Smart PHY release package (`smartphy-installer-<version>.SPA.tgz`) from Cisco.com. The package is approximately 15 GB.
- Step 2** Securely copy the release package to your Staging Environment.
-

Extracting Installation files and Performing Signature Verification

The commands listed in this procedure should be executed from the shell in your Staging Environment.

-
- Step 1** Run the `tar` command against the `smartphy-installer-<version>.SPA.tgz` release package to extract the installation files.

Example:

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

The following files are extracted:

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

- Step 2** Run the `cs-verify.sh` script to verify the signature of the `smartphy-installer-<version>.tgz` installer image.

Example:

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

Example output:

```
Verifying signature
```

```
Signature verification succeeded
```

If the signature verification fail, then delete the previously extracted installation files and the installer image. Re-download the Smart PHY release package from Cisco.com and start this procedure again.

Step 3 Run the `tar` command against the extracted `smartphy-installer-<version>.tgz` installer image to create the staging directory.

Example:

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

A new directory named `smartphy-installer-<version>` has been created. This directory is known as the staging directory.

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

Example:

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

The staging directory `smartphy-installer-<version>` contains the following files and folders:

```
:
smartphy-installer-<version>
├── 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-<version>.tar
│   ├── cee-<version>.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
```

Preparing Cluster Configuration File

This section provides info on how to create the Cluster Configuration file required by Smart PHY's deploy tool. This configuration file stores all of the information the deploy tool needs in order to deploy your Smart PHY cluster.

Sample Configuration Files

During the extraction of the Smart PHY release package an `example` directory is created under the staging directory. The `example` directory contains the following sample configuration files:

- `deploy-sample-config.yaml`—A configuration file with a deployer, but no cluster.
- `aio-smartphy-config.yaml`—A configuration file with a deployer and a single-node Smart PHY cluster.
- `multinode-smartphy-config.yaml`—A configuration file a deployer and multinode Smart PHY cluster.
- `aio-smartphy-standby-config.yaml`—A configuration file with a deployer and a single-node Smart PHY cluster that is configured for standby (without CIN configuration).
- `multinode-smartphy-standby-config.yaml`—A configuration file with a deployer and a multinode Smart PHY cluster configured standby (without CIN configuration).

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

The `environments` section defines a vSphere deployment domain. This environment is referenced in the `deployers` and `clusters` sections, which you define shortly.

```
environments:
  <env-name>:
    server: <value>           # vCenter Server IPv4 Address or name
    username: <value>        # vCenter username, user will be prompted for
                             # the password
    datacenter: <value>      # vCenter Datacenter Name
    datacenter-path: <value> # vCenter Datacenter folder path (optional). When
                             # omitted cluster VMs are created under root of the
                             #
                             # Datastore. For nested folders, the folder name
                             # should be separated by / symbol
    cluster: <value>        # vCenter Cluster Name
```

```

    nics: [ <list> ]           # List of vCenter nics (port groups)
    nameservers: [ <value> ]  # List of DNS Server IPv4 Addresses
    search-domains: [ <value> ] # List of Search domains
    ntp: [ <list> ]           # List of NTP Server IPv4 Addresses or name
    https-proxy: <value>     # Optional HTTPS Proxy
                              (Ex: http://proxyhost.domain.tld:port)
    no-proxy: <value>        # Optional HTTPS Proxy bypass

```

Guidelines for Defining an Environment

- The environment name can have only lowercase letters, digits, and hyphens (-).
- The `nics` list must contain only one network, although the `nics` configuration allows multiple networks. This network is used as the management network by the deployer or cluster that refers to this environment.
- Create multiple `environments` if your vCenter has more than one network that serves as a management network. Create one `environments` for each network. In addition, refer to the corresponding `environments` in the deployer or cluster based on the management network it uses.
- Make sure the `nics`, `nameservers`, and `search-domains` values are structured as YAML lists.

Deployer Configuration

The `deployers` section defines the configuration of a Deployer VM.

```

deployers:
  <deployer-vm-name>:      # Deployer VM name
    environment: <value>  # Environment name defined in the 'Environments'
                          # section.
    address: <value>      # Deployer VM IPv4 Address in CIDR format
    gateway: <value>     # Deployer VM's Gateway IP Address
    ingress-hostname: <value> # Optional FQDN (Ex. "deployer1.example.com")
    username: <value>    # Deployer VM username, you are prompted for
                          # the password
    private-key-file: <value> # Optional SSH private-key-file (.pem) with
                              # path relative to the staging directory. Key is
                              # auto-generated, if one if not provided
    host: <value>        # IPv4 Address of the ESXi Host
    datastore: <value>   # Datastore for the Deployer VM
    datastore-folder: <value> # Optional Datastore folder path. When omitted,
                              # deployer VM is created under root of the
                              # Datastore. For nested folders, the folder
                              # name should be separated by / symbol
    docker-subnet-override: # 'docker-subnet-override' and its values are
                              # optional. It should only be used when you need
                              # to customize the deployer VM's Docker
                              # IP Addressing. When omitted, the Docker bridge
                              # defaults to 172.17.0.0/16
    - pool-name: <value>   # Docker bridge address pool name
      base: <value>       # Docker bridge subnet in CIDR format
      size: <value>      # Docker bridge subnet size: 8-24

```

Guidelines for Defining a Deployer

- The `deployer-vm-name` can have only lowercase letters, digits, and hyphens (-).
- The `private-key-file`, when present, must refer to the SSH private key file (.pem). 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 omitted, the deploy tool generates an SSH private key for the deployer and places it in the `.sec` subdirectory under the staging directory. The filename is `<deployer-vm-name>_auto.pem`.
- The values associated with `docker-subnet-override` are optional. Those values should only be included in the configuration when you need to customize the deployer VM's Docker IP addressing. When omitted, the Docker bridge on the Deployer VM defaults to 172.17.0.0/16.
- To avoid resource contention, do not host the deployer VM on the same ESXi hosts running any of the cluster VMs.
- When you configure a custom `ingress-hostname`, ensure that the following entries are in the DNS:

```

<host.domain.tld>
charts.<host.domain.tld>
files-offline.smi-cluster-deployer.<host.domain.tld>
deployer-ui.smi-cluster-deployer.<host.domain.tld>
cli.smi-cluster-deployer.<host.domain.tld>
restconf.smi-cluster-deployer.<host.domain.tld>
docker.<host.domain.tld>

```

Cluster Configuration

Table 3: Feature History

Feature Name	Release Information	Description
Single NIC for NB and SB communication for AIOs	Cisco Smart PHY, Release 23.2	You can optionally deploy an AIO cluster with a single vNIC.

The `clusters` section defines the type and configuration of the `cluster`. At least one `environment` and one `deployer` must be defined in the cluster configuration file in order for a `cluster` to be deployed.

`Clusters` can be deployed on a single ESXi host or across three ESXi hosts. The single host deployment is known as a single-node deployment, or an All-in-one (AIO), while the three-node deployment is known as a multi-node deployment.

By default, All-in-one clusters are deployed with two vNICs. One vNIC must be logically connected to a management network, while the other must be logically connected to the CIN. If you have a shared management and CIN network, you can optionally deploy an AIO cluster with a single vNIC.

The following `clusters` configuration below shows the mandatory and optional key-value pairs required for a multi-node deployment.

```

clusters:
  <cluster-name>:                                # Cluster name
    type: <value>                                # Cluster type must be 'opshub'
    size: <value>                                # Optional cluster size: 'small' or
'normal'.                                       Defaults to 'small' when not
specified.
  environment: <value>                           # Environment name defined in the
'Environments' section.
  gateway: <value>                                # Cluster Gateway IPv4 Address
  ingress-hostname: <value>                       # Optional FQDN (Ex.
"smartphy.example.com")
  username: <value>                               # Cluster username, User is

```

```

private-key-file: <value>
    prompted for cluster password
    # SSH private-key-file (.pem) including
    path relative to the staging directory
    # Key is auto-generated, if not
    # Management Virtual IPv4 Address in
    # Management Keepalived Virtual Router
    value must be between 0-255
    # Optional. Defaults to false when
    Set to "true" to redirect HTTP
    # The next three key-value pairs are
    They should only be used when you
    customize the cluster's internal
    # K8s Pod subnet in CIDR format. If
    omitted, defaults to: 192.168.0.0/16
    # K8s Service subnet in CIDR format.
    If omitted, defaults to 10.96.0.0/12
    # List of Docker bridge subnets in
    If omitted, defaults to 172.17.0.0/16
    #
    # ESXi Host 1: IPv4 Address
    # ESXi Host 1: List of Mgmt IPv4 addr
    control-plane, etcd, infra, and Ops
    # ESXi Host 1: IPv4 Address for vCenter
    # ESXi Host 2: IPv4 Address
    # ESXi Host 2: List of Mgmt IPv4 addr
    control-plane, etcd, infra, and Ops
    # ESXi Host 2: IPv4 Address for vCenter
    # ESXi Host 3: IPv4 Address
    # ESXi Host 3: List of Mgmt IPv4 addr
    control-plane, etcd, infra, and Ops
    # ESXi Host 3: IPv4 Address for vCenter
    #
    #
    # All of the parameters below are
    #
    # ESXi Host 1: IPv4 Address (Same
    used earlier in the nodes section.)
    # ESXi Host 1: vCenter list of Network

provided.
    primary-vip: <value>
    CIDR format
    vrouter-id: <value>
    ID,
    enable-http-redirect: value
    not specified.
    requests to HTTPS.
    optional.
    need to
    IP Addressing
    pod-subnet: <value>
    service-subnet: <value>
    docker-bridge-subnet: [ <addr-list> ]
    CIDR format.
    nodes:
    - host: <value>
      addresses: [ <addr-list> ]
    assigned to
    VMs respectively
      datastore: <value>
    Datastore
    - host: <value>
      addresses: [ <addr-list> ]
    assigned to
    VMs respectively
      datastore: <value>
    Datastore
    - host: <value>
      addresses: [ <addr-list> ]
    assigned to
    VMs respectively
      datastore: <value>
    Datastore
    apps:
    - smartphy:
    Smart PHY specific
      nodes:
      - host: <value>
    address as
    nics: [ <list> ]

```

```

    for CIN
      ops: # -- The following parameters apply
to Ops VM 1. --
      interfaces: # Ops VM 1: CIN Interface configuration:
        - addresses: [ <addr-list> ] # Ops VM 1: List of CIN IPv4 or
          # IPvv6 Addresses in CIDR format
          vip: [ <vip-list> ] # Ops VM 1: List of CIN Virtual IPv4
          # or v6 Addresses in CIDR format
          vrouter-id: <value> # Ops VM 1: CIN Keepalived Virtual
Router ID, # value must be between 0-255
# Ops VM 1: Optional Route
      routes:
configuration:
  - { dest: [ <list> ], nhop: <value> } # Ops VM 1: Optional list of Destination
# Subnets in <IP address>/<subnet mask>
# format; Next-Hop IP Address
  - { dest: [ <list> ], nhop: <value> } # Ops VM 1: Optional list of Destination
# Subnets in <IP address>/<subnet mask>
# format; Next-Hop IP Address
- host: <value> # ESXi Host 2: IPv4 Address (Same
address # as used earlier in the nodes section.)
# ESXi Host 2: vCenter list of Network
      nics: [ <list> ] # ESXi Host 2: vCenter list of Network
    for CIN
      ops: # -- The following parameters apply
to Ops VM 2. --
      interfaces: # Ops VM 2: CIN Interface configuration:
        - addresses: [ <addr-list> ] # Ops VM 2: List of CIN IPv4 or
          # IPvv6 Addresses in CIDR format
          vip: [ <vip-list> ] # Ops VM 2: List of CIN Virtual IPv4
or # IPv6 Addresses in CIDR format
          vrouter-id: <value> # Ops VM 2: CIN Keepalived Virtual
Router # ID, value must be between 0-255
# Ops VM 2: Optional Route
      routes:
configuration:
  - { dest: [ <list> ], nhop: <value> } # Ops VM 2: Optional list of Destination
# Subnets in <IP address>/<subnet
mask> # format; Next-Hop IP Address
  - { dest: [ <list> ], nhop: <value> } # Ops VM 2: Optional list of Destination
# Subnets in <IP address>/<subnet
mask> # format; Next-Hop IP Address
- host: <value> # ESXi Host 3: IPv4 Address (Same
address as # used earlier in the nodes section.)
# ESXi Host 3: vCenter list of Network
      nics: [ <list> ] # ESXi Host 3: vCenter list of Network
    for CIN
      ops: # -- The following parameters apply
to Ops VM 3. --
      nterfaces: # Ops VM 3: CIN Interface configuration:

```

```

- addresses: [ <addr-list> ] # Ops VM 3: List of CIN IPv4 or
                             # Ipv6 Addresses in CIDR format
  vip: [ <vip-list> ] # Ops VM 3: List of CIN Virtual IPv4
or
                             # IPv6 Addresses in CIDR format
  vrouter-id: <value> # Ops VM 3: CIN Keepalived Virtual
                             # Router ID, value must be between
0-255
  routes: # Ops VM 3: Optional Route
configuration:
- { dest: [ <list> ], nhop: <value> } # Ops VM 3: Optional list of Destination
                                     # Subnets in <IP address>/<subnet mask>
                                     # format; Next-Hop IP Address
- { dest: [ <list> ], nhop: <value> } # Ops VM 3: Optional list of Destination
                                     # Subnets in <IP address>/<subnet mask>
                                     # format; Next-Hop IP Address

```



Note In the preceding `clusters` configuration, the syntax for configuring all instances of `routes` (as shown below) is only applicable for Cisco Smart PHY, Release 23.2 and later. See [Install and Upgrade Guides](#), to view the `clusters` configuration for previous releases of Cisco Smart PHY.

```

routes: # Ops VM VM Number:Optional Route configuration:
- { dest: [ <list> ], nhop: <value> } # Ops VM VM Number: Optional list of Destination
                                     # Subnets in <IP address>/<subnet mask>
                                     # format; Next-Hop IP Address
- { dest: [ <list> ], nhop: <value> } # Ops VM VM Number: Optional list of Destination
                                     # Subnets in <IP address>/<subnet mask>

```

Guidelines for Defining a Cluster

- The `cluster-name` can have only lowercase letters, digits, and hyphens (-).
- When you specify a FQDN in the `ingress-hostname`, ensure that the corresponding entries are configured in your DNS servers:

```

<host.domain.tld>
opscenter.<host.domain.tld>

```

- If you do not specify an FQDN in the `ingress-hostname`, the cluster's `primary-vip` (also known as the Management Virtual IP address) is used to generate an FQDN leveraging `nip.io` as the domain and top-level domain (TLD). For example, if the `primary-vip` is 10.0.0.2, the generated FQDN is 10.0.0.2.nip.io. Your DNS servers must allow the resolution of the `nip.io` domain. If resolution of `nip.io` is blocked, you cannot access the cluster.
- The `private-key-file`, when present, must refer to the SSH private key file (.pem). 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 omitted, the deploy tool 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`.
- If multiple clusters share the same management subnet, the management `vrouter-id` (VRRP ID) of each cluster must be unique.

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 Defining Smart PHY's CIN Interfaces

- The CIN Virtual IP addresses (`vip`) field is mandatory. You can configure up to one IPv4 and one IPv6 addresses per CIN network.
- The CIN Virtual IP addresses (`vip`) and the VRRP ID (`vrouter-id`) fields are used only in multi-node cluster deployments. They are configured on the first node.
- If multiple Smart PHY clusters share a CIN subnet, the VRRP ID (`vrouter-id`) should be unique for each cluster.
- For multi-node clusters, all Ops VMs must have the same number of CIN interfaces. The `nics` or `route` fields must be explicitly mentioned on all Ops VM nodes.
- You can setup a Smart PHY cluster as a backup cluster. To do so, do not include any CIN configuration. The configuration should not have `ops` and `interfaces` under `nodes`.



Note Cisco Smart PHY clusters can connect to multiple CIN networks using multiple network interfaces configured during deployment.

Adding CIN Configuration Without Cluster Reboot

After deploying Smart PHY, you can add new CIN Configuration without restarting the cluster by using the following steps:

1. In the Day-0 config, add additional CIN Configurations details and provide the necessary details of CIN, interface address, vip, and vrouter-id config.
2. Run cluster deployment with the `-np` argument. Example: `./deploy -c day0.yaml -np`

Once deployment is done successfully, the cluster is updated with CIN Configuration information. The cluster does not reboot.

Deploying the Deployer VM and Cisco Smart PHY Cluster

Table 4: Feature History

Feature Name	Release Information	Description
Deploying Clusters on Network Based Storage Devices Using vSAN	Cisco Smart PHY, Release 23.2	You can deploy clusters to remote network-attached storage devices using vSAN.

This section explains how to use the deploy tool to deploy the Deployer VM and Cisco Smart PHY cluster.

From the staging environment, run the deploy tool 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]
  -u               : Cluster chart Upgrade Flag [Optional]
  -s               : Skip Compare Flag [Optional]
  -sc              : Skip Compatibility check during upgrade Flag [Optional]
  -D               : Enable Debug Logs [Optional]
  -np              : Provision new NIC without cluster reboot [Optional]

```

The following options are available in the deploy tool:

- `-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 VMs are rebooted, and the data persisted on disk is retained. You can use this option to modify some of the cluster parameters.

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

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

The deploy tool does the following operations:

- If you're running the deploy tool 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're 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 deploy tool 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 isn't running, the deploy tool installs the deployer VM.
- The deploy tool 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 tool also generates the configuration for each cluster and pushes them to the deployer VM.

- The deploy tool triggers the deployer VM to perform the sync operation for the cluster. The sync operation applies the configuration to the cluster. If you haven't 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 deploy tool triggers the sync operation again. The tool 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 deploy tool 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 deploy tool 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

Deploying Clusters on Network Based Storage Devices Using vSAN

The Operations hub platform and the applications running over it can be deployed on a vSAN cluster. vSAN stands for Virtual Storage Area Network. It is a software-defined storage (SDS) solution that is offered by VMware that virtualizes and abstracts storage resources in a vSphere environment. vSAN combines local storage devices such as hard drives and solid-state drives (SSDs) from multiple hosts into a single shared storage pool.

Prerequisites for Deploying Clusters on vSAN

1. All the ESXi hosts used for the cluster deployment in the day-0 configuration yaml should be part of the same vSAN cluster.
2. The network speed between ESXi hosts in a vSAN cluster should be at least 10 Gbps.
3. On VCenter 7.0, proactive tests are available to ensure the vSAN cluster setup is proper. We recommend that you run proactive tests before installation to ensure the vSAN cluster is healthy. Use the following steps to run the Proactive tests using the VCenter GUI:

- Access **Cluster > Monitor > vSAN** and select **Proactive tests**.

- You can run the **VM Creation Test**, **Network Performance Test**, and **Storage Performance Test**. These tests must pass.
4. We recommend that you enable vSphere HA if there is a host failure. VMs should be moved to the active host. If vSphere HA is not configured, then it is recommended to use **VM/Host groups** and **VM/Host rules** to ensure the **VM-to-host** allocation model is mentioned in the day-0 configuration yaml.
 - ESXi-1 should host Control-plane-1, ops-1, infra-1, etcd-1
 - ESXi-2 should host Control-plane-2, ops-2, infra-2, etcd-2
 - ESXi-3 should host Control-plane-3, ops-3, infra-3, etcd-3

Verifying Installation

After successfully deploying the Cisco Smart PHY application using the deploy tool, 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.

Verifying Installation

After successfully deploying the Cisco Smart PHY application using the deploy tool, 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.

Configuring Dual Stack on External Cluster Interfaces

Table 5: Feature History

Feature Name	Release Information	Description
IPv6 Dual Stack on support on external cluster interfaces	Cisco Smart PHY, Release 22.2	Support that is extended to configure IPv6 dual stack configuration parameters such as mode, gateway, subnet address, and so on, on an external cluster interface.

Prerequisites

1. Dual stack must be configured at the time of cluster creation. It cannot be added to a previously created Smart PHY cluster.
2. The ESXi hosts must be connected to dual stack enabled networks.

To configure dual stack on an external cluster interface, perform the following steps:

1. Navigate to the cluster configuration file in the staging environment.
2. Add the following parameters in the configuration file.

Table 6: IPv6 Dual Stack Parameters

Parameter	Description	Value to set
ipv6-mode (Optional)	Specifies to set whether the IPv6 mode is true or false.	<ul style="list-style-type: none"> • true—To enable the dual stack on the cluster interface. • false—To disable the dual stack on the cluster interface.
primary-vip-ipv6	Specifies the IPv6 address where the Nginx ingress controller binds to.	This parameter is mandatory if the ipv6-mode is set as true.
ipv6-gateway	Specifies the IPv6 gateway address.	This parameter is mandatory if the ipv6-mode is set as true.
pod-subnet-ipv6	Specifies IPv6 subnet address.	This parameter is optional. If no value is specified, then by default the value "fd20::0/112" is assigned.
service-subnet-ipv6	Specifies the IPv6 service subnet address.	This parameter is optional. If no value is specified, then by default the value "fd20::0/112" is assigned. This parameter is valid only if the ipv6-mode is set as true.
addresses-v6	Specifies the virtual machine SSH IPv6 address.	This parameter is mandatory if the ipv6-mode is set as true.

Deployment Limitations

- Modification of cluster parameters such as server names, NTP server configuration details, data store file path, subnets, IP addresses of VMs, and so on, requires a VM restart. You can restart VM using the **deploy-f** command.
- Autodeployer only supports Application Product chart and docker image upgrades. The modification of cluster configuration is not supported as part of the upgrade process.
- When you enable dual-stack, you must redeploy the cluster.
- The Deployer VM must be saved so that you can use the VM while upgrading the cluster.

- Manual provisioning of NIC interfaces on VMs must be performed through vCenter during the SmartPHY installation on top of the running Operations Hub platform.
- Removal of existing NICs requires a VM restart. You can restart VM using the **deploy-f** command.
- Data store folders must be created in vCenter manually before starting the Smart PHY installation.

Smart PHY Client Requirements

Smart PHY supports the following Operating Systems (OS) and Browsers.

Supported Operation Systems

- Mac OS 12 and later.
- Windows 10 and later.

Supported Browsers

- Mozilla Firefox Version 115 and later.
- Google Chrome Version 114 and later.
- Microsoft Edge Version 114 and later.

Windows System Resolution

Smart PHY UI screens are best seen when Windows display resolution is set to 1920 x 1080 and the size of the text and apps is set to 100%.



CHAPTER 2

Performing Cisco Smart PHY In-Place Software Upgrade

Feature History

Feature Name	Release information	Description
Support for Cisco Smart PHY In-Place Software Upgrade	Cisco Smart PHY, Release 22.2	You can perform in-place software upgrade on Cisco Smart PHY. Use the in-place upgrade to update your existing installation to the new version of Cisco Smart PHY, retaining your existing configuration.

Cisco Smart PHY supports in-place software upgrade. Use the in-place upgrade to update your existing installation to the new version of Cisco Smart PHY, retaining your existing configuration.



Note The software upgrade process retains all the application data.

- [Prerequisites for In-Place Upgrade, on page 23](#)
- [Limitations for In-Place Upgrade, on page 25](#)
- [Upgrading Smart PHY, on page 25](#)
- [Troubleshooting Common Error Messages, on page 26](#)

Prerequisites for In-Place Upgrade

Cisco Smart PHY supports in-place software upgrade. Use the in-place upgrade to update your existing installation to the new version of Cisco Smart PHY, retaining your existing configuration. The prerequisites for the In-Place Upgrade are listed below.



Note The software upgrade process retains all the application data.

1. You can upgrade the following Cisco Smart PHY versions to Cisco Smart PHY 23.2.
 - Cisco Smart PHY 23.1
 - Cisco Smart PHY 22.4

- Cisco Smart PHY 22.3.1
- Cisco Smart PHY 22.3

Use the Deployer tool bundled along with the software package.



Note We highly recommend that you upgrade to the latest version of Cisco Smart PHY.

2. Ensure that the cluster configuration file (Day-0 Config File), used during installation is available.
3. Ensure that the SSH user private key for Smart PHY Cluster and Deployer VM is available.



Note

- The SSH user private key is generated at the time of Smart PHY installation and is available in the installation folder.
- SSH to all the cluster VMs should be possible through the cluster private key.
- Ensure the user and group for the private key file is same as the user running the upgrade operation.

4. Ensure that Deployer VM admin password and all application-specific OpsCenter passwords (cee-opscenter, opshub-opscenter, and smartphy-opscenter) are available, and match with the previous deployment.
5. Ensure that a Staging environment having network connectivity with Deployer VM and Cluster VMs is available.
6. Ensure all the cluster VMs are up and running using vCenter or by checking the individual cluster.
7. Ensure that all the clocks of all the ESXi host's involved in the upgrade are synchronized.
8. Ensure that the clocks for the staging server, Deployer VM, and Cluster VM are in sync; preferably pointing to the same NTP server.
9. Ensure that the cluster does not have the following critical alerts that are related to disk usage:
 - **node-disk-running-Low-24hours**
 - **node-disk-running-Low-2hours**
 - **node-disk-running-full-24hours**
 - **node-disk-running-full-2hours**



Note In case you see any of these alerts, contact Cisco Technical support.

10. Ensure that from the Staging environment, the deploy tool must have connectivity to the following resources:
 - Local DNS server

- vCenter server
- NTP server
- All ESXi hosts

Limitations for In-Place Upgrade

- You can use the `./deploy -c <config.yaml> -u` command to upgrade the cluster charts or the application images.
- You cannot use the `./deploy -c <config.yaml> -u` command to modify any cluster parameters and environmental parameters such as NTP server IP, DNS configuration, VM IP, datastore folder, etc. See [Deploying the Deployer VM and Cisco Smart PHY Cluster, on page 17](#) to modify any cluster parameter.

Upgrading Smart PHY

Use the following steps to perform Cisco Smart PHY In-Place software upgrade:

Step 1 Download the latest Smart PHY release package using the link shared by Cisco.

Step 2 Copy the downloaded Smart PHY release package to the Staging environment.

```
cp download-path/smartphy-installer-<new-version>.tgz
staging-server-path/smartphy-installer-<new-version>.tgz
```

Step 3 Extract the image contents.

```
cd staging-server-path && tar -xvfz smartphy-installer-<new-version>.tgz
```

Step 4 Access the new installation directory.

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

Step 5 Copy the cluster configuration file used for cluster installation into the new installation directory.

```
cp filepath/config.yaml smartphy-installer-<new-version>/config.yaml
```

`filepath/config.yaml` is the cluster configuration file that is used in the previous Smart PHY install or upgrade.

Step 6 Copy the SSH user private key file for Deployer VM and Cluster VM into the new installation directory.

Note The private key must be in pem format and the name of the private key must be identical to the cluster configuration file entry.

```
cp filepath/private-key-file.pem smartphy-installer-<new-version>/private-key-file.pem
```

`filepath/private-key-file.pem` is the private key file that is generated during the previous Smart PHY install or upgrade.

Step 7 Upload the application software package to the Deployer VM.

```
./deploy -c <config.yaml> -u -s
```

This step is optional. Use this step to upload the application software package to the Deployer VM and perform the upgrade later using Step 8.

Observe the Deployer tool logs visible on the terminal and check whether the package is successfully uploaded into the Deployer VM. You can also log in to the Deployer VM and check the availability of the software package under the `/data/software/images` directory.

Step 8 Install the application software package software to the target cluster.

```
./deploy -c <config.yaml> -u
```

Step 9 Continue monitoring the log statements in the Deployer tool and provide the password, if prompted. The upgrade process may take several minutes to complete. On successful completion of the upgrade, the message `Upgrade has been done successfully!!` displays on the terminal.

Troubleshooting Common Error Messages

During the upgrade process, you can monitor the Operational Hub Alert Dashboard for upgrade related alerts. The following table describes common errors which may display during a software upgrade.

Table 7: Common Error Messages Displayed During a Software Upgrade

Error Message	Reason for Error Message	Action Performed
The cluster sync failed during the upgrade. Attempting to recover the cluster.	Cluster sync is unsuccessful	The Deployer tool automatically attempts to recover the cluster.
Post-upgrade cluster status check failed, the cluster is likely not healthy	The charts are not deployed or the essential pods fail to load	No user action is needed.
Unable to apply some pre upgrade cluster configuration	The Deployer tool undertakes additional functionality which requires special configuration changes.	The upgrade process automatically retries the configuration changes. No manual intervention is required.

For any other issues, contact [Cisco Technical Support](#)



Note We recommend creating a unique staging environment for each of your clusters. Staging environments contain cluster-specific information including keys and configurations. You shouldn't create or manage more than one cluster from a single staging environment.



CHAPTER 3

Troubleshooting Cisco Smart PHY Installation

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

- [Access the Deployer, on page 27](#)
- [Troubleshooting, on page 27](#)

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:deployer-fqdn/smi-deployer/cli`
- Step 2** Log in with the following user credentials:
- Username: `admin`
 - Password: The password set during deployment of deployer and cluster.
- For more information, see [Deploying the Deployer VM and Cisco Smart PHY Cluster, on page 17](#).

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**

Configuring UCS Servers for Hosting Operations Hub

- [Installing VMware ESXi, on page 29](#)
- [Rebooting the VMware ESXi Host and Setting the Boot Device, on page 29](#)
- [Adding ESXi Hosts to vSphere Virtual Infrastructure, on page 30](#)
- [Configuring VMware ESXi Host Management Networking, on page 30](#)
- [Adding ESXi Hosts to VMware vCenter Server, on page 30](#)
- [Configuring and Enabling ESXi Host Features, on page 31](#)
- [Configuring Virtual Machine Networking, on page 31](#)
- [Preparing Supporting Software Components, on page 31](#)

Installing VMware ESXi

To install VMware ESXi version 7.0.3, perform the following steps:

1. Download the ESXi server software from the VMware website download page.
2. Install the VMware ESXi 7.0.3 version on the M.2 RAID 1 Virtual Drive (Boot Drive).
3. Select a disk to install the VMware ESXi server software.
4. Set a password for the root user during the installation process.
5. Reboot the VMware ESXi host when the installation completes.
6. Add the ESXi server in the production vCenter version 7.0.3.



Note vCenter 8.0 can manage ESXi hosts running version 7.0.

Rebooting the VMware ESXi Host and Setting the Boot Device

When the VMware ESXi host resets and boots into the BIOS mode, you must perform the following steps:

-
- Step 1** Press the F2 key to interrupt the boot process.
- Step 2** In the **Boot Options** tab, set the Boot Option #1 to the UEFI target - *VMware ESXi*.
- Step 3** Disable all other boot options.
- Step 4** Click **Save** and **Exit**. Ensure that the host boots directly into VMware ESXi.
-

Adding ESXi Hosts to vSphere Virtual Infrastructure

1. Configuring VMware ESXi Host Management Networking
2. Adding ESXi Hosts to VMware vCenter Server
3. Configuring and Enabling ESXi Host Features
4. Configuring Virtual Machine Networking

Configuring VMware ESXi Host Management Networking

To configure management network settings for the VMware ESXi host, perform the following steps:

-
- Step 1** Log into the VMware ESXi host from the vSphere page as a root user.
- Step 2** From the VMware vSphere page, choose **Configure** > **Networking** > **Virtual Switches** to open a **Configure Management Network** window.
- Step 3** Edit the following details:
- IP Address Configuration
 - DNS Configuration
 - Custom DNS suffixes
 - VLAN ID (optional)
- Step 4** Click **Save**.
-

Adding ESXi Hosts to VMware vCenter Server

To add ESXi hosts to the VMware vCenter server, use the following steps:

-
- Step 1** From the VMware vCenter page, select the VM cluster, and choose **Add Hosts**.
- Step 2** In the **Add Hosts** window, enter the IP Address or FQDN hostname with credential, and click **Next**.

Step 3 Click **Finish**. The ESXi host is added to the vCenter.

Configuring and Enabling ESXi Host Features

Once the ESXi host is installed, you must configure the following key features:

1. System Time or Clock—Configure time on the host. For this you must enable NTP on the ESXi host.
2. Licenses—Apply the ESXi host licenses.
3. Network settings—Create a new network configuration for the host
4. Datastore—Create a new datastore on the data drive storage device in the ESXi host.

Configuring Virtual Machine Networking

To configure the virtual machine networking, perform the following steps:

-
- Step 1** From the VMware vCenter page, select the ESXi host.
- Step 2** To configure the VMware vCenter management network, choose **Configure > Networking > Virtual Switches > Add Physical Network**.
- Step 3** In the **Add Physical Network** window, enter IP address, Gateway and VLAN ID details.
- Step 4** Click **Configure**. The physical network is configured for VM.
-

Preparing Supporting Software Components

To prepare the Cisco Unified Computing System (UCS) servers for software installation, ensure that you have performed the following tasks:

- Rack mount the Cisco UCS servers and complete the power connections and cabling.
- Configure the servers using [Cisco Integrated Management Controller \(CIMC\)](#).

