CPS vDRA Installation Guide for OpenStack, Release 18.1.0 (Restricted Release)

First Published: 2018-03-16

Last Modified: 2018-03-16

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

Preface

About this Guide

This document describes the sequence of steps taken to launch CPS on the Microservices platform.

This document also describes the assumptions, pre-conditions (in brief), and the steps taken by the system as it performs the low-level startup of the system.

Audience

This guide is best used by these readers:

- Network administrators
- Network engineers
- Network operators
- System administrators

This document assumes a general understanding of network architecture, configuration, and operations.

Additional Support

For further documentation and support:

- Contact your Cisco Systems, Inc. technical representative.

- Call the Cisco Systems, Inc. technical support number.
- Write to Cisco Systems, Inc. at support@cisco.com.
- Refer to support matrix at https://www.cisco.com/c/en/us/support/index.html and to other documents related to Cisco Policy Suite.

Conventions (all documentation)

This document uses the following conventions.

Conventions	Indication
bold font	Commands and keywords and user-entered text appear in **bold** font.
italic font	Document titles, new or emphasized terms, and arguments for which you supply values are in *italic* font.
[ ]	Elements in square brackets are optional.
{x | y | z }	Required alternative keywords are grouped in braces and separated by vertical bars.
[x | y | z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
courier font	Terminal sessions and information the system displays appear in courier font.
< >	Nonprinting characters such as passwords are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

Note Means reader take note. Notes contain helpful suggestions or references to material not covered in the manual.

Caution Means reader be careful. In this situation, you might perform an action that could result in equipment damage or loss of data.

Warning IMPORTANT SAFETY INSTRUCTIONS.

Means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Warning Provided for additional information and to comply with regulatory and customer requirements.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see What's New in Cisco Product Documentation.

To receive new and revised Cisco technical content directly to your desktop, you can subscribe to the What's New in Cisco Product Documentation RSS feed. RSS feeds are a free service.

RESTRICTED RELEASE

Important This is a Short Term Support (STS) release with availability and use restrictions. Contact your Cisco Account or Support representatives for more information.

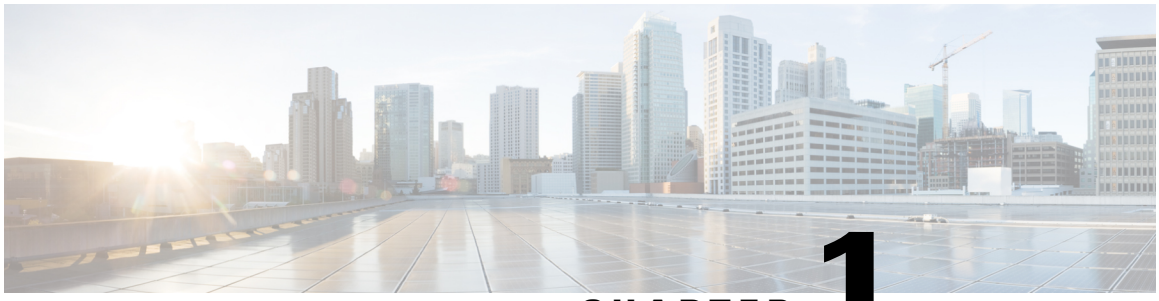

CHAPTER 1

Microservices Platform

Installation Overview

Cisco Policy Suite (CPS) architecture is now designed to use the Docker container technology.

This creates a platform that is more flexible, easier to adopt for new products, easier to deploy, and has much less rigidity with respect to its organization. Other benefits are easier upgrade and patching, and a much crisper separation of the “platform” common components versus the “application” components.

This document is intended to describe how the system is launched, and what inputs are required for the collection of Docker engine VMs to self-organize.

Before You Begin

This document makes the following assumptions:

- OpenStack as a target environment.
- Use of Nova as it is the most basic of all deployment approaches.

For more information about OpenStack, see the OpenStack documentation at https://docs.openstack.org/.

Before you set up CPS vDRA, perform the following steps:

1. Create the Project and User under which the system will be launched.
2. Upload the base VM image into Glance.
3. Upload the distribution ISO into Glance.
4. Create a Cinder volume from the distribution ISO.
5. Create an empty 20 GB Cinder volume for the Master VM.
6. Create one 80 GB/120 GB Cinder volume for each VM that will be supporting “control” services.
   - This will generally require three “control” Cinder volumes.
7. Create all the required Network entities, configured appropriately for the CPS Microservices platform.

Note Any config server addition or scale up must be done in maintenance window only as the operation may cause some timeouts in production.

VM Roles in CPS vDRA

The following table describes the VMs and their role in CPS vDRA:

Table 1: VM Names and Roles

VM Name	Role
Master	master
Control	control
DRA Director	dra-director
DRA Worker	dra-worker
Mongo DB (binding)	mongo-node - persistent router, Persistent DB.

Hardware Requirements for VMs

The following table describes the hardware requirements for binding and non-binding CPS vDRA VMs:

Table 2: Hardware Requirements

Role	CPU	RAM (GB)
master	8	32
control-0	8	32
control-1	8	32
dra-director	16	64
dra-worker	8	32
persistence-router	8	32
persistence-db	8	64

A standard deployment of CPS vDRA includes the following VMs:

- For vDRA VNF: one cluster manager, two management VMs, nine DRA Directors, eight DRA Workers.
- For Binding vDRA VNF: one cluster manager, two management VMs, five Persistent Routers, 12 DB VMs.

The number of directors, workers, routers, and database VMs varies, depending upon deployment needs.

Note vDRA supports the failure of one virtual machine in the set of master, control-0, and control-1 virtual machines. Therefore, these virtual machines cannot reside on the same ESXi Server.

Network Requirements For VMs

The following table describes the network requirements for each VM:

Table 3: Network Requirements for VMs

VM	Networks Required
Cluster Manager	Internal Management Replication
Management	Internal Management

VM	Networks Required
DRA Directors	Internal Management/Diameter Traffic VLAN
DRA Workers	Internal
Persistent Routers	Internal
Persistent Database	Internal Replication
Arbiter	Internal Management Replication

Protocols and Port Ranges

For each tenant in OpenStack, configure a security group and the protocol with a port range.

The following table describes the port range for each protocol:

Table 4: Port Ranges for Protocols

Protocol	Port Range
TCP	22 (SSH)
TCP	80 (HTTP)
TCP	443 (HTTPS)
TCP	636
TCP	2024
TCP	2375-2376
TCP	3375-3376
TCP	3868
TCP	5000
TCP	5001
TCP	5003

Protocol	Port Range
TCP	6443
TCP	6783
TCP	7443
TCP	7946
TCP	8000
TCP	8008
TCP	8080
TCP	8400
TCP	8500
TCP	8888
TCP	9100
TCP	9210
TCP	9212
TCP	9213
TCP	9443
TCP	12375
TCP	27017-27047
UDP	4789
UDP	6783-6784
UDP	7946
TCP	12375

Launching VMs

To bring up the system you must launch the VMs.

CPS vDRA includes the following two types of VMs:

- Master VM : hosts critical central system services such as the Registry and the Orchestrator.
- Docker engine VM (also sometimes called a Worker)

All VMs are launched with cloud-init configuration injected into them. This is done with the “config drive” mechanism in OpenStack, and it conveys all the information needed for the VM to start successfully.

Launch Master VM

The Master VM is the main VM for the whole system and requires a cloud-init configuration file.

Some of the configuration details in a cloud-init configuration are:

- Users: non-root administrative user details for the ‘cps’ user that is created and assigned to the ‘docker’ group. The user also has an SSH key injected. The SSH key is user-supplied and should be unique for every installation.
- Password: password for the ‘cps’ user. The password may be removed for all production installations.
- Default path for cps.pem is /etc/puppet/modules/qps/templates/certs
- Configuration details in JSON format (swarm.json) for the system to start successfully.
- Informational file: details of the product and VM within that product they are working with when they log in via SSH.

For an example of a cloud-init configuration file for a Master VM, see Master Cloud Init Example, on page 13.

swarm.json

The following table describes the configuration information that is included in a swarm.json file:

Table 5: JSON Parameters

Parameters	Description
role	Defines the “role” that this VM will be playing in the system. The role is used to map into the “deployment plan” to select the “scheduling slots” that are available on this particular VM.
identifier	This value combines the “role” with an “index number” on a per-role basis to create a unique identifying string for the deployed VM.
init	Name of the product-specific “initialization image”. This is the image to be launched by the boot scripts when the platform is established and it is time to start the application.

Parameters	Description
deployment_name	Name specific to a given deployment instance, used to identify the specific deployment, useful in the event that multiple instances of a given product (eg: PCRF) are deployed in the same environment. (This is also used by ESC.)
master	IP address of the Master VM, and must be communicated to all VMs
network	CIDR for the internal network.
registry	IP and port for the cluster-internal Docker registry. This is used by VMs to access the images deployed for the system.
scheduler	Style of scheduler to be used by the Orchestrator. Acceptable values are “ha” and “aio”, which launch the Orchestrator in High-Availability or All-In-One modes respectively. This defaults to “ha” if no orchestrator is specified.
reinitialize_data	Flag which, if set to 1, erases all existing data (primarily on the ‘control’ VMs) so that the system is started fresh. In most normal non-development situations this would be set to 0.
zing	Indicates whether, for those VMs where it is appropriate, the Azul Zing JVM should be used rather than the normal JVM.
weavePw	Password used to encrypt the Weave traffic.
tenant	(ESC only) Identifies the OpenStack tenant under which the system is being launched
esc	(ESC only) IP address of the ESC VM.
escUser	(ESC only) Username required to access the ConfD data on the ESC VM.
escPw	(ESC only) Password required to access the ConD data on the ESC VM.

Nova Boot Command

The Nova boot command to launch the Master VM (in this example) looks like the following:

```
nova boot --config-drive true --user-data=node-master-0.cfg \
  --flavor=cps.medium --image=docker-host-1.13.1-2 \
```

```
  --nic net-id="$Internal,v4-fixed-ip=172.16.2.11" \
  --nic net-id="$Management,v4-fixed-ip=172.18.11.121" \
  --block-device id=${iso},source=volume,dest=volume,device=/dev/vdb \
  --block-device id=${volume},source=volume,dest=volume,device=/dev/vdc \
  --availability-zone nova \
  --security-groups esc-security-group \
  docker-pcrf-master-0
```

The variables in the command are:

- $Internal: The OpenStack UUID of the “Internal” network.
- $Management: The OpenStack UUID of the “Management” network.
- $iso: The OpenStack UUID of the Cinder volume which contains the deployment ISO data.
- $volume: The OpenStack UUID of the Cinder volume to be used for persistent storage.

Launch Engine VM

The launching of a Docker engine VM (aka: Worker or Engine) is identical, though with less information injected into the ‘swarm.json’ file. Worker VMs might or might not have fixed IP addresses, and there is an assumption that if there isn’t a need for a fixed address that the interface will get it’s address via DHCP. This is primarily true for interfaces on the “internal” network, but there is no prohibition against doing it on other networks.

An example of the cloud-init file for an example Worker VM can be seen in Worker Cloud Init Example, on page 14. The values included in the worker ‘swarm.json’ file are:

- role
- identifier
- master
- network

In the simplest case a VM is launched exclusively on the “internal” network, assumed to DHCP, and has no Cinder volumes. The Nova boot command for that situation looks like:

```
nova boot --config-drive true --user-data=node-pcrf-0.cfg \
  --flavor=cps.medium --image=docker-host-1.13.1-2 \
  --nic net-id="$Internal" \
  --availability-zone nova \
```

Deployment Matrix

The following tables describe the minimum deployment for the DRA and DRA Binding applications.

The instances are the default choices for a minimal deployment. For those instances that have “scalable” set to YES, new instances can be launched to additional capacity. Where Cinder volumes are indicated, it is the case that each instance will have its own Cinder volume as defined in the table.

DRA

Table 6: Minimum Deployment for DRA

Role	Instances	Cinder Volumes	Scalable
master	master-0	/mnt/iso - /dev/vdb /mnt/install - /dev/vdc /mnt/install - /dev/vdd	NO
control	control-0 control-1	/data - /dev/vdb	NO
diameter-endpoint	diameter-endpoint-0 diameter-endpoint-1	None	YES
binding	binding-0 binding-1	None	YES

DRA Binding

Table 7: Minimum Deployment for DRA Binding - Microservices Platform

Role	Instances	Cinder Volumes	Scalable
master	master-0	/mnt/iso - /dev/vdb /mnt/install - /dev/vdc /mnt/install - /dev/vdd	NO
control	control-0 control-1	/data - /dev/vdb	NO
mongo-node	router-0 router-1	None	YES
mongo-node	db-app-0 db-app-1 db-app-2	None	YES

Startup Sequence

All VMs can be started at any time, in any order. The Worker VMs pause until the Master VM is ready to field their requests. Therefore, no matter the mechanism or VNFM used to launch the system ordering is not an issue. During the boot process normal OpenStack bootup activities occur. This includes all boot-time initialization, the execution of cloud-init, and the triggering of SystemD units. Those are all standard activities and the description of the startup sequence will begin upon the execution of the SystemD unit which is responsible for kicking off the CPS Microservices system. That unit is named “cpsinstall.service” and is staged to occur at the proper time in the Linux startup activities.

The “cpsinstall.service” is extremely simplistic, and performs one operation. It invokes the “/root/bootstrap.sh” script. This is true for Worker VMs as well as for the Master VM. Reading through this script will show you that it performs the following actions:

- Reads the values from the ‘swarm.json’ and sets them as environment variables
- Sets the queuing for RPS
- Determines if the VM is a Master VM, or Worker
- If the VM is a Master, do the following:
  - Mount the ISO from /dev/vdb
  - Mount the persistent Cinder volume from /dev/vdc
  - Perform ‘fsck’ and ‘resize2fs’ operations on the Cinder volume
  - Invoke the ‘install-master.py’ script
- If the VM is a Worker, do the following:
  - Wait for the Master to indicate it is alive
  - Mount the persistent Cinder volume from /dev/vdb
  - Perform ‘fsck’ and ‘resize2fs’ operations on the Cinder volume
  - Pull the ‘install-worker.py’ script from the Master with curl
  - Invoke the ‘install-worker.py’ script

The two install Python scripts perform the various install/startup operations for their VMs. Each of those actions will be detailed in its own section.

install-master.py

This script is specific to only the Master VM, and performs tasks that put the Master in play. The tasks will be outlined below as a sequence of bullet points.

- Condition the target system
  - If there is a /dev/vdc device first determine if it is formatted, and if not format the volume to be an EXT4 filesystem.

    - Remove all data from the mounted file system (/mnt/install) because the only time this runs on a Master VM is when it is booting, and that always calls for a fresh installation.
    - Copy all ISO data to the writable filesystem (/mnt/install).
    - Create the 'cps-app' system user.
- Load the configuration from the '/root/swarm.json' file.
    - During loading combine the values from '/mnt/iso/release.json' file
- Write the updated configuration to '/mnt/install/swarm.json'
- Stop the timesync daemon.
- Configure the Docker Engine daemon appropriately for operation on the Master.
- Restart the Docker service.
- Set the VM's hostname to match the VM's system-internal identifier.
- Install a locally-provided image from the saved image file 'cps-registry-latest.tgz'.
- Launch the registry container.
- Launch the registry upgrade container.
- Launch Weave.
- Launch the Engine Proxy container.
- Launch the Personality service.
- Launch the product-specific initialization container identified by the 'init' value in 'swarm.json'.

At this point control for the behavior of the system passes out of the startup processes and into the hands of the initialization container. For now, this ultimately leads to the invocation of the product specific Orchestrator. It is the Orchestrator that is responsible for launching all the remaining containers which make up the deployed application.

install-worker.py

This script is run for each of the Worker VMs that are launched. Operations on the Worker are very different from the Master, and that will be reflected in the steps detailed below.

- Load configuration information from the '/root/swarm.json' file for the VM.
- Read the release information from the Master VM, and incorporate it into an updated '/root/swrm.json' file.
- Determine the interface and IP for the Internal network
- If an 'identifier' value is not set in the provided '/root/swarm.json' file, create a unique engine identifier value.
- Create the 'cps-app' system user.
- If the role is 'control' then :

    - If there is a /dev/vdc device first determine if it is formatted, and if not format the volume to be an EXT4 filesystem.
    - Mount /dev/vdc on /data.
    - If 'reinitialize_data' is set, clear out the /data filesystem.

- If 'zing' is set and the role is 'policy', enable Zing operation.
- Set the hostname to match the 'identifier'.
- Stop the timesync daemon.
- Configure the Docker engine to include information about this Worker instance (it's identifier and it's network interface) so that the Orchestrator can know how to use the Worker in the system.
- Launch Weave

At this point the Worker VM is ready to accept containers. This will conclude the low-level startup process. From this point forward additional startup is handled by the Orchestrator and can vary from product to product.

External Port Matrix

The following table lists the services and ports are available to external users and applications in CPS vDRA .

It is recommended that connectivity to these ports be granted from the appropriate networks that require access to the below services.

Service	Port
Ssh	22/tcp
https	443/tcp
xinuexpansion4	2024/tcp
Upnp	5000/tcp
http	8008/tcp
http-proxy	8080/tcp

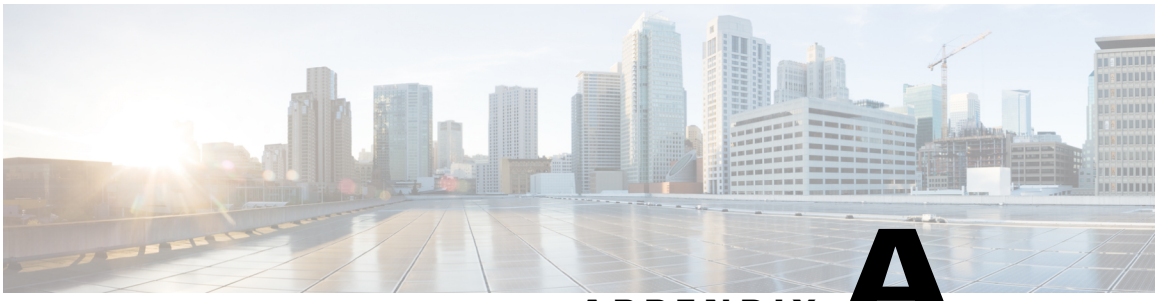

APPENDIX A

INIT Examples

Onboarding CPS vDRA on OpenStack Using ESC

Before configuring ESC to launch CPS vDRA VNF, ensure that below requirements are fulfilled.

- The qcow2 image has been onboarded on OpenStack.
- Identify the names of deployment in data model, VM groups, flavors and image name.
- Network – Identify the network (internal, external, management, IPs, subnets, interface IDs, DHCP pool etc.) that are going to be used for various VMs, and make sure that these networks are configured in OpenStack.
- The configuration data – Identify the initial configuration for a given type of VM, and nature of the VM (e.g. master, control etc.). Also, from where the configuration is to be loaded.

Master Cloud Init Example

```
#cloud-config
debug: True
output: {all: '| tee -a /var/log/cloud-init-output.log'}

users:
  - name: cps
    sudo: ['ALL=(ALL) NOPASSWD:ALL']
    groups: docker
    ssh-authorized-keys:
      - ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQDzjJjndIvUiBt
a4VSIbd2gJmlMWcQ8wtejgAbiXtoFZdtMdo9G0ZDEOtxHNNDPwWujMiYAkZh
ZWX/zON9raavU8lgD9+YcRopWUtujIC71YjtoxIjWIBBbrtqtPlUXMUXQsi91
RQbUtslENP+tSatS3awoQupyBMMSutyBady/7Wq0UTwFsnYs5Jfs8jIQuMfV
Q9uJ4mNn7wJ0N+Iaf27rE0t3oiY5DRN6j07WhauM6lCnZ1JDlzqmTnTHQkg
J3uKmQa5x73tJ1OW89Whf+R+dfslVn/yUwK/vf4extHTn32Dtsxkjz7kQeED
gCe/y7owimaEFcCIfEWEaj/50jegN cps@root-public-key
chpasswd:
  list: |
```

```
    cps:cisco123
  expire: False
write_files:
  - path: /root/swarm.json
    content: |
      {
        "role": "master",
        "identifier": "master-0",
        "master": "172.16.2.11",
        "network": "172.16.2.0/24",
        "registry": "172.16.2.11:5000",
        "consul": "172.16.2.11:8500",
        "reinitialize_data": "1",
        "zing": "1",
        "tenant": "esc",
        "weavePw": "cisco123",
        "esc": "172.18.11.122",
        "escUser": "admin",
        "escPw": "admin",
        "scheduler": "ha",
        "deployment_name": "docker-pcrf",
        "init": "cisco-policy-pcrf/init",
        "cluster_id": "cluster-1",
        "system_id": "system-1"

      }
    owner: root:root
    permissions: '0644'
  - path: /etc/update-motd.d/20-cps-text
    content: |
      #!/bin/sh
      product=`jq ".product" /mnt/install/swarm.json  | tr -d '"'`
      identifier=`jq ".identifier" /mnt/install/swarm.json  | tr -d '"'`

      printf "\n"
      printf " * CPS Microservices - ${product}\n"
      printf " * CPS Docker Engine - ${identifier}\n"
      printf "\n"
    owner: root:root
    permissions: '0755'
```

Worker Cloud Init Example

```
#cloud-config
debug: True
output: {all: '| tee -a /var/log/cloud-init-output.log'}

users:
  - name: cps
    sudo: ['ALL=(ALL) NOPASSWD:ALL']
    groups: docker
    ssh-authorized-keys:
      - ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQDzjJjndIvUi
Bta4VSIbd2gJmlMWcQ8wtejgAbiXtoFZdtMdo9G0ZDEOtxHNNDPwWujMi
YAkZhZWX/zON9raavU8lgD9+YcRopWUtujIC71YjtoxIjWIBBbrtqtPlU
XMUXQsi91RQbUtslENP+tSatS3awoQupyBMMSutyBady/7Wq0UTwFsnYs
5Jfs8jIQuMfVQ9uJ4mNn7wJ0N+Iaf27rE0t3oiY5DRN6j07WhauM6lCnZ
1JDlzqmTnTHQkgJ3uKmQa5x73tJ1OW89Whf+R+dfslVn/yUwK/vf4extHT
n32Dtsxkjz7kQeEDgCe/y7owimaEFcCIfEWEaj/50jegN
cps@root-public-key
chpasswd:
  list: |
    cps:cisco123
  expire: False
write_files:
  - path: /root/swarm.json
    content: |
      {
        "role": "pcrf",
        "identifier": "pcrf-0",
```

```
      "master": "172.16.2.11",
      "network": "172.16.2.0/24",
      "registry": "172.16.2.11:5000",
      "consul": "172.16.2.11:8500",
      "reinitialize_data": "1",
      "zing": "1"
    }
  owner: root:root
  permissions: '0644'
- path: /etc/update-motd.d/20-cps-text
  content: |
    #!/bin/sh
    product=`jq ".product" /root/swarm.json  | tr -d '"'`
    identifier=`jq ".identifier" /root/swarm.json  | tr -d '"'`

    printf "\n"
    printf " * CPS Microservices - ${product}\n"
    printf " * CPS Docker Engine - ${identifier}\n"
    printf "\n"
  owner: root:root
  permissions: '0755'
```