



CloudCenter Suite 5.3 Documentation

First Published: December 10, 2021

Last Modified: December 23, 2021

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

1. Install 5.3 Home	2
1.1 Suite Installer 5.3 Home	3
1.1.1 Suite Admin 5.3.0 Release Notes	4
1.1.1.1 Suite Admin 5.3.1 Release Notes	14
1.1.2 Suite Architecture	23
1.1.3 Self-Hosted Installation	25
1.1.3.1 Installer Overview	26
1.1.3.2 Installer Virtual Appliances	27
1.1.3.2.1 Virtual Appliance Overview	28
1.1.3.2.2 OpenStack Appliance Setup	30
1.1.3.2.3 VMware vSphere Appliance Setup	33
1.1.3.3 Prepare Infrastructure	43
1.1.3.4 New Cluster Installation	46
1.1.3.4.1 VMware vSphere Installation	47
1.1.3.4.2 OpenStack Installation	71
1.1.3.5 Existing Cluster Installation	76
1.1.3.6 Upgrade Kubernetes Cluster	80
1.1.3.6.1 Upgrade Approach	81
1.1.3.6.2 OpenStack Upgrade	91
1.1.3.6.3 VMware vSphere Upgrade	97
1.1.3.7 Air Gap Installation	102
1.1.3.8 Upgrade Offline Repository	110
1.1.3.9 Backup and Restore	114
1.1.3.9.1 Public Cloud	115
1.1.3.9.2 Private Cloud	154
1.1.3.10 Troubleshooting	172
1.1.4 Suite Admin Workflow	177
1.1.5 Initial Administrator Setup	179
1.1.6 Kubernetes Cluster Management	181
1.1.6.1 Cluster Status	182
1.1.6.2 Manage Clusters	184
1.1.7 Configure Smart Licenses	186
1.1.8 Module Lifecycle Management	196
1.1.8.1 Install Module	197
1.1.8.2 Update Module	202
1.1.8.3 Monitor Modules	207

Suite Installer 5.3 Home

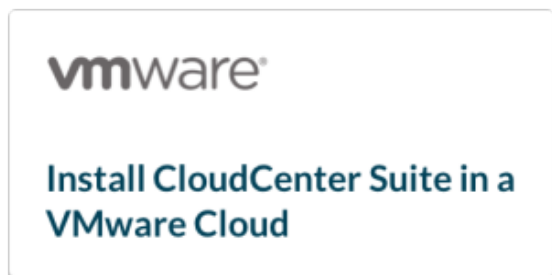
Self-Hosted 5.3 Documentation

System Announcements

Cisco released the following Suite Admin releases:

- [Suite Admin 5.3.0](#) released on Dec 10, 2021

Installation Process



Search Suite Installer 5.3 Documentation

The selected root page could not be found.

Recent Updates

[Private Cloud](#)
updated Nov 08, 2021
[view change](#)
[Private Cloud](#)
updated Jul 15, 2021
[view change](#)
[Suite Installer 5.2 Home](#)
updated Apr 16, 2021
[view change](#)

Back to: [CloudCenter Suite Home](#)

Suite Admin 5.3.0 Release Notes

Suite Admin 5.3.0 Release Notes

- [Release Date](#)
- [Helm Chart Upgrade from 5.2.3 to 5.2.4](#)
- [Architecture](#)
- [Public Clouds](#)
- [Administration](#)
- [Module Management](#)
- [Smart Software Licensing](#)
- [Suite Admin Dashboard](#)
- [User Tenant Management](#)
- [Cluster Management](#)
- [Security Management](#)
- [Suite UI](#)
- [Deprecated](#)
- [API](#)
 - [New API Calls](#)
 - [Updated API Calls](#)
- [Documentation](#)
- [Known Issues](#)
- [Resolved Issues](#)

Release Date

First Published: December 10, 2021

Updating Modules

CloudCenter Suite5.3.0 is only supported as an upgrade from Suite Admin 5.2.4. Please ensure that you follow the procedural steps below to correctly update your CloudCenter installation.

1. You must upgrade the Suite Admin module to version 5.3.0 before you upgrade any other CloudCenter Suite module.
2. After the Suite Admin upgrade has completed, upgrade the other CloudCenter modules to a supported configuration.
 - a. Workload Manager required version - 5.5.1
 - b. Cost Optimizer required version - 5.5.1
 - c. Action Orchestrator required version - 5.2.4
3. After all CloudCenter modules have been upgraded, follow the directions below to upgrade your Kubernetes cluster.

Note

Before updating any module, verify that you have twice the required CPU/Memory in your cluster. A module-update scenario requires additional resources for the old pod to continue running until the new pod initializes and takes over. This additional resource requirement is temporary and only required while a module update is in progress. After the module is updated, the additional resources are no longer needed.

Update only one module at a time. If you simultaneously update more than one module, your update process may fail due to limited resource availability. See [Prepare Infrastructure](#) for additional context.

Backup and Restore

There are two ways to upgrade the Kubernetes cluster. Ensure you take backup of the Kubernetes cluster before upgrading the cluster.

1. Upgrade the existing Kubernetes cluster or
2. Spin up a new Kubernetes cluster with the Suite 5.3.0 installer. On this new cluster [restore](#) the backup.

Please follow the instructions documented here to [backup](#) before upgrading the cluster.

Kubernetes Cluster Upgrade

Customers running CCS v5.2.4 can now upgrade their Kubernetes cluster using v5.3.0 installer. However they need to follow a few required steps to start upgrading the Kubernetes version to 1.18.12.

1. IMP: Follow documented steps to take [a backup of the cluster](#) before proceeding with the kubernetes upgrade.
2. Upgrade Suite Admin version to v5.3.0 using the UI. This is a required step to continue the upgrade, if Suite Admin is not upgraded users will see this error message on the installer: "
Cluster is NOT running Common-Framework Suite Admin v5.3.x. Please upgrade the Suite Admin chart to v5.3.x before upgrading kubernetes."
3. **(Only on vSphere)** This release addresses the security issue of encrypting ETCD secrets in Kubernetes. To do this users need to run the documented workaround script to enable ETCD Encryption of secrets in the kubernetes clusters. New cluster installed using v5.3.0 installer has ETCD encryption enabled by default. The contents of the `etcdEncrypt_master_1.sh` script is shown [here](#).

Important: File name of the script in the first master must be `etcdEncrypt_master_1.sh`

```
#!/bin/bash

# if encryption is not enabled, enable it.
# Because etcd disk is shared between old and new master node,
# we have to turn on encryption on the old master, and encrypt
# all secrets first. After this, when new master node
# is booted, its own etcd will have encrypted data and kube api
# process will then use encryption key to decrypt secrets stored
# in etcd.
function check_if_running_sudo {
    if [ "$EUID" -ne 0 ]
    then echo "This script needs to run as sudo. Please retry with:"
    echo "sudo bash $0"
    exit
    fi
}

function check_if_vsphere_cloud {
    if kubectl get cm k8s-mgmt.cloudaccount --kubeconfig=/home/${users}/.kube/config -n cisco ; then
        echo "continue with the upgrade..."
    else
        echo "Workaround is not required for this cloud. Please continue with the upgrade from installer UI."
        exit
    fi
    kubectl get cm k8s-mgmt.cloudaccount --kubeconfig=/home/${users}/.kube/config -n cisco -o jsonpath="{.data.data}" | base64 -d > k8s-mgmt.cloudaccount.json
    sleep 60
    CLOUD_TYPE=$(cat k8s-mgmt.cloudaccount.json | jq -r '.cloudType')
    if [[ $CLOUD_TYPE == *"vsphere"* ]]; then
        echo "Vshere cloud, continue with the workaround."
    else
        echo "$CLOUD_TYPE cloud, workaround is not required for this cloud. Please continue with the upgrade from installer UI"
        exit
    fi
}

function check_for_encryption_file {
    #Function to check if this file is executed only Once, if first time, create a lock
    FILE=/etc/kubernetes/pki/etcd/encryption.conf
    if [ -f "$FILE" ]; then
        echo -e "-----"
        echo "$FILE already exists.
        Looks like workaround is already performed on this master node.
        Please do not run the script multiple times on same master. Continue with other master nodes."
        echo -e "-----"
        exit
    fi
}

function backup_resources {
    mkdir -p backup && cd backup || exit
    echo -e "Backing up all the certificates to backup folder $PWD /backup"
    kubectl get -o yaml --all-namespaces issuer,clusterissuer,certificates,secrets --kubeconfig=/home/${users}/.kube/config > cert-manager-backup.yaml
    kubectl get cm k8s-mgmt.cloudaccount --kubeconfig=/home/${users}/.kube/config -n cisco -o jsonpath="{.
```

```

data.data}" | base64 -d > k8s-mgmt.cloudaccount.json
echo -e "Backing up all the secrets to resources_backup.yaml"
kubectl get -o yaml --all-namespaces secrets --kubeconfig=/home/$(users)/.kube/config >
resources_backup.yaml
for n in $(kubectl get -n cisco -o=name pvc,certificate,configmap,serviceaccount,secret,ingress,
service,deployment,statefulset,hpa,job,cronjob)
do
mkdir -p $(dirname "$n")
kubectl get -n cisco --kubeconfig=/home/$(users)/.kube/config -o=yaml --export "$n" > "$n".yaml 2>
/dev/null
done
cd ..
echo -e "-----"
}

ETCD_ENCRYPTION_SECRET=$(head -c 32 /dev/urandom | base64)

function generate_encryption_conf_file {
#!/bin/bash

OUTPUT_FILE=/etc/kubernetes/pki/etcd/encryption.conf

echo -e "\nCreating ETCD_ENCRYPTION_SECRET and ETCD_ENCRYPTION_KEY"

ETCD_ENCRYPTION_SECRET=$(head -c 32 /dev/urandom | base64)
ETCD_ENCRYPTION_KEY=$(echo ccp-key-$(echo "$ETCD_ENCRYPTION_SECRET" | base64 -d | sha256sum | cut -c1-
10))

echo "Writing content to file $OUTPUT_FILE"

cat <<EOF > $OUTPUT_FILE
apiVersion: apiserver.config.k8s.io/v1
kind: EncryptionConfiguration
resources:
- resources:
- secrets
providers:
- aescbc:
keys:
- name: $ETCD_ENCRYPTION_KEY
secret: $ETCD_ENCRYPTION_SECRET
- identity: {}
EOF
}

check_if_running_sudo
check_if_vsphere_cloud
check_for_encryption_file
backup_resources
echo -e "\nCreating etcd Encryption configuration file for all the master nodes"
generate_encryption_conf_file
echo -e "\nSaved etcd Encryption configuration file at $OUTPUT_FILE"
cp $OUTPUT_FILE encryption.conf
echo -e "\nSaved etcd Encryption configuration file's copy at encription.conf\n\n"

#-----
#-----

function check_if_running_sudo {
if [ "$EUID" -ne 0 ]
then echo "This script needs to run as sudo. Please retry with:"
echo "sudo bash $0"
exit
fi
}

function sleep_function {
echo "Sleeping for $1 minutes..."
while true;do echo -n .;sleep 1;done &
sleep "$1" # or do something else here
}

```

```

kill $!; trap 'kill $!' SIGTERM
echo "Done"
echo -e "\n\n-----"
}
function check_if_vsphere_cloud {
    if kubectl get cm k8s-mgmt.cloudaccount --kubeconfig=/home/${users}/.kube/config -n cisco ; then
        echo "continue with the upgrade..."
    else
        echo "Workaround is not required for this cloud. Please continue with the upgrade from installer UI."
        exit
    fi
    kubectl get cm k8s-mgmt.cloudaccount --kubeconfig=/home/${users}/.kube/config -n cisco -o jsonpath="{.data.data}" | base64 -d > k8s-mgmt.cloudaccount.json
    sleep 60
    CLOUD_TYPE=$(cat k8s-mgmt.cloudaccount.json | jq -r '.cloudType')
    if [[ $CLOUD_TYPE == *"vsphere"* ]]; then
        echo "Vshere cloud, continue with the workaround."
    else
        echo "$CLOUD_TYPE cloud, workaround is not required for this cloud. Please continue with the upgrade from installer UI"
        exit
    fi
}

function check_for_encryption_file {
    #Function to check if this file is executed only Once, if first time, create a lock
    FILE=/etc/kubernetes/pki/etcd/encryption.conf
    if [ -f "$FILE" ]; then
        echo -e "-----"
        echo "$FILE already exists."
        Looks like workaround is already performed on this master node.
        Please do not run the script multiple times on same master. Continue with other master nodes."
        echo -e "-----"
        exit
    fi
}

check_if_running_sudo
check_if_vsphere_cloud

OUTPUT_FILE=/etc/kubernetes/pki/etcd/encryption.conf

function copy_encryption_conf_file {
    #!/bin/bash

    OUTPUT_FILE=/etc/kubernetes/pki/etcd/encryption.conf
    echo "Copying script to $OUTPUT_FILE"
    cat <<EOF > $OUTPUT_FILE
$(sed -n -e '/^ETCD_ENCRYPTION_FILE_START$/ ,/^ETCD_ENCRYPTION_FILE_END$/ { /^ETCD_ENCRYPTION_FILE_START$/d; /^ETCD_ENCRYPTION_FILE_END$/d; p; }' "${0}")
EOF
cp $OUTPUT_FILE encryption.conf
echo -e "\nSaved etcd Encryption configuration file's copy at encription.conf\n\n"
}

# File content which updates kube-apiserver.yaml and waits for kubserver to start
function check_kube_api {
    API_PID=$(ps -ef | awk '/kube-ap[i]server/{print $2}')
    echo -e "-----"
    echo "check_kube_api to verify that kube-apiserver is restarted: $API_PID ..."
    while [ $API_PID = $1 ]; do
        echo "wait for kube api server to exist.."
        sleep_function 2m
        API_PID=$(ps -ef | awk '/kube-ap[i]server/{print $2}')
    done

    while ! ps -ef | grep kube-ap[i]server; do
        echo "wait for kube api server to restart.."
        sleep_function 2m
    done
}

```



```

echo -e "\n=====
echo -e "Waiting. Please DO NOT INTERRUPT."
sleep 300
echo -e "Successfully updated kubernetes manifests"
echo -e "\n=====
}

function update_kube_apiserver_manifest {
    API_PID=$(ps -ef | awk '/kube-ap[is]server/{print $2}')

    echo -e "Adding ETCD encryption resources to kube-apiserver, it might take some time. Please DO NOT
    INTERRUPT."

    API_CONF="/etc/kubernetes/manifests/kube-apiserver.yaml"
    ENC_CONF="/etc/kubernetes/pki/etcd/encryption.conf"

    sudo python - <<EOF
import yaml
name = 'k8s-encryption'

with open("$API_CONF", "r") as f:
    data = yaml.safe_load(f)
    data['spec']['volumes'].append({'hostPath': {'path': "$ENC_CONF", 'type': ''}, 'name': name})
    data['spec']['containers'][0]['command'].append("--encryption-provider-config=$ENC_CONF")
    data['spec']['containers'][0]['volumeMounts'].append({'mountPath': "$ENC_CONF", 'name': name})

with open("$API_CONF", "w") as f:
    f.write(yaml.dump(data, default_flow_style=False, indent=2))
EOF

    sleep_function 2m
}

function patch_kubeconfig_secret {

    KUBECONFIG_SECRET_NAME=$(kubectl get secret -n ccp --kubeconfig=/home/$(users)/.kube/config | grep
    kubeconfig | awk '/-/{print $1}')
    echo "\nCreating patch file for kubeconfig secret $KUBECONFIG_SECRET_NAME"
    echo "{\"data\":{\"etcdEncryptionKey\": \"${ETCD_ENCRYPTION_SECRET}\"}}\" > patch_kubeconfig_secret.
    json
    kubectl -n ccp patch secret "$KUBECONFIG_SECRET_NAME" --kubeconfig=/home/$(users)/.kube/config --patch
    "$(cat patch_kubeconfig_secret.json)"
    echo -e "\nSuccessfully patched kubeconfig secret, sleeping for 2 mins, please DO NOT INTERRUPT.\n\n"
    sleep_function 2m
    kubectl get secrets --all-namespaces -o json --kubeconfig=/home/$(users)/.kube/config | kubectl
    replace -f - --kubeconfig=/home/$(users)/.kube/config
    echo -e "\n\nUpdating the secret, please DO NOT INTERRUPT."
    sleep_function 5m

    echo -e
"\n\n=====
echo -e "=====
echo -e "\nWorkaround completed."
echo -e "\nRestart all 3 Master nodes one by one."
echo -e "\nGive significant time for all services to become RUNNING before proceeding with next master
node.(Recommended ~5 mins)"
echo -e
"\n\n=====
echo -e "=====
}

#-----
#-----

API_PID=$(ps -ef | awk '/kube-ap[is]server/{print $2}')
echo -e "\n\nPrinting encryption file, please copy it to other masters and execute\n"

```

```

echo ": " > etcdEncrypt_master_2.sh
echo ": " > etcdEncrypt_master_3.sh
echo "ETCD_ENCRYPTION_FILE_START" >> etcdEncrypt_master_2.sh
echo "ETCD_ENCRYPTION_FILE_START" >> etcdEncrypt_master_3.sh
cat $OUTPUT_FILE >> etcdEncrypt_master_2.sh
cat $OUTPUT_FILE >> etcdEncrypt_master_3.sh
echo "ETCD_ENCRYPTION_FILE_END" >> etcdEncrypt_master_2.sh
echo "ETCD_ENCRYPTION_FILE_END" >> etcdEncrypt_master_3.sh
echo " " >> etcdEncrypt_master_2.sh
echo " " >> etcdEncrypt_master_3.sh

sed -n '107,245p' etcdEncrypt_master_1.sh >> etcdEncrypt_master_2.sh
sed -n '107,245p' etcdEncrypt_master_1.sh >> etcdEncrypt_master_3.sh
echo "check_for_encryption_file" >> etcdEncrypt_master_2.sh
echo "check_for_encryption_file" >> etcdEncrypt_master_3.sh
echo "copy_encryption_conf_file" >> etcdEncrypt_master_2.sh
echo "copy_encryption_conf_file" >> etcdEncrypt_master_3.sh
echo "update_kube_apiserver_manifest" >> etcdEncrypt_master_2.sh
echo "update_kube_apiserver_manifest" >> etcdEncrypt_master_3.sh
echo -e "export -f check_kube_api \ntimeout 300s bash -c check_kube_api \"\$API_PID\" " >>
etcdEncrypt_master_2.sh
echo -e "export -f check_kube_api \ntimeout 300s bash -c check_kube_api \"\$API_PID\" " >>
etcdEncrypt_master_3.sh
echo -e "ETCD_ENCRYPTION_SECRET=\"\$ETCD_ENCRYPTION_SECRET\">>etcdEncrypt_master_3.sh
echo "patch_kubeconfig_secret" >> etcdEncrypt_master_3.sh

update_kube_apiserver_manifest
export -f check_kube_api
timeout 300s bash -c check_kube_api "$API_PID"

echo -e "\nWorkaround completed on Master-1, please continue with Master-2 and Master-3"
echo -e "\n\n======"
echo -e "\nCopy etcdEncrypt_master_2.sh to second master and run sudo bash etcdEncrypt_master_2.sh"
echo -e "\nCopy etcdEncrypt_master_3.sh to third master and run sudo bash etcdEncrypt_master_3.sh"
echo -e "======"

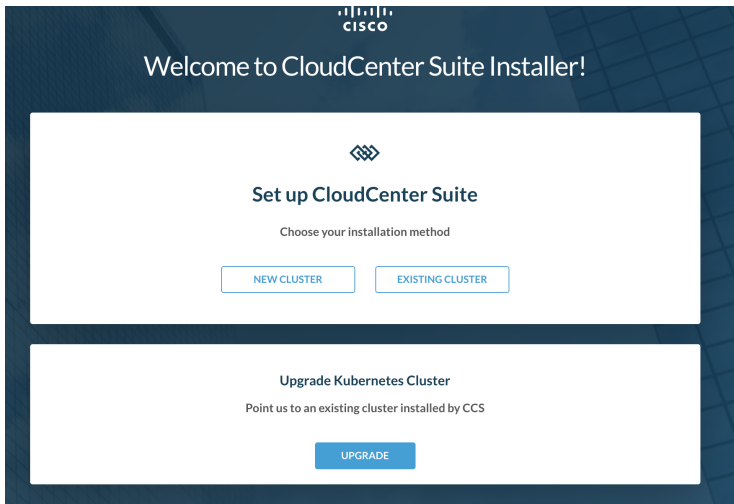
```

The script is mandatory, not optional for the upgrade to work.

IMPORTANT: DO NOT run the workaround script TWICE on any of the masters, this can bring the cluster since script interacts with kube-apiserver of the cluster.

4. Use v5.3.0 installer to upgrade the kubernetes cluster & cert-manager
5. After upgrading the cluster, users can use the "Take me to Suiteadmin" link to go back to the Suite Admin UI. Please note that this address might change for DHCP based clusters.

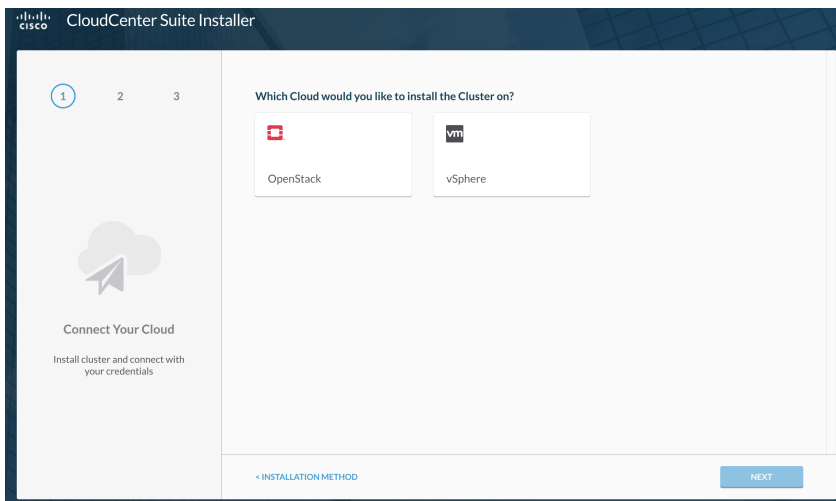
In the upgraded cluster, if you upgrade Workload-manager, Cost-optimizer, or Action Orchestrator and see this error on suite admin UI:
 "failed to wait for command /bin/helm, err: exit status 1, msg: Error: [unable to recognize "": no matches for kind "Certificate" in version "certmanager.k8s.io/v1alpha1", unable to recognize "": no matches for kind "Issuer" in version "certmanager.k8s.io/v1alpha1"], please follow workaround steps mentioned in [Update Module](#).



Kubernetes Cluster Installation

CCS Installer drops support for Public clouds (AWS, GKE & Azure) for new cluster installation, however, the "existing cluster" mode of installation is still supported. Users can bring a GKE/Azure based Kubernetes cluster and install suite admin using v5.3.0 installer. ([CPSGCORE-4305-Authenticate](#) to see issue details)

CCS v5.3.0 does not have installer images for public clouds. Only OVA and QCOV2 appliances are published.



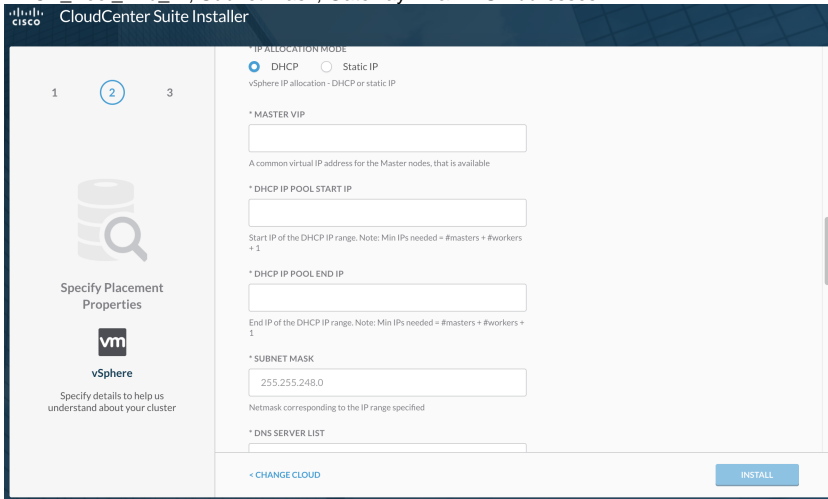
CCS Installer upgrades Supported Kubernetes version and cert-manager

- CCS v5.2.3 (Previous release) supported Kubernetes v1.16.3 and Cert-manger v0.10.1
- CCS v5.3.0 Upgraded Kubernetes to v1.18.12 and Cert-manger to v1.0.2

CCS Installer changes how DHCP based clusters are installed:

- CCS v5.2.3 required users to only provide only a master VIP to create a cluster, but

- CCS v5.3.0 now requires users to provide all the network information like static IP case to create a cluster & its resources. DHCP_Pool_Start_IP, DHCP_Pool_End_IP, Subnet mask, Gateway IP & DNS Addresses.



Helm Chart Upgrade from 5.2.3 to 5.2.4

If you are upgrading from SA 5.2.3 or older to 5.2.4, perform the following procedure after upgrade:

1. Run an SQL command in the suite-postgresqldatabase.
2. Log in to common-framework-suite-postgresql-0

Renew your license to continue

Your evaluation license has expired. Contact your administrator to renew your Reporting for Confluence license.

3. Run the following bash command.

```
PGDATABASE=suite-samlpsql-c"UPDATEpublic.saml_infra_configSETcert_source='cert_manager' WHERE id=1;"
```

4. Optionally regenerate the SSO certificate.

```
kubectldelate secret suite-saml-ss0-tls-n cisco
```

5. Restart the suite-samlpod.

```
kubectldelate pod suite-saml-pod -n cisco
```

6. If your customer uses SSO, reconfigure SSO.

Architecture

No updates

Public Clouds

The following public clouds with existing kubernetes clusters are supported:

- **Azure AKS:**
 - Supported Version: 1.20.9
- **Google GKE:**
 - Supported Version: 1.20.11
- See [Existing Cluster Installation](#) for additional details
- Kubernetes version upgrade on existing clusters from older version 1.13, 1.16.x is not supported for public clouds.

Administration

No updates

Module Management

No updates

Smart Software Licensing

No updates

Suite Admin Dashboard

No updates

User Tenant Management

No updates

Cluster Management

No updates

Security Management

No updates

Suite UI

No updates.

Deprecated

- CCS Installer drops support for Public clouds(AWS, GKE & Azure) for new cluster installation

API

New API Calls

Updated API Calls

No updates.

Documentation

The following documentation changes were implemented in CloudCenter Suite5.3.0:

- [Support for Two Cloud Types](#)
- [Support for New Kubernetes Version 1.18.x](#)
- [Restrictions on Containers](#)
- [New IP Allocation Mode](#)
- [Helm Chart Configuration after k8s Upgrade](#)
- [Update Backup](#)
- [SSO Configuration](#)
- [VMware Folder Name Change](#)
- [Cluster Upgrade](#)

Known Issues

CloudCenter Suite5.3.0 has the following known issues:

- When the Suite Admin chart is upgraded from version 5.2.4 to 5.3.0 in a few elasticsearch pods, note the following:
 - The busybox image is not upgraded from version 1.29.3 to 1.33.1.
 - The elasticsearch image is not upgraded from elasticsearch version 6.8.12 to elasticsearch version 6.8.13-cisco2.
 - Two elasticsearch client pods are running after performing a Suite Admin upgrade.

Resolved Issues

CloudCenter Suite 5.3.0 has the following resolved issues:

- **CSCvs03618:** Update base installer images to address multiple Vulnerabilities in ntp//snmp/libopts
- **CSCvs03618:** Apache Camel XML External Entity Injection Vulnerability. An Apache Camel XML External Entity Injection vulnerability appeared in versions of Suite Admin prior to version 5.3.0.
- **CSCvu73902:** After starting to upgrade from Suite Admin version 5.1 to 5.2, the upgrade begins, but halts after completing only 2 percent of the task because of Helm Chart version mismatch.
- **CSCvx68373:** After configuring CloudCenter Suite, the Suite Admin module with email settings using TLS shows the emails are not sent from Suite Admin when using password reset. option.
- **CSCwa44980:** A custom role name in Action Orchestrator displays with the "null" string incorrectly appended to the name.
- **CSCwa43511:** Action Orchestrator user interface. After clicking on the Uninstall Action Orchestrator button, the interface incorrectly opened in a new tab displaying a prompt to uninstall.
- **CSCwa44925:** Using the IPAM/DHCP allocation method, CloudCenter Suite cannot create a NodePort cluster.
- **NO CDETS:** Add support for upgrading DHCP vSphere clusters from CloudCenter Suite 5.2.4 kubernetes version 1.16 to CloudCenter Suite 5.3.0 kubernetes 1.18.12
- **CSCwa44974:** Google has removed Helm Chart version 2.16.3 when upgrading from version 2.16.3 to 2.16.12.
- **NO CDETS:** The walkme object is still present in Suite Admin, although it should be removed.
- **CSCvt19754:** A violation of an information disclosure agreement could allow an authenticated Elasticsearch user to improperly view details.
- **CSCwa44971:** When the reloader image used by Quicksilver was updated, an imagepullbackoff session occurred when the cluster tried to pull the new reloader image.
- **CSCvs03620:** Google Guava Eager Allocation Remote Denial of Service Vulnerability
- **CSCwa44964:** Unexpected stop actions occurred after performing the k8s-mgmt Anchore scans.
- **CSCwa44972:** A CloudCenter Suite cluster runs out of space on a disk with Action Orchestrator version 5.2.4 installed.
- **CSCwa44958:** Details of a previously created user incorrectly display in the creation form of another user in Suite Admin.
- **CSCwa44959:** When Suite Admin logs into the Suite Admin dashboard and enables the log archive to configure a storage location on AWS, the validator indicates only two fields are missing when several other fields are not displaying.
- **CSCwa44981:** When enabling a filter in the filter pane, the Suite Admin improperly indicates the filter has invalid criteria.
- **CSCwa44960:** libgrypt vulnerabilities.
- **CSCwa44961:** After logging in to the Suite Admin dashboard and enabling the log archive, you are unexpectedly unable to disable the archive. Also, the archive does not establish a valid connection to AWS.
- **CSCwa44973:** After setting the forge_personallInfoAcknowledged setting to false and clicking the Acknowledge button, an error is unexpectedly generated.
- **CSCwa44969:** When attempting to access public storage in a Suite Admin private cluster, the storage bucket configuration fails indicating the server could not find the resource.
- **CSCwa44962:** Address OpenSSL CVEs for Alpine and Ciscoj.
- **CSCwa44970:** When a Suite Admin installation failed for an existing cluster on Azure, the Start Over button does not redirect you to the installer Welcome page.
- **CSCwa44977:** The Download KubeConfig file option is incorrectly unlisted in the Suite Admin user interface.
- **CSCwa44963:** After clicking the Startover button to delete virtual machines on OpenStack, the installer user interface incorrectly goes to the main installer page instead of going directly to the cluster deletion page.
- **CSCwa4336/4343/4266/4258/4260/4277/4280/4286/4289/4290:** Minerva requires sw updates: suite-k8s-mgmt-5.2.4-RC1.0,api-6.1.1,busybox-1.29.3,kubectl-1.13.2,ccp-6.1.1,suite-gateway-5.2.4-RC1.6,suite-jwt-keys-5.2.4-RC1.6,suite-license-RC1.6,suite-res-mgmt-5.2.4-RC1.6,ui-5.2.4-RC1.0

Suite Admin 5.3.1 Release Notes

Suite Admin 5.3.1 Release Notes

- [Release Date](#)
- [Helm Chart Upgrade from 5.2.3 to 5.2.4](#)
- [Architecture](#)
- [Public Clouds](#)
- [Administration](#)
- [Module Management](#)
- [Smart Software Licensing](#)
- [Suite Admin Dashboard](#)
- [User Tenant Management](#)
- [Cluster Management](#)
- [Security Management](#)
- [Suite UI](#)
- [Deprecated](#)
- [API](#)
 - [New API Calls](#)
 - [Updated API Calls](#)
- [Documentation](#)
- [Known Issues](#)
- [Resolved Issues](#)

Release Date

First Published: December 23, 2021

Updating Modules

CloudCenter Suite5.3.1 is only supported as an upgrade from Suite Admin 5.2.4. Please ensure that you follow the procedural steps below to correctly update your CloudCenter installation.

1. You must upgrade the Suite Admin module to version 5.3.1 before you upgrade any other CloudCenter Suite module.
2. After the Suite Admin upgrade has completed, upgrade the other CloudCenter modules to a supported configuration.
 - a. Workload Manager required version - 5.5.1
 - b. Cost Optimizer required version - 5.5.1
 - c. Action Orchestrator required version - 5.2.4
3. After all CloudCenter modules have been upgraded, follow the directions below to upgrade your Kubernetes cluster.



Note: As Suite Installer 5.3.0 versions are deprecated avoid upgrade to 5.3.1 from 5.2.4

Before updating any module, verify that you have twice the required CPU/Memory in your cluster. A module-update scenario requires additional resources for the old pod to continue running until the new pod initializes and takes over. This additional resource requirement is temporary and only required while a module update is in progress. After the module is updated, the additional resources are no longer needed.

Update only one module at a time. If you simultaneously update more than one module, your update process may fail due to limited resource availability. See [Prepare Infrastructure](#) for additional context.

Backup and Restore

There are two ways to upgrade the kubernetes cluster. Ensure you take backup of the kubernetes cluster before upgrading the cluster.

1. Upgrade the existing kubernetes cluster or
2. Spin up a new kubernetes cluster with the Suite 5.3.1 installer. On this new cluster [restore](#) the backup.

Please follow the instructions documented here to [backup](#) before upgrading the cluster.

Kubernetes Cluster Upgrade

Customers running CCS v5.2.4 can now upgrade their Kubernetes cluster using v5.3.1 installer. However they need to follow a few required steps to start upgrading the Kubernetes version to 1.18.12.

1. IMP: Follow documented steps to take a [backup of the cluster](#) before proceeding with the kubernetes upgrade.
2. Upgrade Suite Admin version to v5.3.1 using the UI. This is a required step to continue the upgrade, if Suite Admin is not upgraded users will see this error message on the installer: "
Cluster is NOT running Common-Framework Suite Admin v5.3.x. Please upgrade the Suite Admin chart to v5.3.x before upgrading kubernetes."
3. **(Only on vSphere)** This release addresses the security issue of encrypting ETCD secrets in Kubernetes. To do this users need to run the documented workaround script to enable ETCD Encryption of secrets in the kubernetes clusters. New cluster installed using v5.3.1 installer has ETCD encryption enabled by default. The contents of theetcdEncrypt_master_1.shscript is shown below and it needs to be run on the first kubernetes master node. Post that follow the instructions from the output of script execution. Keep the filenames as mentioned in this document.

```
#!/bin/bash

# if encryption is not enabled, enable it.
# Because etcd disk is shared between old and new master node,
# we have to turn on encryption on the old master, and encrypt
# all secrets first. After this, when new master node
# is booted, its own etcd will have encrypted data and kube api
# process will then use encryption key to decrypt secrets stored
# in etcd.
function check_if_running_sudo {
if [ "$EUID" -ne 0 ]
then echo "This script needs to run as sudo. Please retry with:"
echo "sudo bash $0"
exit
fi
}

function check_if_vsphere_cloud {
if kubectl get cm k8s-mgmt.cloudaccount --kubeconfig=/home/$(users)/.kube/config -n cisco ; then
echo "continue with the upgrade..."
else
echo "Workaround is not required for this cloud. Please continue with the upgrade from installer UI."
exit
fi
kubectl get cm k8s-mgmt.cloudaccount --kubeconfig=/home/$(users)/.kube/config -n cisco -o jsonpath="{.data.data}" | base64 -d > k8s-mgmt.cloudaccount.json
sleep 60
CLOUD_TYPE=$(cat k8s-mgmt.cloudaccount.json | jq -r '.cloudType')
if [[ $CLOUD_TYPE == *"vsphere"* ]]; then
echo "Vshere cloud, continue with the workaround."
else
echo "$CLOUD_TYPE cloud, workaround is not required for this cloud. Please continue with the upgrade from installer UI"
exit
fi
}

function check_for_encryption_file {
#Function to check if this file is executed only Once, if first time, create a lock
FILE=/etc/kubernetes/pki/etcd/encryption.conf
if [ -f "$FILE" ]; then
echo -e "-----"
echo "$FILE already exists.
Looks like workaround is already performed on this master node.
Please do not run the script multiple times on same master. Continue with other master nodes."
echo -e "-----"
exit
fi
}

function backup_resources {
mkdir -p backup && cd backup || exit
echo -e "Backing up all the certificates to backup folder $PWD /backup"
kubectl get -o yaml --all-namespaces issuer,clusterissuer,certificates,secrets --kubeconfig=/home/$(users)/.kube/config > cert-manager-backup.yaml
kubectl get cm k8s-mgmt.cloudaccount --kubeconfig=/home/$(users)/.kube/config -n cisco -o jsonpath="{.data.data}" | base64 -d > k8s-mgmt.cloudaccount.json
echo -e "Backing up all the secrets to resources_backup.yaml"
kubectl get -o yaml --all-namespaces secrets --kubeconfig=/home/$(users)/.kube/config > resources_backup.yaml
for n in $(kubectl get -n cisco -o=name pvc,certificate,configmap,serviceaccount,secret,ingress,
```



```

service,deployment,statefulset,hpa,job,cronjob)
do
    mkdir -p $(dirname "$n")
    kubectl get -n cisco --kubeconfig=/home/${users}/.kube/config -o=yaml --export "$n" > "$n".yaml 2>
/dev/null
done
cd ..
echo -e "-----"
}

ETCD_ENCRYPTION_SECRET=$(head -c 32 /dev/urandom | base64)

function generate_encryption_conf_file {
    #!/bin/bash

OUTPUT_FILE=/etc/kubernetes/pki/etcd/encryption.conf

echo -e "\nCreating ETCD_ENCRYPTION_SECRET and ETCD_ENCRYPTION_KEY"

ETCD_ENCRYPTION_SECRET=$(head -c 32 /dev/urandom | base64)
ETCD_ENCRYPTION_KEY=$(echo ccp-key-$(echo "$ETCD_ENCRYPTION_SECRET" | base64 -d | sha256sum | cut -c1-10))

echo "Writing content to file $OUTPUT_FILE"

cat <<EOF > $OUTPUT_FILE
apiVersion: apiserver.config.k8s.io/v1
kind: EncryptionConfiguration
resources:
  - resources:
    - secrets
    providers:
    - aescbc:
      keys:
        - name: $ETCD_ENCRYPTION_KEY
          secret: $ETCD_ENCRYPTION_SECRET
  - identity: {}
EOF
}

check_if_running_sudo
check_if_vsphere_cloud
check_for_encryption_file
backup_resources
echo -e "\nCreating etcd Encryption configuration file for all the master nodes"
generate_encryption_conf_file
echo -e "\nSaved etcd Encryption configuration file at $OUTPUT_FILE"
cp $OUTPUT_FILE encryption.conf
echo -e "\nSaved etcd Encryption configuration file's copy at encryption.conf\n\n"

#-----
#-----

function check_if_running_sudo {
if [ "$EUID" -ne 0 ]
    then echo "This script needs to run as sudo. Please retry with:"
echo "sudo bash $0"
    exit
fi
}
function sleep_function {
    echo "Sleeping for $1 minutes..."
    while true;do echo -n .;sleep 1;done &
    sleep "$1" # or do something else here
    kill $!; trap 'kill $!' SIGTERM
    echo "Done"
    echo -e "\n\n-----"
}
function check_if_vsphere_cloud {

```

```

    if kubectl get cm k8s-mgmt.cloudaccount --kubeconfig=/home/$(users)/.kube/config -n cisco ; then
        echo "continue with the upgrade..."
    else
        echo "Workaround is not required for this cloud. Please continue with the upgrade from installer UI."
        exit
    fi
    kubectl get cm k8s-mgmt.cloudaccount --kubeconfig=/home/$(users)/.kube/config -n cisco -o jsonpath="{.
data.data}" | base64 -d > k8s-mgmt.cloudaccount.json
    sleep 60
    CLOUD_TYPE=$(cat k8s-mgmt.cloudaccount.json | jq -r '.cloudType')
    if [[ $CLOUD_TYPE == *"vsphere"* ]]; then
        echo "Vshere cloud, continue with the workaround."
    else
        echo "$CLOUD_TYPE cloud, workaround is not required for this cloud. Please continue with the upgrade
from installer UI"
        exit
    fi
}

function check_for_encryption_file {
    #Function to check if this file is executed only Once, if first time, create a lock
    FILE=/etc/kubernetes/pki/etcd/encryption.conf
    if [ -f "$FILE" ]; then
        echo -e "-----"
        echo "$FILE already exists.
Looks like workaround is already performed on this master node.
Please do not run the script multiple times on same master. Continue with other master nodes."
        echo -e "-----"
        exit
    fi
}
check_if_running_sudo
check_if_vsphere_cloud

OUTPUT_FILE=/etc/kubernetes/pki/etcd/encryption.conf

function copy_encryption_conf_file {
    #!/bin/bash

    OUTPUT_FILE=/etc/kubernetes/pki/etcd/encryption.conf
    echo "Copying script to $OUTPUT_FILE"
    cat <<EOF > $OUTPUT_FILE
$(sed -n -e '/^ETCD_ENCRYPTION_FILE_START$/ ,/^ETCD_ENCRYPTION_FILE_END$/ { /^ETCD_ENCRYPTION_FILE_START$/
/d; /^ETCD_ENCRYPTION_FILE_END$/d; p; }' "${0}")
EOF
cp $OUTPUT_FILE encryption.conf
echo -e "\nSaved etcd Encryption configuration file's copy at encription.conf\n\n"
}

# File content which updates kube-apiserver.yaml and waits for kubserver to start
function check_kube_api {
    API_PID=$(ps -ef | awk '/kube-ap[i]server/{print $2}')
    echo -e "-----"
    echo "check_kube_api to verify that kube-apiserver is restarted: $API_PID ..."
    while [ $API_PID = $1 ]; do
        echo "wait for kube api server to exist.."
        sleep_function 2m
        API_PID=$(ps -ef | awk '/kube-ap[i]server/{print $2}')
    done

    while ! ps -ef | grep kube-ap[i]server; do
        echo "wait for kube api server to restart.."
        sleep_function 2m
    done

    echo -e "\n=====
echo -e "Waiting. Please DO NOT INTERRUPT."
sleep 300
echo -e "Successfully updated kubernetes manifests"

```

```

echo -e "\n=====
}

function update_kube_apiserver_manifest {
    API_PID=$(ps -ef | awk '/kube-apiserver/{print $2}')

    echo -e "Adding ETCD encryption resources to kube-apiserver, it might take some time. Please DO NOT
    INTERRUPT."

    API_CONF="/etc/kubernetes/manifests/kube-apiserver.yaml"
    ENC_CONF="/etc/kubernetes/pki/etcd/encryption.conf"

    sudo python - <<EOF
import yaml
name = 'k8s-encryption'

with open("$API_CONF", "r") as f:
    data = yaml.safe_load(f)
    data['spec']['volumes'].append({'hostPath': {'path': "$ENC_CONF", 'type': ''}, 'name': name})
    data['spec']['containers'][0]['command'].append("--encryption-provider-config=$ENC_CONF")
    data['spec']['containers'][0]['volumeMounts'].append({'mountPath': "$ENC_CONF", 'name': name})

with open("$API_CONF", "w") as f:
    f.write(yaml.dump(data, default_flow_style=False, indent=2))
EOF

    sleep_function 2m
}

function patch_kubeconfig_secret {

    KUBECONFIG_SECRET_NAME=$(kubectl get secret -n ccp --kubeconfig=/home/$(users)/.kube/config | grep
    kubeconfig | awk '/-/{print $1}')
    echo "\nCreating patch file for kubeconfig secret $KUBECONFIG_SECRET_NAME"
    echo "{\"data\":{\"etcdEncryptionKey\": \"${ETCD_ENCRYPTION_SECRET}\"}}\" > patch_kubeconfig_secret.
    json
    kubectl -n ccp patch secret "$KUBECONFIG_SECRET_NAME" --kubeconfig=/home/$(users)/.kube/config --patch
    "$(cat patch_kubeconfig_secret.json)"
    echo -e "\nSuccessfully patched kubeconfig secret, sleeping for 2 mins, please DO NOT INTERRUPT.\n\n"
    sleep_function 2m
    kubectl get secrets --all-namespaces -o json --kubeconfig=/home/$(users)/.kube/config | kubectl
    replace -f - --kubeconfig=/home/$(users)/.kube/config
    echo -e "\nUpdating the secret, please DO NOT INTERRUPT."
    sleep_function 5m

    echo -e
"\n\n=====
    echo -e "=====
    echo -e "\nWorkaround completed."
    echo -e "\nRestart all 3 Master nodes one by one."
    echo -e "\nGive significant time for all services to become RUNNING before proceeding with next master
    node.(Recommended ~5 mins)"
    echo -e
"\n\n=====
    echo -e "=====
}

#-----
#-----

API_PID=$(ps -ef | awk '/kube-apiserver/{print $2}')
echo -e "\n\nPrinting encryption file, please copy it to other masters and execute\n"
echo ": ' > etcdEncrypt_master_2.sh
echo ": ' > etcdEncrypt_master_3.sh
echo "ETCD_ENCRYPTION_FILE_START" >> etcdEncrypt_master_2.sh
echo "ETCD_ENCRYPTION_FILE_START" >> etcdEncrypt_master_3.sh
cat $OUTPUT_FILE >> etcdEncrypt_master_2.sh

```

```

cat $OUTPUT_FILE >> etcdEncrypt_master_3.sh
echo "ETCD_ENCRYPTION_FILE_END" >> etcdEncrypt_master_2.sh
echo "ETCD_ENCRYPTION_FILE_END" >> etcdEncrypt_master_3.sh
echo " " >> etcdEncrypt_master_2.sh
echo " " >> etcdEncrypt_master_3.sh

sed -n '107,245p' etcdEncrypt_master_1.sh >> etcdEncrypt_master_2.sh
sed -n '107,245p' etcdEncrypt_master_1.sh >> etcdEncrypt_master_3.sh
echo "check_for_encryption_file" >> etcdEncrypt_master_2.sh
echo "check_for_encryption_file" >> etcdEncrypt_master_3.sh
echo "copy_encryption_conf_file" >> etcdEncrypt_master_2.sh
echo "copy_encryption_conf_file" >> etcdEncrypt_master_3.sh
echo "update_kube_apiserver_manifest" >> etcdEncrypt_master_2.sh
echo "update_kube_apiserver_manifest" >> etcdEncrypt_master_3.sh
echo -e "export -f check_kube_api \ntimeout 300s bash -c check_kube_api "\$API_PID" " >>
etcdEncrypt_master_2.sh
echo -e "export -f check_kube_api \ntimeout 300s bash -c check_kube_api "\$API_PID" " >>
etcdEncrypt_master_3.sh
echo -e "ETCD_ENCRYPTION_SECRET="\$ETCD_ENCRYPTION_SECRET">>etcdEncrypt_master_3.sh
echo "patch_kubeconfig_secret" >> etcdEncrypt_master_3.sh

update_kube_apiserver_manifest
export -f check_kube_api
timeout 300s bash -c check_kube_api "$API_PID"

echo -e "\nWorkaround completed on Master-1, please continue with Master-2 and Master-3"
echo -e "\n\n=====
echo -e "\nCopy etcdEncrypt_master_2.sh to second master and run sudo bash etcdEncrypt_master_2.sh"
echo -e "\nCopy etcdEncrypt_master_3.sh to third master and run sudo bash etcdEncrypt_master_3.sh"
echo -e "=====

```

The script is mandatory, not optional for the upgrade to work.

IMPORTANT: DO NOT run the workaround script TWICE on any of the masters, this can bring the cluster since script interacts with kube-apiserver of the cluster.

4. Use v5.3.1 installer to upgrade the kubernetes cluster & cert-manager
5. After upgrading the cluster, users can use the "Take me to Suiteadmin" link to go back to the Suite Admin UI. Please note that this address might change for DHCP based clusters.



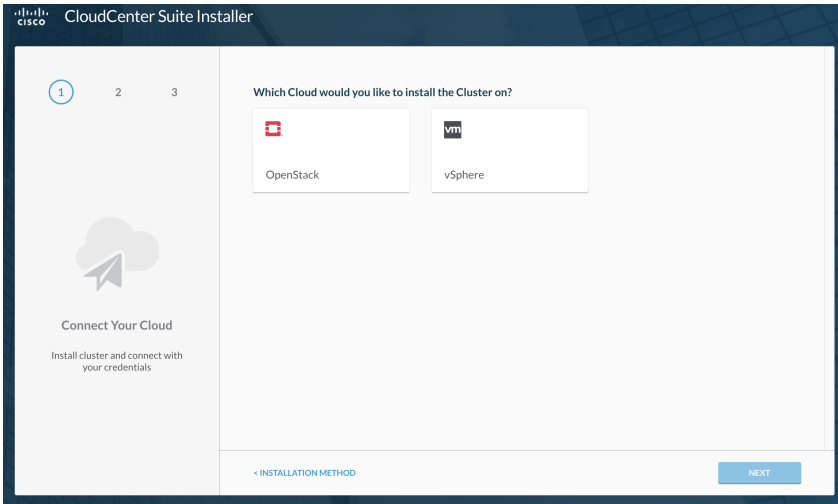
In the upgraded cluster, if you upgrade Workload-manager, Cost-optimizer, or Action Orchestrator and see this error on suite admin UI: "failed to wait for command /bin/helm, err: exit status 1, msg: Error: [unable to recognize "": no matches for kind "Certificate" in version "certmanager.k8s.io/v1alpha1", unable to recognize "": no matches for kind "Issuer" in version "certmanager.k8s.io/v1alpha1"]", please follow workaround steps mentioned in [Update Module](#).



Kubernetes Cluster Installation

CCS Installer drops support for Public clouds (AWS, GKE & Azure) for new cluster installation, however, the "existing cluster" mode of installation is still supported. Users can bring a GKE/Azure based Kubernetes cluster and install suite admin using v5.3.1 installer. (CPSGCORE-4305-Authenticate to see issue details)

CCS v5.3.1 does not have installer images for public clouds. Only OVA and QCOV2 appliances are published.

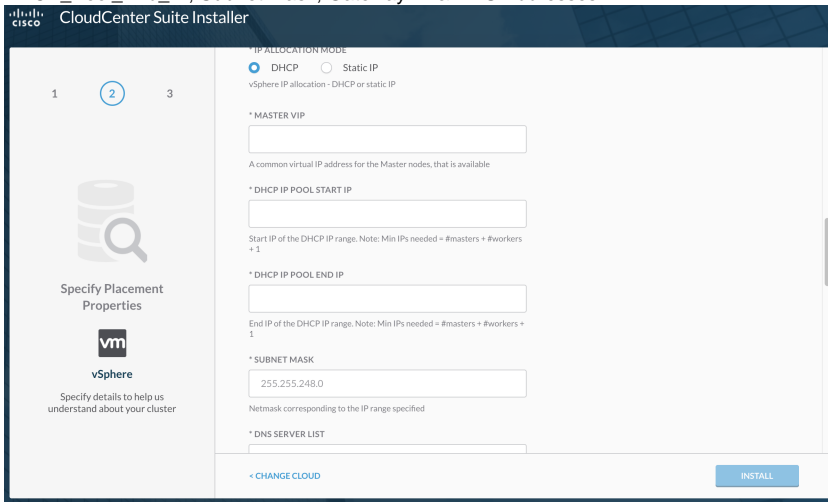


CCS Installer upgrades Supported Kubernetes version and cert-manager

- CCS v5.2.3 (Previous release) supported Kubernetes v1.16.3 and Cert-manager v0.10.1
- CCS v5.3.1 Upgrades Kubernetes to v1.18.12 and Cert-manager to v1.0.2

CCS Installer changes how DHCP based clusters are installed:

- CCS v5.2.3 required users to only provide only a master VIP to create a cluster, but
- CCS v5.3.1 now requires users to provide all the network information like static IP case to create a cluster & its resources. DHCP_Pool_Start_IP, DHCP_Pool_End_IP, Subnet mask, Gateway IP & DNS Addresses.



Helm Chart Upgrade from 5.2.3 to 5.2.4

If you are upgrading from SA 5.2.3 or older to 5.2.4, perform the following procedure after upgrade:

1. Run an SQL command in the suite-postgresql database.
2. Log in to common-framework-suite-postgresql-0

Renew your license to continue

Your evaluation license has expired. Contact your administrator to renew your Reporting for Confluence license.

3.Run the following bash command.

```
PGDATABASE=suite-samlpsql-c"UPDATEpublic.saml_infra_configSETcert_source='cert_manager' WHERE id=1;"
```

4.Optionally regenerate the SSO certificate.

```
kubectldeldelete secret suite-saml-sso-tls-n cisco
```

5. Restart the suite-samlpod.

```
kubectldeldelete pod suite-saml-pod -n cisco
```

6. If yourcustomeruses SSO, reconfigure SSO.

Architecture

No updates

Public Clouds

The following public clouds with existing kubernetes clusters are supported:

- **Azure AKS:**
 - Supported Version: 1.20.9
- **Google GKE:**
 - Supported Version: 1.20.11
- See [Existing Cluster Installation](#) for additional details
- Kubernetes version upgrade on existing clusters from older version 1.13,1.16.x is not supported for public clouds.

Administration

No updates

Module Management

No updates

Smart Software Licensing

No updates

Suite Admin Dashboard

No updates

User Tenant Management

No updates

Cluster Management

No updates

Security Management

No updates

Suite UI

No updates.

Deprecated

- CCS Installer drops support for Public clouds(AWS, GKE & Azure) for new cluster installation

API

New API Calls

Updated API Calls

No updates.

Documentation

The following documentation changes were implemented in CloudCenter Suite5.3.1:

- [Support for Two Cloud Types](#)
- [Support for New Kubernetes Version 1.18.x](#)
- [Restrictions on Containers](#)
- [New IP Allocation Mode](#)
- [Helm Chart Configuration after k8s Upgrade](#)
- [Update Backup](#)
- [SSO Configuration](#)
- [VMware Folder Name Change](#)
- [Cluster Upgrade](#)

Known Issues

CloudCenter Suite 5.3.1 has the following known issues:

- When the Suite Admin chart is upgraded from version 5.2.4 to 5.3.1 in a few elasticsearch pods, note the following:
 - The busybox image is not upgraded from version 1.29.3 to 1.33.1.
 - The elasticsearch image is not upgraded from elasticsearch version 6.8.12 to elasticsearch version 6.8.13-cisco2.
 - Two elasticsearch client pods are running after performing a Suite Admin upgrade.

Resolved Issues

CloudCenter Suite 5.3.1 has the following resolved issues:

- **CSCwa47349**: Vulnerability in Apache Log4j Library

Suite Architecture

Suite Architecture

- [Overview](#)
- [The Suite Architecture](#)
- [Port Requirements](#)
- [The Suite Admin](#)
- [The Modules](#)

Overview

The %ccsis Cisco's hybrid cloud deployment platform. This platform takes a unique approach to install, configure, and maintain hybrid cloud environments that are often encountered by Information Technology (IT) departments to adopt business agility and improve time-to-market solutions within an enterprise. As a cloud-based organization, your enterprise can choose from multiple cloud (*multicloud*) providers depending on your location, policies, permissions, security requirements, and governance regulations for both traditional and modern IT requirements.

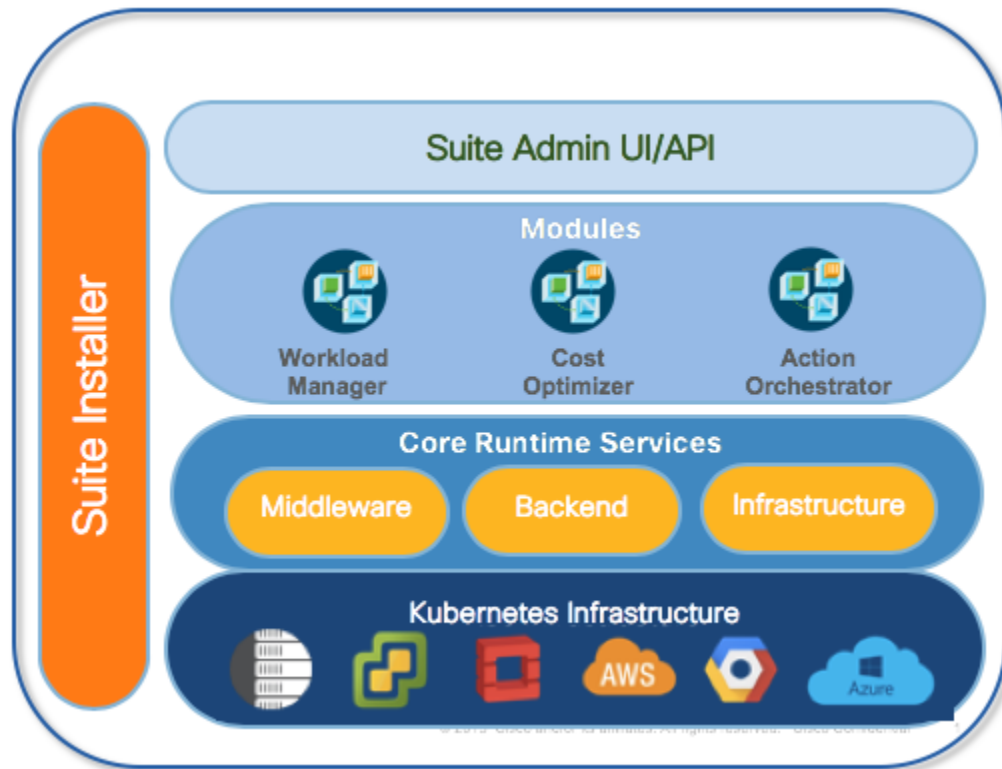
The %ccs provides a solution that is cloud agnostic, works with diverse workloads, provides cross-domain orchestration, supports cost-optimization, and integrates easily in an agile world.

The Suite Architecture

The %ccsis made up of the following components:

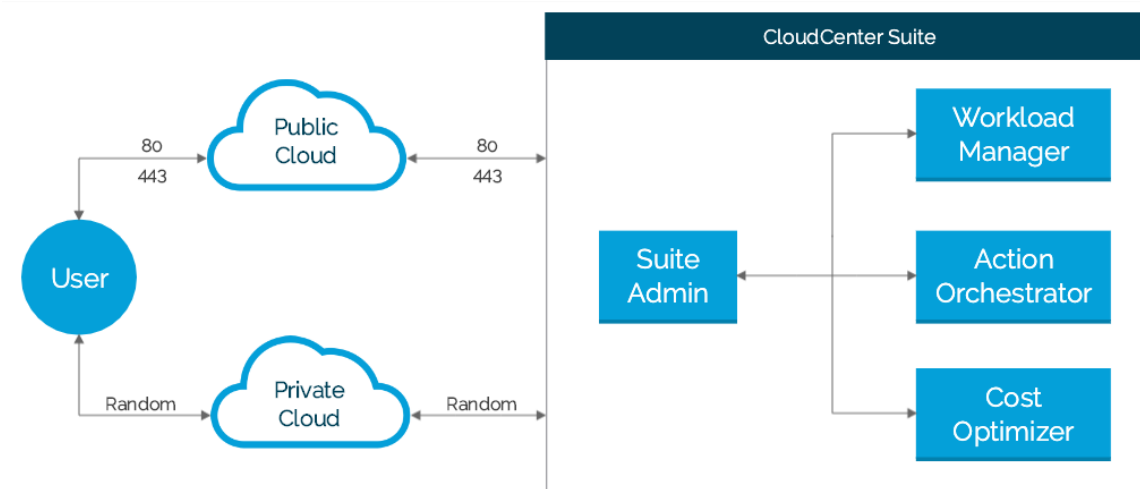
- **Suite Installer** Installs the Suite Admin. See [Installer Overview](#) for additional details.
- **Suite Admin** Installs and launches a suite of modules. See *The Suite Admin* section below for additional details.
- **Modules** The workload manager, the Cost Optimizer, and the Action Orchestrator. See *The Modules* section below for additional details.
- **Core Runtime Platform and Kubernetes Infrastructure** A Kubernetes-based platform that allows you to launch each module on a new or existing Kubernetes cluster.

The following image displays the Suite Admin architecture.



Port Requirements

The following image identifies the ports that must be open for the CloudCenter Suite to function as designed.



The Suite Admin

When you download and install the Suite Installer, the [Suite Admin](#) is **already installed**! You have the option to use the Suite Admin UI to perform the following tasks:

- Install additional, available modules based on the list available in the Dashboard.
- Upgrade the Suite Admin or other installed modules when a new version becomes available.

The Modules

The Suite Admin facilitates the installation of the following modules:

- **workload manager:**
 - This module allows IT organizations to provide management for clouds (public/private/container), applications, VMs/pods, governance policies with centralized visibility and permission control for enterprise environments.
 - See [Workload Manager](#) for additional details.
- **Action Orchestrator:**
 - This module allows IT organizations to use cross-domain orchestration to automate a process that has multiple, complex steps with a specific order and implemented across different technical domains.
 - See [Action Orchestrator](#) for additional details.
- **Cost Optimizer:**
 - This module allows IT organizations to use cost optimization in a pay-per-use environment to avoid consumption that does not add value.
 - See [Cost Optimizer](#) for additional details.

Each module in the CloudCenter Suite is independent and allows access to additional gateways or endpoints so you can add on module-specific components on supported clouds.

Self-Hosted Installation

Self-Hosted Installation

- [Installer Overview](#)
- [Installer Virtual Appliances](#)
- [Prepare Infrastructure](#)
- [New Cluster Installation](#)
- [Existing Cluster Installation](#)
- [Upgrade Kubernetes Cluster](#)
- [Air Gap Installation](#)
- [Upgrade Offline Repository](#)
- [Backup and Restore](#)
- [Troubleshooting](#)

Installer Overview

Installer Overview

- [Overview](#)
- [Supported Clouds](#)
- [Installer Appliance Download Location](#)

Overview

The CloudCenter Suite provides a new way to install, configure, and maintain multiple modules that jointly make up the suite. The CloudCenter Suite has a common installer to install, upgrade, and integrate all modules included in the suite.

You can install the CloudCenter Suite by using installer appliance images provided by Cisco. As part of the installation process, the CloudCenter Suite installs the Suite Admin. Once authenticated, each user can access the CloudCenter Suite using valid credentials created by the Suite Administrator.

Supported Clouds

Cisco supports the following private clouds for the CloudCenter Suite:

- [VMware vSphere 6.5](#)
- [OpenStack Queens](#)



All supported clouds are visible and enabled for private cloud installers.

This includes both the functionality and the CloudCenter Suite UI.

Installer Appliance Download Location

Major releases include installer appliances for the following components and cloud providers.

You can download these files from software.cisco.com.

The [Virtual Appliance Overview](#) section provides more details on these files.

Installer Virtual Appliances

Installer Virtual Appliances

- [Virtual Appliance Overview](#)
- [OpenStack Appliance Setup](#)
- [VMware vSphere Appliance Setup](#)

Virtual Appliance Overview

Virtual Appliance Overview

- [Virtual Appliance Overview](#)
- [General Virtual Appliance Approach](#)
- [Cloud-Specific Setup](#)



Virtual Appliance Overview

The only way to install the CloudCenter Suite is to use the virtual appliance Installer method. Cisco builds these appliances on CentOS 7.x base images.

General Virtual Appliance Approach


To prepare infrastructure for the appliance approach, follow this process.

1. Review and ensure that you have met the requirements to [Prepare Infrastructure](#) before installing the CloudCenter Suite.
2. Review the list of [Supported Suite Installer](#) to verify the supported Virtual Appliances.
3. Navigate to software.cisco.com to download virtual appliances for each supported cloud.
4. Follow directions as specified in the table below to obtain and import each image.

Cloud	Image Type	Description
OpenStack	Downloaded Virtual Appliance (QCOW2)	Import the QCOW2 image file using the OpenStack client. Refer to the OpenStack Documentation for additional context.
VMware vSphere	Downloaded Virtual Appliance (OVA)	<p>Follow this procedure:</p> <ol style="list-style-type: none">a. Download the OVA image.b. Import the OVA to your vSphere environment by using the vSphere client<ol style="list-style-type: none">i. When you import the OVA as a VM, ensure that it is powered off on vSphere.ii. If your environment requires a static IP, use a VMware Customization Spec to manually configure the static IP for the installer VM.c. A default password is required to ensure access to the VM using the console (in case the SSH has issues). <div data-bbox="522 1144 1484 1341" style="border: 1px solid green; padding: 10px;"><p> If you provide a default password or public-key, be aware of the following requirements:</p><ul style="list-style-type: none">• The login user is the cloud-user.• If you configure a default password or public key in the VM, you must also configure the default instance ID and hostname fields as they are dependent and required fields.• Use this password to access the VM via vSphere console.• You cannot use this password to SSH into the launched VMs.</div> <ol style="list-style-type: none">d. Select the required Network for the interface to be connected.e. Convert the VM to a template. <div data-bbox="522 1432 1484 1541" style="border: 1px solid orange; padding: 10px;"><p> You <i>must</i> convert the VM to template and then create a VM from this template, so that the template can be used when installing a VMware data center. If you do not provide the template name when installing a VMware data center, your installation will fail.</p></div> <ol style="list-style-type: none">f. Select the template created in the previous step and <i>clone to Virtual Machine</i>, to launch the installer VM. This template will also be used as the value for the <i>vSphere Template Name</i> cloud setting, in the installer UI.g. After the VM is created from the template, power it on. To access the UI, go to the newly created VM IP using HTTPS protocol in a supported browser (see Browser Compatibility).

5. Launch the installer instance using the image.

Cloud-Specific Setup

 The per-cloud setup procedures are only listed below to serve as sample setup scenarios.

- [OpenStack Appliance Setup](#)

- [VMware vSphere Appliance Setup](#)

OpenStack Appliance Setup


OpenStack Appliance Setup

To setup infrastructure for OpenStack clouds, follow this process.



The exact VM size really depends on the instance type configuration in your environment! See [Prepare Infrastructure](#) > *Resource Requirements for CloudCenter Suite Modules* for additional details.

1. Download the CloudCenter Suite QCOW2 file to your local machine.
2. Login into your OpenStack datacenter to perform this task.
 - a. Click **Images**.
 - b. Click the **Create Image** button.
 - c. Enter a valid name.
 - d. Click the **File Browse** button.
 - e. Select the QCOW2 file stored in your local machine as displayed in the following screenshot.



3. In the **Format** dropdown, select QCOW2.
4. To share this image with other users, select **Public** in the **Image Sharing Visibility** field.

5. Click **Next** and then click the **Create Image** button as displayed in the following screenshot.

Create Image

Specify an image to upload to the Image Service.

Image Name
installer985

Image Description

Image Source

Source Type
File

File
Browse... suite-v0.0.978.qcow2

Format
QCOW2 - QEMU Emulator

Image Requirements

Kernel
Choose an image

Ramdisk
Choose an image

Architecture

Minimum Disk (GB)
0

Minimum RAM (MB)
0

Image Sharing

Visibility
Public Private

Protected
Yes No

Cancel < Back Next > Create Image



The image import will take some time depending on the network speed. During this time, do not close the browser/application/tab.

6. Create the instance for each component using the imported images:

- Follow the standard OpenStack procedure to create the instance from an image.
- Create the security group(s) with Port 80 and 443 (optionally 22 if you need SSH access) open for Ingress and Outbound communication.
- You may need to assign floating IP to your VM after you create the VM is created.

7. Select a new or existing key pair to log into each instance if multiple key pairs are available, you must *select oneto* be used for the CloudCenter instance as displayed in the following screenshot.



If you do not select a key pair, you will not be able to log into the component VM!

Import Key Pair

Key Pairs are how you login to your instance after it is launched. Choose a key pair name you will recognize and paste your SSH public key into the space provided.

Key Pair Name
cliqr_user-key_1

Public Key

```
ssh-rsa
AAAAAB3NzaC1yc2EAAAADAQABAAQAC4x93DDQBAwT5D54aQrKdUHQNaakudda
...
z9gucsWNgAtNoD12ua0YpMeBX020QWiLAZ6g7/GkrijSF0iH2BfeYIAsc8aAOP7DngcYI
HJIDYDjFqrCILgvZqQ76J
```

Cancel Import Key Pair

You have now setup the installer for an OpenStack cloud.

VMware vSphere Appliance Setup

VMware vSphere Appliance Setup

To setup infrastructure using CloudCenter appliances for VMware vSphere clouds, follow this process.

1. **Configure Network Time Protocol (NTP) on the VMware ESXi hosts this is important as the CloudCenter Suite installation can fail, if NTP is not configured or if it is wrongly configured.**
See https://kb.vmware.com/s/article/57147?lang=en_US for additional details.

i Note the value that you enter in this field for later use. You will need to enter the same values for the **NTP Servers** or **NTP Pools** fields in the Placement Properties page (see [VMware vSphere Installation](#) > Advanced Installation Process > Step 6).

Identical NTP values are required to ensure that the NTP communication between the installer and CloudCenter Suite master/worker VMs are in sync so the certificates generated by the installer for CloudCenter Suite are also in sync.

2. Download the OVA image file from software.cisco.com to your local machine.

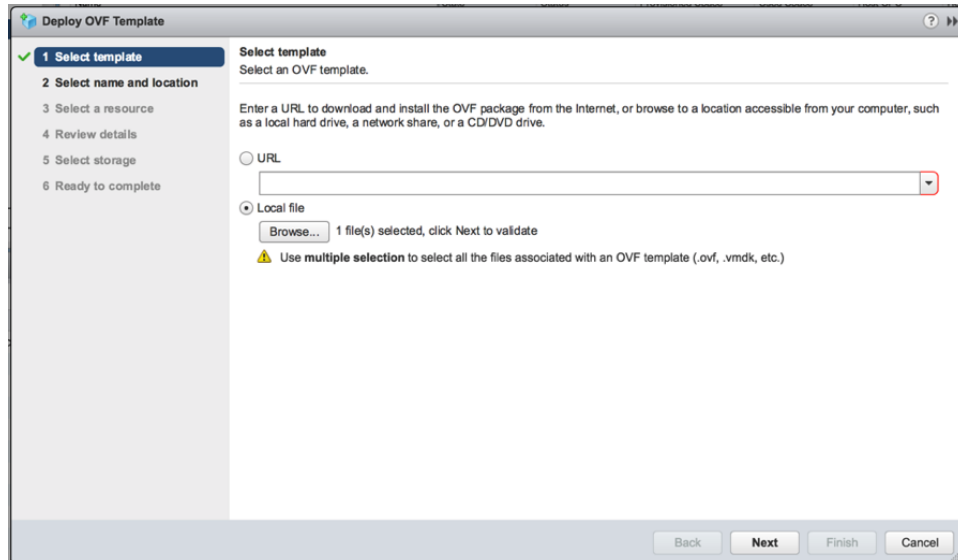
✓ The installer appliance has/requires a minimum resource requirement of 4 vCPUs and 75 GB storage (root disk).

3. Log into the VMware Datacenter console and click on the **VMs and Templates** section.
4. Deploy an OVA template (right-click and select Deploy OVA Template option).

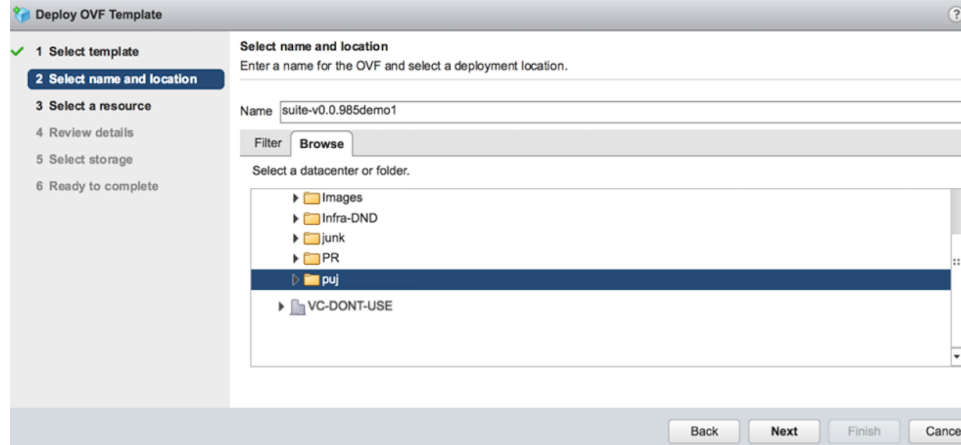
- a. If DHCP **is** installed, follow these steps.

Follow these steps **ONLY** if DHCP **is** installed.

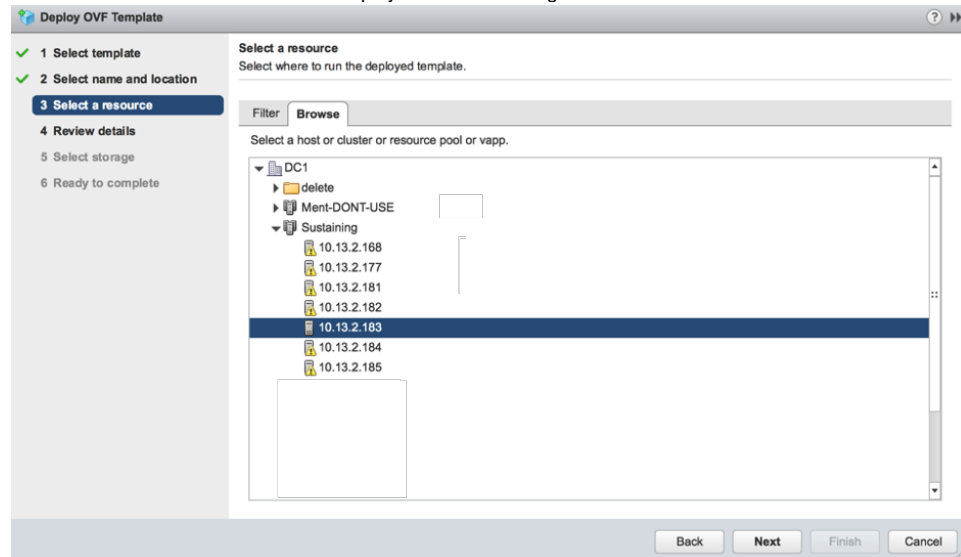
- i. Click the **Local file** option, click **Browse** to provide the location for the downloaded OVA file, ensure the file is selected, and then click **Next** as displayed in the following screenshot.



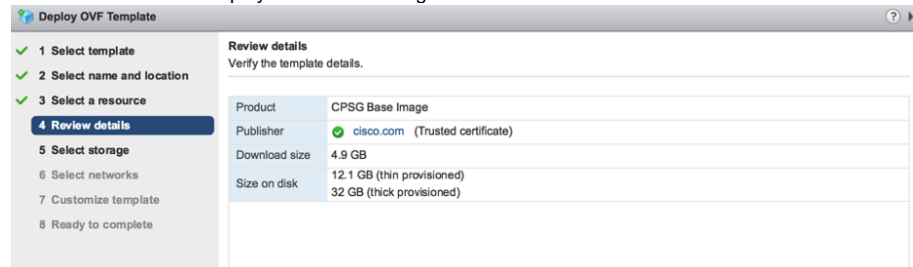
- ii. Provide a suitable name and select the target folder where you need to create the Template as displayed in the following screenshot.



- iii. Select a suitable host and cluster as displayed in the following screenshot.



- iv. Review the details as displayed in the following screenshot.



- v. Select the storage location as displayed in the following screenshots.



Use **Thin Provision** as the storage format so it has the flexibility to optimize the storage location. The following screenshots displays views from two different datacenters to provide a point of context.

Select storage

Select the datastore in which to store the configuration and disk files

- Same format as source
- Thick Provision Lazy Zeroed
- Thick Provision Eager Zeroed
- Thin Provision**

Configure per disk

Select virtual disk format:

VM Storage Policy: Keep existing VM storage policies

Name	Capacity	Provisioned	Free	T
Storage Compatibility: Compatible				
datastore26-1	7.26 TB	370.6 GB	7.09 TB	

Compatibility

✓ Compatibility checks succeeded.

CANCEL BACK NEXT

Select storage

Select the datastore in which to store the configuration and disk files

Select virtual disk format:

- Thin Provision
- Same format as source
- Thick Provision Lazy Zeroed
- Thick Provision Eager Zeroed
- Thin Provision

VM storage policy:

The following datastores are available. Select the destination datastore for the virtual machine configuration files.

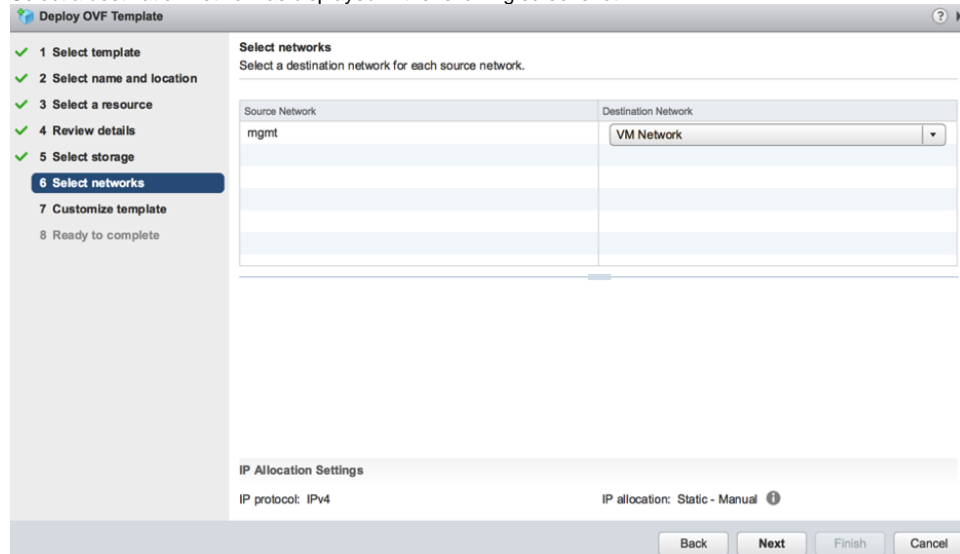
Name	Capacity	Provisioned	Free	Type	Cluster
hx-scale	128 TB	14.24 TB	121.89 TB	NFS v3	
SpringpathDS-WZP223202DA	216 GB	9.89 GB	206.11 GB	VMFS 5	

Advanced >>


Compatibility

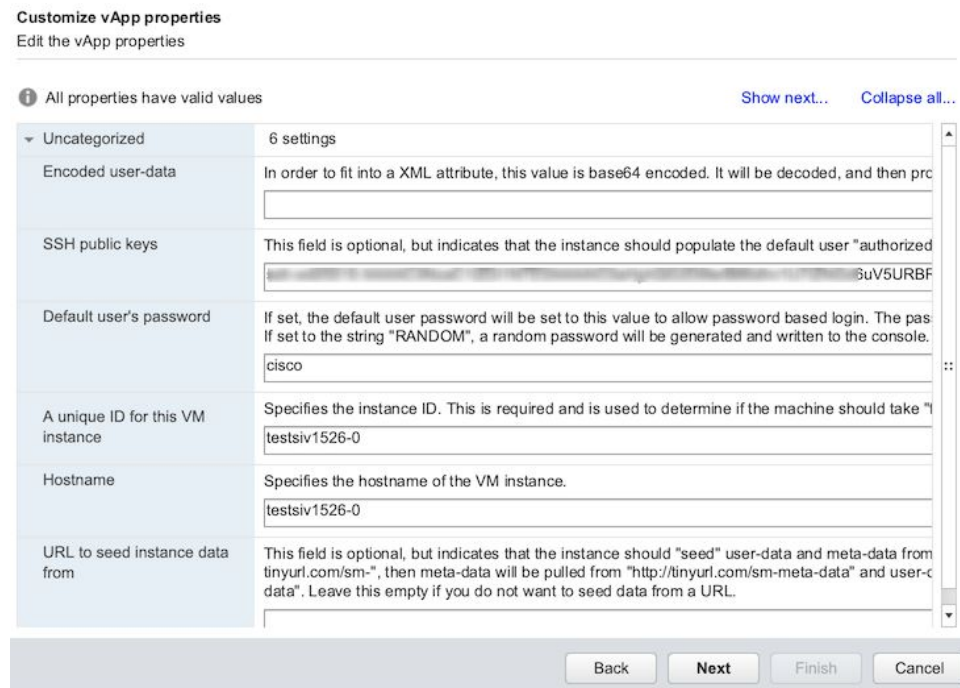
Back Next Finish Cancel

- vi. Select a destination network as displayed in the following screenshot.




- vii. Enter the information identified below in the Customize vApp Properties page displayed in the following screenshot.


 Do not customize your setup credentials at this point or any other point during the installation. You can do so after you complete the installation process.



1. The public SSH key.
2. The default user's password to SSH from the vSphere console.
3. The unique ID and hostname ensure that these credentials are unique to avoid duplication issues.

 Use lowercase characters when providing the installer hostname in the Customize vApp Properties page.

- viii. Customize the template as required for your environment and review the completed information as displayed in the following screenshot.



- ix. Click **Finish** to start deploying the VM from the template inside the target folder.

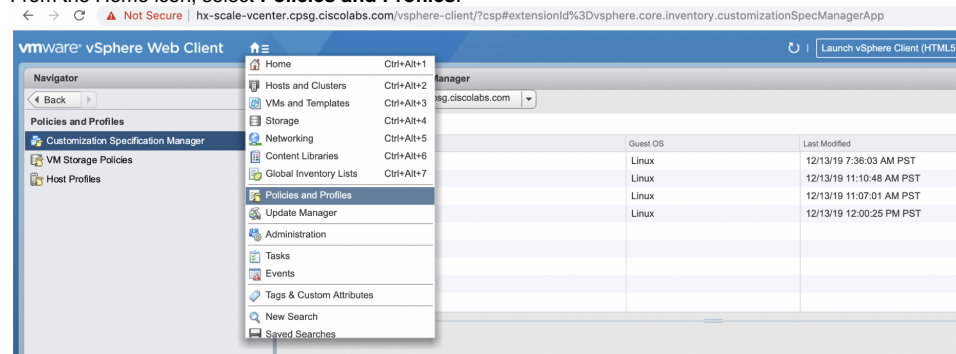
- b. If DHCP **is not** installed, follow these steps

Follow these steps **ONLY** if DHCP **is not** installed use your static IP as the VMware customization specification is needed to attach the IP to the installer VM.

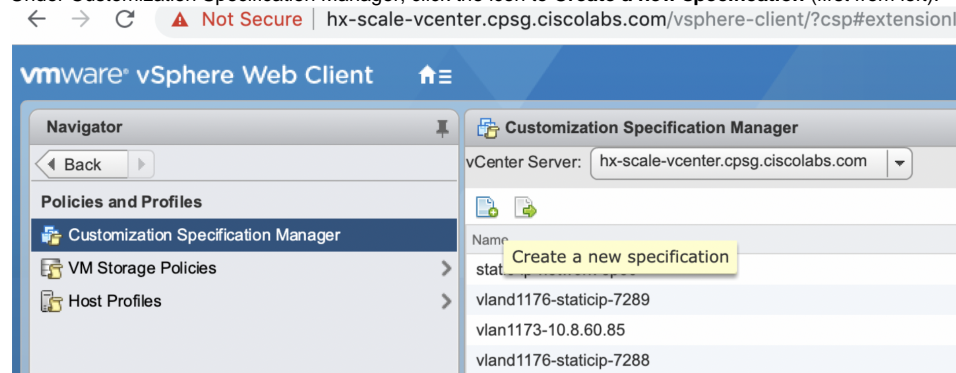
The details attached in the Customization *Specification* (term specific to vSphere), like the IP, DNS, Gateway, and so forth are assigned to the VM, when it is powered on.

IPs cannot be attached to the VM when it is Powered ON automatically and you must follow the instructions provided below to create an installation VM using the *Customization Specification* (specific to vSphere) which is used to create a template or custom profile with IP details, when attached to the VM.

- i. Login to vSphere.
- ii. From the Home icon, select **Policies and Profiles**.



- iii. Under Customization Specification Manager, click the icon to **Create a new specification** (first from left).



- iv. For the *Target VM OS*, select **Linux**.
- v. Set the *Computer Name* to any suitable name.

vi. Enter cpsg.ciscolabs.com as the *Domain name*.


The screenshot shows the 'New VM Guest Customization Spec' wizard at the 'Set Computer Name' step. The left sidebar shows steps 1 through 6, with 'Set Computer Name' selected. The main area is titled 'Computer Name' and contains the following options:

- Enter a name:
Input field: `compname-spec-test1571-1`
The name cannot exceed 63 characters.
 Append a numeric value to ensure uniqueness
The name will be truncated if combined with the numeric value, it exceed 63 characters.
- Use the virtual machine name
If the name exceeds 63 characters, it will be truncated.
- Enter a name in the Clone/Deploy wizard
- Generate a name using the custom application configured with the vCenter Server
Argument:

Domain name:

Buttons: Back, Next, Finish, Cancel

vii. To configure the network, select the button to **Manually select custom settings** for to ensure Static IP allocation so that you can manually enter the Static IP details.

 Select the option to **use standard network**. *if you are using a DHCP setup.*

The screenshot shows the 'New VM Guest Customization Spec' wizard at the 'Configure Network' step. The left sidebar shows steps 1 through 6, with 'Configure Network' selected. The main area is titled 'Configure Network' and contains the following options:

- Use standard network settings for the guest operating system, including enabling DHCP on all network interfaces
- Manually select custom settings


Below the options is a table with columns: Description, IPv4 Address, and IPv6 Address.

Description	IPv4 Address	IPv6 Address
NIC1	Use DHCP	Not used

Buttons: Back, Next, Finish, Cancel

viii. Enter other details in subsequent screens to complete the wizard requirements.

ix. Wait for the installer VM to start when it does, the Static IP assigned by the custom specification will be assigned to the VM.

 Currently, an existing VMware issue does not save the check box setting. To workaround this issue, click the **Edit** settings on the VM, and check it again, and save your changes to assign the static IP.

x. Click the **Edit** Button.

NIC1 - Edit Network

IPv4
Specify IPv4 settings for the virtual network adapter.

IPv6

Use DHCP to obtain an IP address automatically.

Prompt the user for an address when the specification is used

Use an application configured on the vCenter Server to generate an IP address

Argument:

Use the following IP settings:

IP Address:

Subnet Mask:

Default Gateway:

Alternate Gateway:

OK **Cancel**

xi. Click **OK** and then click the **Enter DNS and Domain Settings**.

xii. In the DNS search path enter **cpsg.ciscolabs.com** and click **OK**.

New VM Guest Customization Spec

Enter DNS and Domain Settings
Enter the DNS and domain information for this new virtual machine.

Primary DNS:

Secondary DNS:

Tertiary DNS:

DNS Search Path

Add **Delete** **Move Up** **Move Down**

Back **Next** **Finish** **Cancel**

xiii. Click **Next** and then **Finish**.

VMware vSphere Web Client

Customization Specification Manager

vCenter Server: hx-scale-vcenter.cpsg.ciscolabs.com

Name	Guest OS	Last Modified
static-ip-network-spec	Linux	12/13/19 7:36:03 AM PST
vland1176-staticip-7289	Linux	12/13/19 11:10:48 AM PST
vlan1173-10.8.60.85	Linux	12/13/19 11:07:01 AM PST
vland1176-staticip-7288	Linux	12/13/19 12:00:25 PM PST
Custom-Spec-test1571-1	Linux	12/18/19 10:35:48 AM PST

- xiv. **Create a New Installer VM using this customization spec.** Start creating the VM installer from the installer template, in the wizard section Select the **Clone** option, make sure to check the **Customize the Operating System** box so that you can select the custom specification in the next screen.

1 Edit settings

- ✓ 1a Select a name and folder
- ✓ 1b Select a compute resource
- ✓ 1c Select storage
- ✓ 1d Select clone options
- 1e Customize guest OS
- 1f Customize hardware
- 1g Customize vApp properties
- 2 Ready to complete

Select clone options
Select further clone options

- Customize the operating system
- Customize this virtual machine's hardware (Experimental)
- Power on virtual machine after creation

Back Next Finish Cancel

1 Edit settings

- ✓ 1a Select a name and folder
- ✓ 1b Select a compute resource
- ✓ 1c Select storage
- ✓ 1d Select clone options
- 1e Customize guest OS
- 1f Customize hardware
- 1g Customize vApp properties
- 2 Ready to complete

Customize guest OS
Customize the guest OS to prevent conflicts when you deploy the virtual machine

Operating System: CentOS 4/5 or later (64-bit)

Name	Guest OS	Last Modified
static-ip-network-spec	Linux	12/13/19 7:36:03 AM PST
vland1176-staticip-7289	Linux	12/13/19 11:10:48 AM PST
vlan1173-10.8.60.85	Linux	12/13/19 11:07:01 AM PST
vland1176-staticip-7288	Linux	12/13/19 12:00:25 PM PST
Custom-Spec-test1571-1	Linux	12/18/19 10:35:48 AM PST

Back Next Finish Cancel

- xv. **Select the specification that you need and click Next.**

1 Edit settings

- ✓ 1a Select a name and folder
- ✓ 1b Select a compute resource
- ✓ 1c Select storage
- ✓ 1d Select clone options
- ✓ 1e Customize guest OS
- 1f Customize hardware
- 1g Customize vApp properties
- 2 Ready to complete

Customize hardware
Configure the virtual machine hardware

Virtual Hardware VM Options SDRS Rules

- CPU: 4
- Memory: 8192 MB
- Hard disk 1: 75 GB
- SCSI controller 0: VMware Paravirtual
- Network adapter 1: vlan1176 (hx-scale-vm-net) Connect...
- CD/DVD drive 1: Client Device Connect...
- Video card: Specify custom settings
- VMCI device
- Other Devices

New device: ----- Select ----- Add

Compatibility: ESXi 5.5 and later (VM version 10)


Back Next Finish Cancel

- xvi. Enter other details in subsequent screens, to complete the wizard. Wait for the installer VM to start, the Static IP assigned by the custom specification will be assigned to the VM.
- xvii. Wait for the installer VM to start when it does, the Static IP assigned by the custom specification will be assigned to the VM.

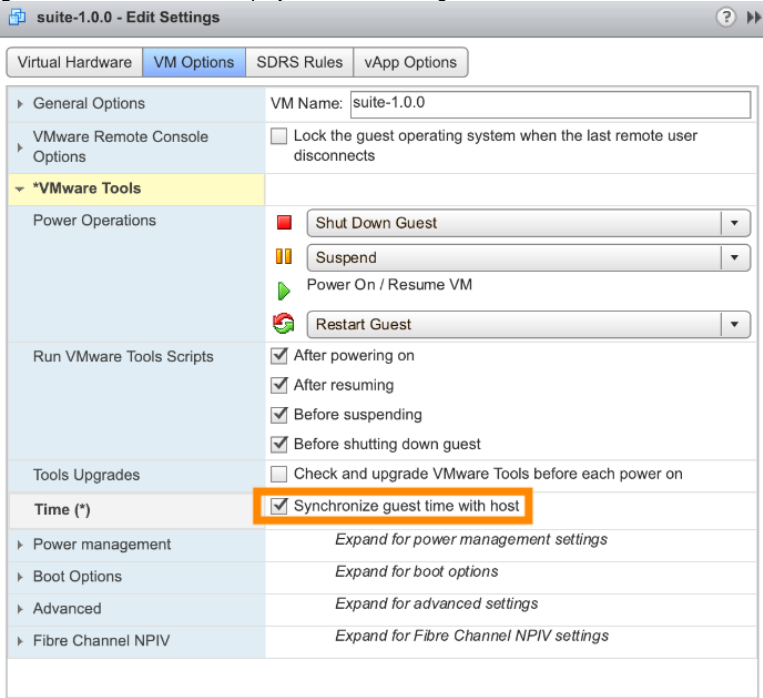


Currently, an existing VMware issue does not save the check box setting. To workaround this issue, click the **Edit** settings on the VM, and check it again, and save your changes to assign the static IP.


- Wait for some time so the VM is cloned and created, then refresh the VM page to view the powered off VM. The OVA is imported as a VM (powered off) on vSphere.

 When you import the OVA as a VM, ensure that it is powered off on vSphere.


- Right-click to edit the VM Settings for the powered off VM. Click the *VM Options* tab. Under *VMware Tools*, select the checkbox to **Synchronize guest time with host** as displayed in the following screenshot.




- Clone the VM to a template using the **Convert to template...** option (a sample of this template is displayed in the following screenshot).



- Once the VM is converted to template, it should appear as identified by the orange arrow in the following screenshot.

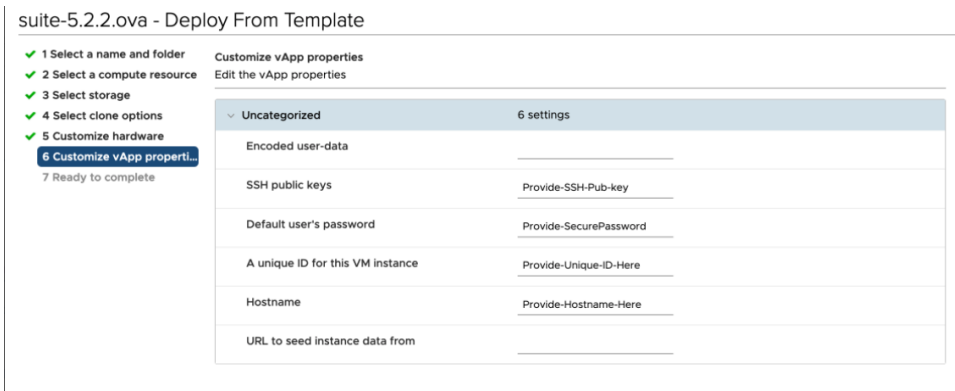


- Right click this template name and select the **New VM from This Template** option as displayed in the following screenshot this template will also be used as the value for the *vSphere Template Name* cloud setting, when you provide the details to install the Suite Admin.



- After the VM is created from the template, power it on.
- Edit the **1eCustomize vApp properties** to ensure that the VM has unique values for **A unique ID for this VM instance, Hostname, Default user's password, and SSH public keys** for this VM instance.

! For the password and/or the public key to take effect when deploying the VMware OVA for the CloudCenter Suite installer, you **must** change the **default-instance-id** to something else than *default-instance-id* or *the hostname*!



- Use this IP address to access the CloudCenter Suite UI (displayed in the following screenshot), go to the newly created VM's IP using HTTPS protocol in a supported browser (see [Browser Compatibility](#)).

You have now setup the installer for a VMware cloud.

Prepare Infrastructure

Prepare Infrastructure

- [General Compatibility](#)
- [Resource Requirements for CloudCenter Suite Modules](#)
- [Number of VMs](#)
- [IP Pool Requirements](#)
- [NTP Requirements](#)
- [The Suite Installer Dashboard](#)
- [Without Internet Access](#)

General Compatibility

See [Browser Compatibility](#) and the Suite Admin [Compatibility Matrix](#) for additional details. %ccs supports Kubernetes 1.16.3 for new installations.



For existing installations:

- For public clouds, Kubernetes support parallels the popular version supported by the major public cloud providers.
- For private clouds, the previous Kubernetes version (1.14) is also supported.

The CloudCenter Suite requires Tiller v2.16.3 to be installed. Refer to the [Helm documentation](#) for additional details.



Installers are already incorporated in the CloudCenter Suite SaaS offer, see [SaaS Access](#) for additional details.

Resource Requirements for CloudCenter Suite Modules

The following table lists the minimum resource requirements assuming that you install all available modules.

Module ^{1,2}	Public Cloud ⁵			Private Cloud ³		
	vCPU	Memory (GB)	Storage (GB)	vCPU	Memory (GB)	Storage (GB)
Suite Admin	16	37	300	16	37	300
Workload Manager ⁴ and Cost Optimizer	15	68	230 ⁶	15	68	230 ⁶
Action Orchestrator ⁷	20	30	750	20	30	750
Kubernetes Cluster (3 primary servers)	na	na	na	9	24	120
Total	51	135	1280	60	159	1400

¹ Update only one module at a time. If you simultaneously update more than one module, your update process may fail due to limited resource availability.

² Before updating any module, verify that you have un-allocated CPU/Memory in your cluster to ensure that your environment has free CPU/Memory a module-update scenario requires additional resources for the old pod to continue running until the new pod initializes and takes over. This additional resource requirement is temporary and only required while a module update is in Progress. After the module is updated, the additional resources are no longer needed.

³ On private clouds (vSphere and OpenStack), each of the 3 primary server instances require 3 vCPU and 8 GB memory and 40 GB storage (root disk), hence the difference in the additional requirement of 9 vCPU, 24 GB memory, and 120 GB storage (root disk). See the Number of VMs section below for additional details. Similarly, each worker instances require 3 vCPU and 8 GB memory and 40 GB storage (root disk) however, the number of workers changes dynamically at install time. Installer VMs require a minimum of 4 vCPUs and 8 GB RAM.

⁴ Workload Manager numbers include considerations for 4 Cloud Regions in the same instance. To support additional cloud regions, you must scale your cluster by adding Kubernetes worker nodes. You will need 1 CPU and 3 GB memory for each additional region. For regions without Cloud Remote, you will need 1.5 GB memory and 0.5 CPU when using Workload Manager 5.2.

⁵ Public clouds do not support auto-scaling the number of nodes might differ if scaled on an auto-scaling enabled node group.

⁶ The storage is 230 GB just to enable [StatefulSet](#) migration. In reality, only 115 GB is being used for operation of services.

⁷ Effective [Action Orchestrator 5.2.0](#). The Action Orchestrator also requires 3 worker nodes to proceed with the installation.

Number of VMs

A CloudCenter Suiteinstallation launches a highly available Kubernetes cluster which consists of primary server(s) and worker(s) instances.



The number of worker nodes (for both private and public cloud) vary based on the instance type selected during the installation process.

For private clouds, a redundant cluster requires a minimum of 2 out of 3 primary server nodes to be running at any point, so the cluster can function as designed.



If you plan to scale up at a later date, be aware that the worker instance type selected at installation time will also be used for the scaled nodes.

The CloudCenter Suite requires that the underlying disks for Kubernetes disk attachments be redundant and available. Most public clouds already provide built-in redundancy for their block disks (AWS EBS, GCP Persistent Disks, and so forth). Be sure to verify that the Datastores/Datastore Clusters are also on redundant, non-local storage (NFS, NetApp) before you begin the installation process.

IP Pool Requirements

You must select IP address to ensure that each IP endpoints is available, accessible, and not used by any other resource.

When configuring or modifying you pool of IP addresses, be aware of the following requirements:

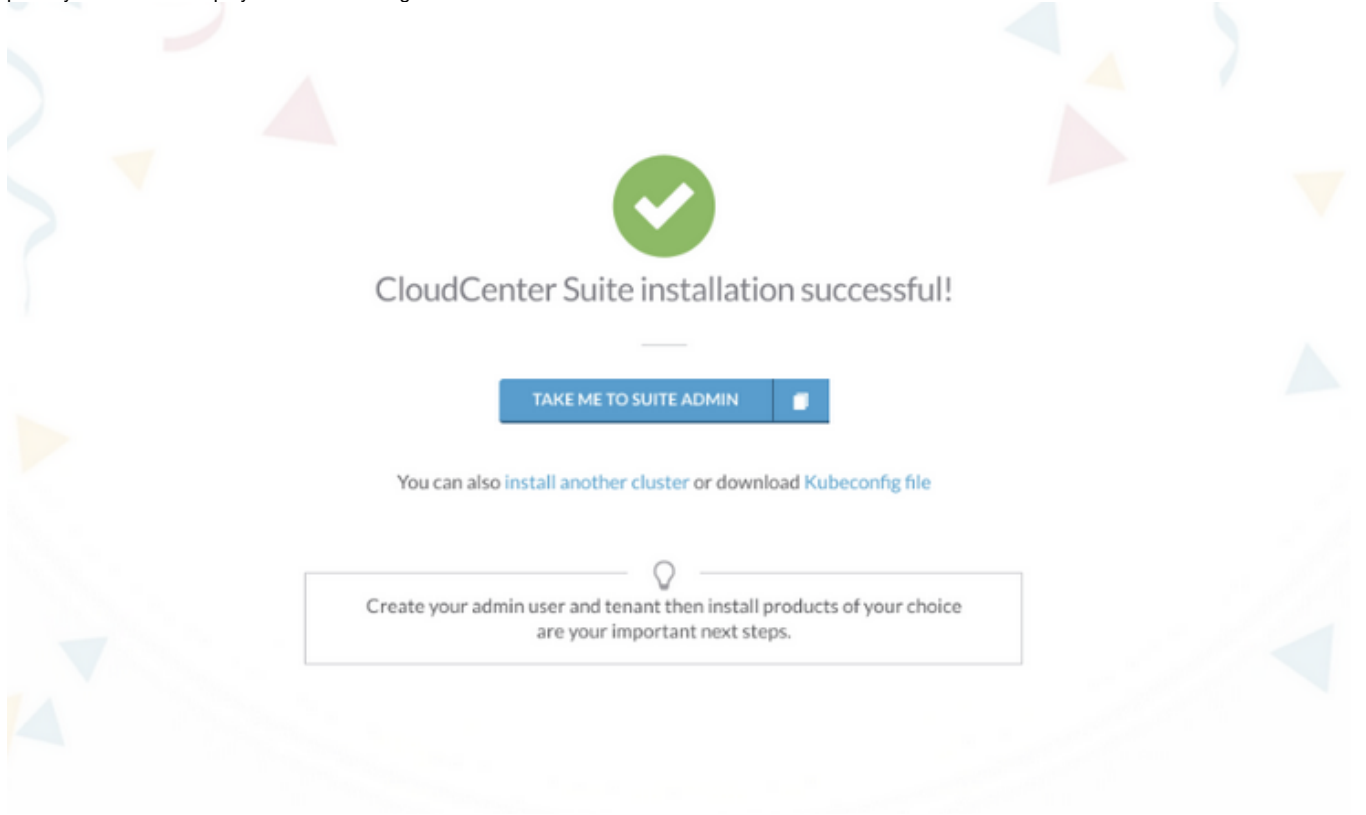
- Verify if the IP pool can accommodate additional workloads.
- Select your instance type according to the following dependencies based on your instance type selection, the installer displays the error or success information in the UI.
 - The CloudCenter Suitesetup requires 3 primary servers.
 - The CloudCenter Suite dynamically calculates the number of application VMs (workers).
- Do not use 172.18.0.1/16 for the installer instance as this IP address is used by the Docker/Kubernetes setup.
- NodePort: If you set the type field to NodePort, the Kubernetes control plane allocates a port from a range specified by service-node-port-range flag (default: 30000-32767). Refer to <https://kubernetes.io/docs/concepts/services-networking/service/> for additional details.

NTP Requirements

You must either set the Network Time Protocol (NTP) time at the datacenter level or at the time of installation.

If set at installation time, then verify that the network can access the NTP server.

The time for all worker and primary server nodes is synced with the *primary* controller node. The *primary* controller node is the instance used to launch the CloudCenter Suite identified by the link that takes you to the Suite Admin UI (Take Me to Suite Admin). This link contains the IP address of the primarycontroller as displayed in the following screenshot.



The Suite Installer Dashboard

After launching the installer, navigate to the IP address of your VM in a supported browser. This presents the Suite Installer Dashboard. The Suite Installer Dashboard has the following options:

- [New Cluster Installation](#)
- [Existing Cluster Installation](#)
- [Upgrade Kubernetes Cluster](#)

Without Internet Access

The Cisco Repository is used to host Cisco-related files and packages for various purposes. You may need to install the CloudCenter Suite in an environment that does not have internet access. If so, you need to set up an [Air Gap Installation](#).

New Cluster Installation

- [VMware vSphere Installation](#)
- [OpenStack Installation](#)

Install the CloudCenter Suite on a New Kubernetes Cluster

Once you access the Suite Installer Dashboard (see [Prepare Infrastructure](#)), you can install a new cluster and launch nodes for the new Kubernetes cluster.

Navigating to the Cluster IP When "Take Me Home" Page Does not Work

After upgrading your Kubernetes version to 1.18.12, the Take Me to the Suite Admin button does not navigate to the cluster-ip. To navigate to the cluster-ip, perform the following workaround:

1. Run the following command to display the allocation mode:

```
echo $(kubectl get cmk8s-mgmt.cluster-n cisco -ojsonpath="{.data.data}") | base64 -d | grep "vsphere_ip_allocation_mode"
```

2. From the output, look for the static IP or DHCP allocation mode.

For the static IP case:

Run the following command and note the value of the External-IP and port 443 mapping:

```
kubectl get svc -n cisco | grep common-framework-nginx-ingress-controller
```

Note the Suite Admin URL will be `https://<external_ip_address>:<443_port>`. For example, a Suite Admin URL of `https://10.10.124.157:30037` would appear in the command this way:

```
kubectl get svc -n cisco | grep common-framework-nginx-ingress-controller common-framework-nginx-ingress-controller LoadBalancer 10.99.116.114 10.10.124.157 80:32165/TCP,443:30037/TCP 17d
```

Option 1 for the DHCP case: The service type for common-framework-nginx-ingress-controller is NodePort;

Find the Master IP and note the master IP address. The Master_VIP is stored in the `cmk8s-mgmt.cluster-n cisco` namespace. To obtain the master IP address, run this command:

```
echo $(kubectl get cmk8s-mgmt.cluster-n cisco -ojsonpath="{.data.data}") | base64 -d | grep "master_vip"
```

Option 2 for the DHCP case: The master IP address is the external IP address of the first master node.

You can find the master IP address by logging into the vSphere/OpenStack console and running the following command.

```
kubectl get svc -n cisco | grep common-framework-nginx-ingress-controller
```

Note the Suite Admin URL is `https://<master_vip>:<443_port>`

For example, a Suite Admin URL of `https://10.12.104.77:30750` would appear this way:

```
kubectl get svc -n cisco | grep common-framework-nginx-ingress-controller common-framework-nginx-ingress-controller NodePort 10.111.85.45 <none> 80:30719/TCP,443:30750/TCP 20h
```

VMware vSphere Installation

VMware vSphere Installation

Overview

These instructions outline the end-to-end steps for installing CloudCenter Suite in a vSphere environment. In order to ensure successful installation, please take special care to review and understand the required prerequisites below in **PART 1** and **PART 2** of *Prepare/Verify the Installation Environment and Infrastructure*.

Installation Process

- [Prerequisites: Prepare/Verify the Installation Environment and Infrastructure - PART 1](#)
- [Import the Suite Installer into vSphere](#)
- [Prerequisites: Prepare/Verify the Installation Environment and Infrastructure - PART 2](#)
- [Deploy CloudCenter Suite into vSphere](#)

Prerequisites: Prepare/Verify the Installation Environment and Infrastructure - PART 1

In order to ensure a successful installation of CloudCenter Suite into a vSphere environment, the following steps can be used to verify and/or appropriately configure the environment and infrastructure.

1. Ensure the vSphere Datastore being used for installation meets the following requirements:

- The Datastore should be directly under the vSphere Datacenter.



The Datastore should **NOT** be part of a Datastore Cluster.

- The Datastore should be reachable from the workers and primary servers in the CloudCenter Suite cluster.
- Verify that the network and IP assigned to workers and primary servers in the CloudCenter Suite cluster can reach this datastore.
- The Datastore should have adequate permissions to be managed by the previously created user.
- Ideally, the Datastore utilized for the VM Installer and Tenant image should be the same to ensure the quickest possible installation.

2. The installation process requires a vSphere User with specific Permissions. For users who do not want to use the default administrator, use the following steps to create a new Role and User for the installation.

Step 1: In vSphere, login into vSphere as an administrator user. Navigate to **Home > Administration > Roles** and create a Role by providing the following privileges to this role -

- Datastore.Allocate space
- Datastore.Browse datastore
- Datastore.Low level file operations
- Datastore.Remove file
- Folder. Create folder
- Global.Manage Custom Attributes
- Global.Set custom attribute
- Network.Assign network
- Resource.Apply recommendation
- Resource.Apply vApp to resource pool
- Resource.Apply virtual machine to resource pool
- Storage views. View
- Tasks.Create task
- Tasks.Update task
- Virtual machine (Check all the permissions under this Privilege).
- vApp.Import
- vApp.Power off
- vApp.Power on
- vApp.Suspend
- vApp.vApp application configuration
- vApp.vApp instance configuration
- vApp.vApp managed By configuration
- vApp.vApp resource configurationIn

Step 2: Navigate to **Home > Administration > User and Groups**. Click on the + icon and create a new user. Remember the username and password - these will be used in subsequent steps.

Step 3: Click on **Global Permissions**. Click on the + icon to open *Global Permission Root - Add Permission*. Click on **Add** to map the previously created user to the Role created in Step 1 - make sure to click **Propagate to children**.

3. The Suite Installer requires a single IP address. For environments without support for DHCP, users will need to create a VM Customization Specification to assign a Static IP to the Suite Installer.

Step 1: In vSphere, login into vSphere as an administrator user. Navigate to **Home > Policies and Profiles**. Click on the + icon to create a new *VM Customization Specification*.

Step 2: In the *New VM Customization Specification* wizard, enter a name and then select **Linux** for the *Guest OS*. Click **Next** to proceed.

New VM Customization Specification

1 Name and target OS | Name and target OS
Specify a unique name for the VM customization specification and select the OS of the target VM.

2 Computer name
3 Time zone
4 Network
5 DNS settings
6 Ready to complete

VM Customization Specification

Name:

Description:

vCenter Server:

Guest OS

Target guest OS: Windows Linux

Use custom SysPrep answer file
 Generate a new security identity (SID)

CANCEL BACK NEXT

Step 3: For the *Computer name* step of the wizard, ensure **Use the virtual machine name** is checked, and enter the **Domain name** if applicable. Click **Next** to proceed.

New VM Customization Specification

- ✓ 1 Name and target OS
- 2 Computer name**
- 3 Time zone
- 4 Network
- 5 DNS settings
- 6 Ready to complete

Computer name
Specify a computer name that will identify this virtual machine on a network.

Use the virtual machine name ⓘ
 Enter a name in the Clone/Deploy wizard
 Enter a name

Append a unique numeric value. ⓘ

Generate a name using the custom application configured with the vCenter Server
 Argument _____

Domain name

CANCEL BACK **NEXT**

Step 4: For the *Time zone* step of the wizard, select the appropriate time zone and then click **Next** to proceed.

Step 5: For the *Network* step of the wizard, select **Manually select custom settings** and then click on the three dots to

New VM Customization Specification

- ✓ 1 Name and target OS
- ✓ 2 Computer name
- ✓ 3 Time zone
- 4 Network**
- 5 DNS settings
- 6 Ready to complete

Network
Specify the network settings for the virtual machine.

Use standard network settings for the guest operating system, including enabling DHCP on all network interfaces
 Manually select custom settings

Description	IPv4 Address	IPv6 Address
<input checked="" type="checkbox"/> <div style="border: 1px solid gray; padding: 2px; display: inline-block;"> Edit Delete </div>	Use DHCP	Not used

CANCEL BACK **NEXT**

Step 6: Once the *Edit Network* wizard appears, select **Use custom settings** and then input the Static IPv4 IP address, including the appropriate subnet and gateway - this is the IP address the user will use to access the Suite Installer post-installation. Click **OK** to proceed, and then click **Next** to move onto the next step of the wizard.

Step 7: For the *DNS settings* step of the wizard, input the necessary information for DNS. Click **Next** to proceed.

Step 8: Verify the configuration and then click **Finish** to complete the creation of the VM Customization Specification.

Name	suite-5.2.0-RC1.3
OS type	Linux
Computer name	Use Virtual Machine name
Domain name	galaxy.cisco.com
Time zone	US/Eastern
Hardware clock	Set to UTC
Network type	Custom
NIC1 IPv4	10.2.1.60
NIC1 IPv6	Not used
Primary DNS server	10.2.1.172


- The CloudCenter Suite installation process recommends that the Suite Installer uses the same NTP server as the ESX cluster. The NTP server can be retrieved from the ESX host by navigating to **Configure > System > Time Configuration**. Remember the IP address of the NTP server - it will be used in subsequent steps.

This completes **PART 1** of the **Prepare/Verify the Installation Environment and Infrastructure**.

Import the Suite Installer into vSphere

- Download the Installer OVA from software.cisco.com.

2. Login into vSphere as an administrator or with a user with the appropriate permissions as outlined above in *PART 1 of Prepare/Verify the Installation Environment and Infrastructure*. Click on **VM and Templates**, and then select the vSphere Datacenter where the Installer needs to be uploaded. Right-click and select **Deploy OVF Template...**



3. In the *Deploy OVF Template* wizard, select **Local File** and open the previously downloaded OVA from your computer's file browser. Click **Next** to proceed.
4. For the *Select name and folder* step of the wizard, select a folder directly underneath the Datacenter - see below screenshot for an example. Click **Next** to proceed.



You **MUST** select an installation folder, however do **NOT** select a sub-folder. This requirement is the same for uploading the Suite Installer, as well as selecting an installation directory during the installation of CloudCenter Suite. This behavior applies to CloudCenter Suite 5.2.1 and earlier versions.

Effective 5.2.2, CloudCenter Suite supports the following changes:

- VMware environments can configure Clusters, DataStores, and/or Networks under a sub-folder. For example, sub-folder /Cluster , sub-folder/Datastore , sub-folder/Network
- You can install a CloudCenter Suite cluster under any sub-folder

Deploy OVF Template

1 Select an OVF template Select a name and folder
2 Select a name and folder Specify a unique name and target location
3 Select a compute resource
4 Review details
5 Select storage
6 Ready to complete

Virtual machine name:


Select a location for the virtual machine.

Select an **Installation Folder** →


- ✓ GFFA-vCenter.galaxy.cisco.com
 - ✓ GFFA-DC
 - ✓ **1_Install-Here**
 - > 2_Not-Here **Do NOT select a sub-folder**
 - > CCP-ONE
 - > CCP-UCS-6
 - > CCP-UCS-GPU
 - > CCS-UCS
 - > Cisco CP
 - > CliqrTemplates
 - > CliqrUser-1
 - > CloudCenter Suite
 - > Discovered virtual machine
 - > HX1-CCP
 - > HX1-VMs
 - > HX2-CCP
 - > HX2-Infrastructure
 - > HX2-VMs

CANCEL BACK NEXT

5. For the *Select resource* step of the wizard, select an ESX Host from the Cluster. Click **Next**, and wait for the validation checks to complete. Click **Next** again to proceed to the *Select storage* step of the wizard.



6. For the *Select storage* step of the wizard, select an Datastore with necessary permissions as outlined above in *PART 1 of Prepare/Verify the Installation Environment and Infrastructure*. Click **Next** to proceed.

 **Reminder:** The Suite Installer does **NOT** support Datastore Clusters.

 **Recommendation:** Select Thin for the Virtual Disk Format.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- 5 Select storage**
- 6 Select networks
- 7 Customize template
- 8 Ready to complete

Select storage
Select the storage for the configuration and disk files

Encrypt this virtual machine (Requires Key Management Server)

Select virtual disk format: **Thin Provision** Recommended

VM Storage Policy: Datastore Default

Name	Capacity	Provisioned	Free
GFFA-Mgmt-HX1-Datastore-1	1.46 TB	777 GB	863.03 GB
GFFA-Mgmt-HX1-Datastore-2	1.46 TB	171 TB	653.25 GB
SpringpathDS-WZP2201M10	216 GB	715 GB	208.85 GB
SpringpathDS-WZP2201M14	216 GB	715 GB	208.85 GB
SpringpathDS-WZP22020D89	216 GB	715 GB	208.85 GB

Reminder:
Datastore Clusters are **NOT** supported

Compatibility
✓ Compatibility checks succeeded.

7. For the *Select networks* step of the wizard, from the drop-down select the appropriate network for the installer management interface - if necessary, this can be modified later. Click **Next** to proceed.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 Select storage
- 6 Select networks**
- 7 Customize template
- 8 Ready to complete

Select networks
Select a destination network for each source network.

Source Network	Destination Network
vlan1004	gffa-mgmt-401

1 items


IP Allocation Settings

IP allocation: Static - Manual

IP protocol: IPv4


8. For the *Customize template* step of the wizard, use the following table to complete the form:

Field	Description	Condition
Unique ID	This value must be unique within the vSphere networking domain. This field will be used to generate the hostname.	Required


Password	This value will be used to allow password-based authentication to the Installer VM via the vSphere Console.	Recommended
SSH Public Key	<p>This value will be used to allow key-based authentication with the Installer VM via SSH. The encryption formats supported are ecdsa and ed25519.</p> <div style="border: 1px solid green; padding: 5px;"> For additional information - including instructions on how to generate a SSH key - please consider visiting SSH Documentation.</div>	Recommended
Hostname	This value must be unique within the vSphere networking domain. This field will be used to generate the hostname.	Required

▼ Uncategorized	6 settings
Encoded user-data	In order to fit into a XML attribute, this value is base64 encoded. It will be decoded, and then processed normally as user-data. _____
SSH public keys Recommended	This field is optional, but indicates that the instance should populate the default user "authorized_keys" file with this value. _____
Default user's password Recommended	If set, the default user password will be set to this value to allow password based login. The password will be good for only a single login. If set to the string "RANDOM", a random password will be generated and written to the console. _____
A unique ID for this VM instance REQUIRED Must be Unique	Specifies the instance ID. This is required and is used to determine if the machine should take "first boot" actions. default-instance-id _____
Hostname REQUIRED Must be Unique	Specifies the hostname of the VM instance. default-hostname _____
URL to seed instance data from	This field is optional, but indicates that the instance should "seed" user-data and meta-data from the given URL. If set to "http://tinyurl.com/sm-", then meta-data will be pulled from "http://tinyurl.com/sm-meta-data" and user-data from "http://tinyurl.com/sm-user-data". Leave this empty if you do not want to seed data from a URL.

- Click **Next** and then **Finish** to proceed. The OVA will start uploading - this will take approximately 5-10 minutes.
- Once the OVA is finished uploading, create a VM Template from the uploaded installer image. This template can be used in future installations. Right-click on the OVA and select **Template > Convert to Template**. Click **Yes** to confirm. Once the wizard is complete, the convert will take approximately 5-10 minutes.



11. Once the converting completes, right-click on the OVF template and select **New VM from this Template...**



i The following steps are similar to **Steps 4-6**. Remember that the following behavior applies to CloudCenter Suite 5.2.1 and earlier versions:

- You **MUST** select an installation folder, however do **NOT** select a sub-folder.
- Select the same Datacenter Cluster or Host as the Suite Installer.
- The Suite Installer does **NOT** support Datastore Clusters.

Effective 5.2.2, CloudCenter Suite supports the following changes:

- VMware environments can configure Clusters, DataStores, and/or Networks under a sub-folder. For example, sub-folder /Cluster , sub-folder/Datastore , sub-folder/Network
- You can install a CloudCenter Suite cluster under any sub-folder



For environments **NOT** using DHCP, select **Customize the operating system**.

For the *Select clone options* step of the wizard, check **Power on virtual machine after creation**. Click **Next** to proceed.

suite-5.2.0-RC1.3-template - Deploy From Template

- ✓ 1 Select a name and folder
- ✓ 2 Select a compute resource
- ✓ 3 Select storage
- 4 Select clone options**
- 5 Customize guest OS
- 6 Ready to complete

Select clone options

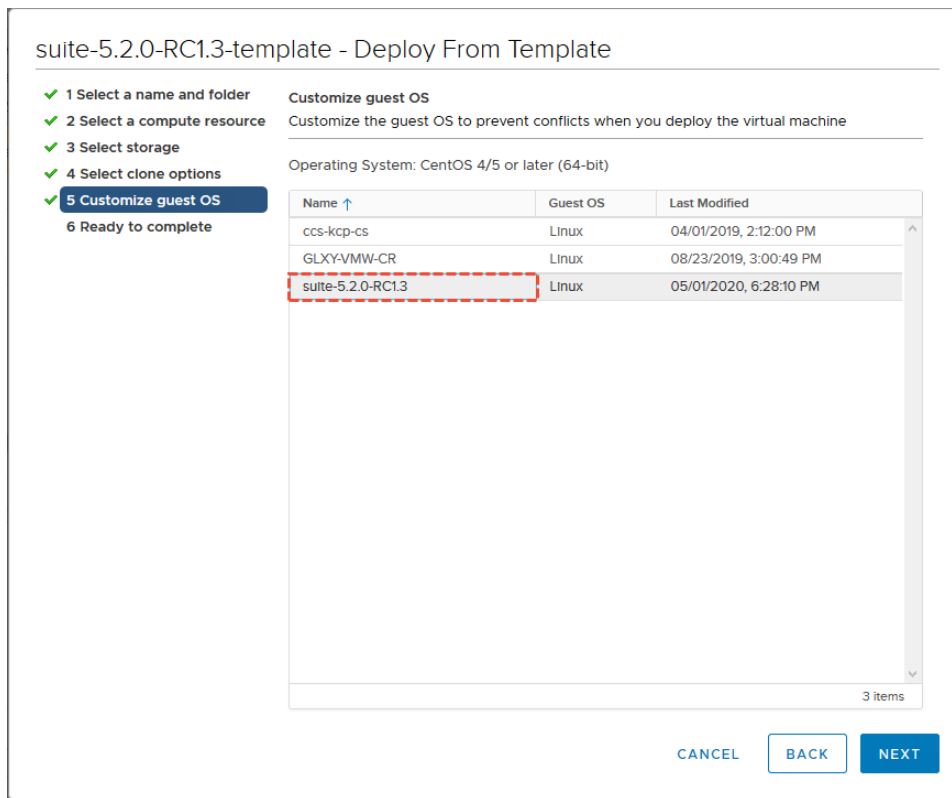
Select further clone options

- Customize the operating system **Optional: For non-DHCP enabled environments**
- Customize this virtual machine's hardware
- Power on virtual machine after creation

CANCEL BACK NEXT

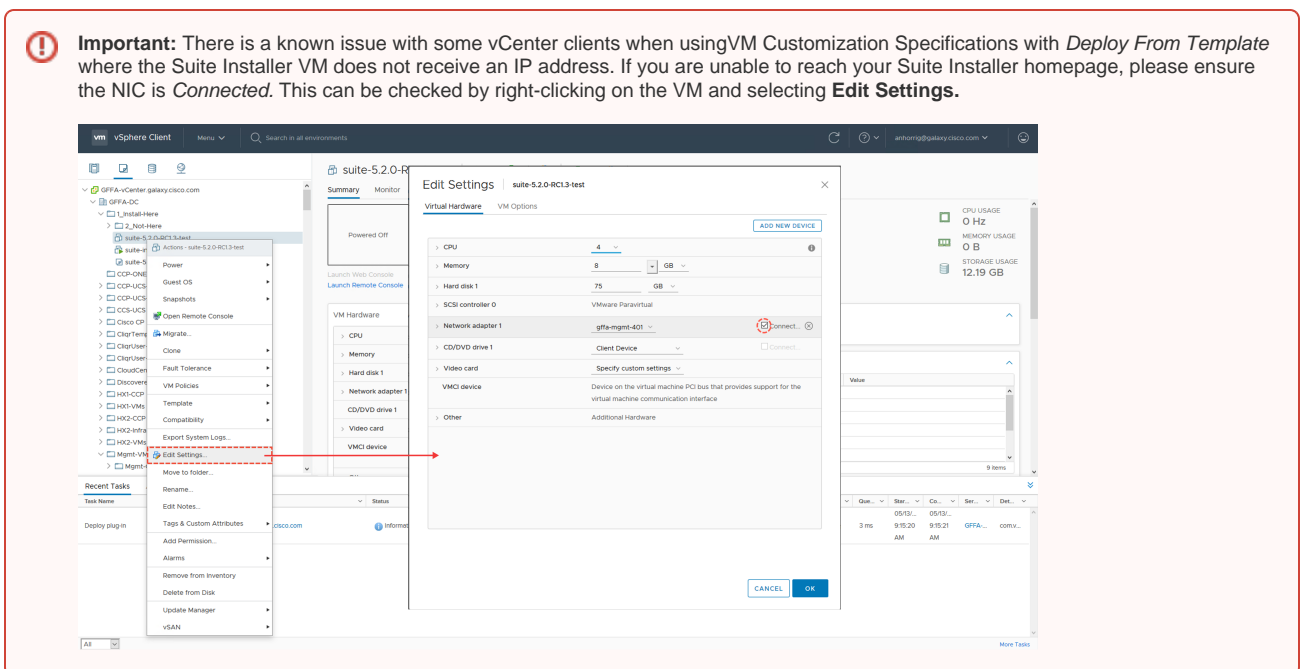
13. The Suite Installer requires a single IP address. For environments without support for DHCP, users will need to attach a *VM Customization Specification* to assign a Static IP to the VM Installer. The creation of the VM Customization Specification was previously outlined above in *PA RT 1 of Prepare/Verify the Installation Environment and Infrastructure*.

Step 1: For the *Customize guest OS* part of the wizard, select the previously created VM Customization Specification and click **Next** to proceed.



14. Review the details of the wizard and then click **Finish** to proceed with the creation of the Suite Installer VM. The creation of the VM will take approximately 5-10 minutes.

Important: There is a known issue with some vCenter clients when using VM Customization Specifications with *Deploy From Template* where the Suite Installer VM does not receive an IP address. If you are unable to reach your Suite Installer homepage, please ensure the NIC is *Connected*. This can be checked by right-clicking on the VM and selecting **Edit Settings**.



suite-5.2.0-RC1.3-template - Deploy From Template

- ✓ 1 Select a name and folder
- ✓ 2 Select a compute resource
- ✓ 3 Select storage
- ✓ 4 Select clone options
- ✓ 5 Customize guest OS
- 6 Ready to complete**

Ready to complete
Click Finish to start creation.

Provisioning type	Deploy from template
Source template	suite-5.2.0-RC1.3-template
Virtual machine name	suite-installer-5.2.0-RC1.3
Folder	1_Install-Here
Cluster	GFFA-MgmtHX1-Cluster
Datastore	GFFA-MgmtHX1-Datastore-1
Disk storage	Same format as source
Guest OS customization specification	suite-5.2.0-RC1.3

CANCEL BACK FINISH

This completes the import/upload of the Suite Installer into VMware.

Prerequisites: Prepare/Verify the Installation Environment and Infrastructure - PART 2

In order to ensure a successful installation of CloudCenter Suite into a vSphere environment, the following steps can be used to verify and/or appropriately configure the environment and infrastructure.

- The installation process assumes internet connectivity to certain domains. When installing CloudCenter Suite into environments residing behind a proxy, please ensure the following domains are entirely accessible. Remember the proxy information - this will be used during the installation of CloudCenter Suite.

Note: The Installer VM supports HTTP and HTTPS proxies, with or without username and password. The proxy must support TLS 1.2.

Warning: Several of the following links might perform redirects. Please ensure your proxy and firewall are configured to allow redirects of the following URLs.


Proxy URL	Description
https://devhub.cisco.com http://devhub.cisco.com https://devhub-docker.cisco.com http://devhub-docker.cisco.com	Repository for Cisco CloudCenter Suite Docker Charts
https://gcr.io http://gcr.io	Repository for Cisco CloudCenter Suite Helm Charts

https://storage.googleapis.com http://storage.googleapis.com	Repository for Cisco CloudCenter Suite Tiller Image
Other	The Suite Installer may require additional connections to the installation environment (for example, vCenter, Hyperflex Data Platform, AWS Console, and so forth) Please ensure your cloud target is reachable via the proxy!

A Note on Offline Clusters

While CloudCenter Suite 5.2 offers a completely air gapped environment, your CCS cluster will require access to the URLs in the above table if your internet access is via a proxy environment. However, as the offline solution is a completely air gapped environment and you do not need to add URLs to your acceptable list of URLs when using the [Air Gap Installation](#) approach.

Users can use an existing Linux VM to test their proxy configurations. The following steps outline how to test a proxy on an Ubuntu VM.

 **Note:** These steps may vary depending on the user's installation environment and proxy configuration.

Step 1: Configure the proxy on the VM:

```
export http_proxy=http://<proxy value> (HTTP Proxy)
export https_proxy=https://<proxy value> (HTTPS Proxy)
export http_proxy=http://<username>:<password>@<proxy value> (HTTP w/ Authentication)
export https_proxy=https://<username>:<password>@<proxy value> (HTTPS w/ Authentication)
```


Step 2: Use this command to login to the CloudCenter Suite docker registry. If this command fails, there might be an issue with the proxy configuration:

```
docker login -u "multicloudsuite.gen" -p
"AKCp5aTvLmuvA2dleRkiehsSAYSuWZiyEv76bczZWzHe7bq5W96drHsmUzKus6v2ZsYXqMFje"devhub-docker.cisco.com
/multicloudsuite-release
```


Step 3: Use this command to download a docker image from the CloudCenter Suite registry. If this command fails, there might be an issue with the proxy configuration:

```
sudo docker pull gcr.io/kubernetes-helm/tiller:v2.12.3
```

- In vSphere environments with more than one Datacenter, users are required to create a [Resource Pool](#). This is true for both uploading the Suite Installer, as well as picking an installation environment.

 The Resource Pool should **NOT** be "nested" and part of another Resource Pool.


- In order to improve installation time, it is also recommended to upload the Tenant Image to the same Datastore and Datacenter as the installation target. The Tenant Image is cloned and used to deploy the Worker Nodes in the Kubernetes control-plane. The Tenant Image can be downloaded from software.cisco.com. When the Tenant Image is not uploaded prior to installation, the Suite Installer will attempt to upload clone and upload a copy of the image from the Installer OVA.

 The name of uploaded OVA **MUST** have a prefix of **"CCS"**.

This completes **PART 2** of the [Prepare/Verify the Installation Environment and Infrastructure](#)

Deploy CloudCenter Suite into vSphere

- Once the Suite Installer VM finishes deploying and powering on, navigate to appropriate IP address to start the installation process of CloudCenter Suite. For DHCP-enabled deployments, the IP address can be found on the VMware console. Click on **New Cluster** to proceed.

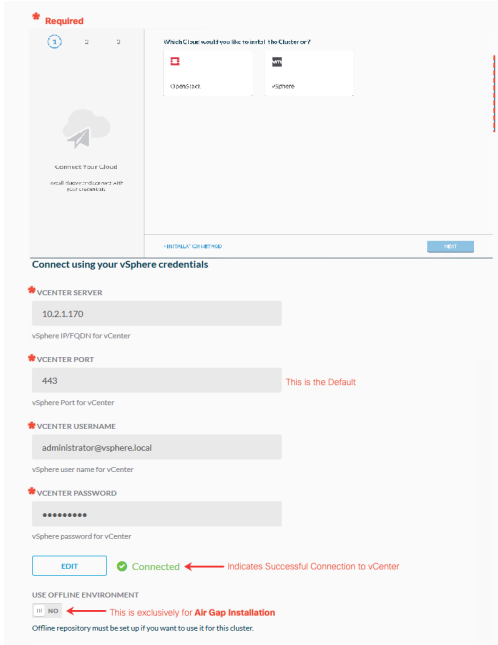
 **Note:** Depending on the browser, users may need to dismiss a self-signed certificate error before proceeding.



2. Select **vSphere** and then complete the wizard with the required information (IP address and login credentials). Click **Connect** verify the connectivity. If the connectivity check successfully completes, click **Next** to continue.

Reminder: Use the account created above in **Step 2of Prepare/Verify the Installation Environment and Infrastructure - PART 1** to connect to vCenter.

Note: Depending on the deployment environment, the selections/inputs for the following steps may vary.



Cisco CloudCenter Suite supports Air Gap Installations. However, in order to deploy CloudCenter Suite into environments without internet connectivity, users need to first setup an [Offline Repository](#). Once the repository is setup, users can select **Yes** for *Use Offline Repository*. Provide the login information and then click **Validate** to proceed.

*** Required**

USE OFFLINE ENVIRONMENT

YES Indicates an **Air Gap Environment**

Offline repository must be set up if you want to use it for this cluster.

*** OFFLINE REPOSITORY FQDN / IP ADDRESS**

*** PASSWORD**

VALIDATE

3. On the next step of the installer, select/input the necessary *vSphere Placement Properties* for your environment.

vSphere Configuration			
Field	Input	Condition	Notes
Datacenter	Select the vSphere Datacenter for installation	Required	
Cluster	Select the vSphere Cluster for installation	Required	
Resource Pool	Select the vSphere Resource Pool for allocation of resources	Optional	<div style="border: 1px solid red; padding: 5px;"> Reminder: This field is required for environments with more than one VMware Datacenter. </div>
Datastore	Select the vSphere Datastore for installation	Required	<div style="border: 1px solid blue; padding: 5px;"> Recommendation: Select the same Datastore as the Suite Installer and Tenant Image. </div>
Network	Select the vSphere Network for installation and connectivity between the various nodes/services of the Kubernetes cluster	Required	
CCS VM Tenant Image	Select the installation image used to create the Kubernetes cluster	Optional	<p>The Suite Installer includes a default Kubernetes cluster image (<i>CCS-version-Base-Image</i>). This image will be automatically used whenever this field is left empty.</p> <div style="border: 1px solid blue; padding: 5px;"> Recommendation: For slow environments, upload the Tenant Image to same folder as the Suite Installer - ensure the name of the image is prefaced with "CCS-". This was previously outlined above in Step 3 of <i>Prepare /Verify the Installation Environment and Infrastructure - PART 2</i>. </div>
Cluster Folder	Select the installation directory	Required	

The following screenshot is an example. Selections and values may differ between different installation environments.

*** Required**

What are your placement properties?

*** DATACENTER**

GFFA-DC

vSphere datacenter where the kubernetes cluster will be launched

*** CLUSTER**

GFFA-MgmtHX1-Cluster

vSphere cluster where the kubernetes cluster will be launched

RESOURCE POOL

Optional

Required for environments with more than one Datacenter

vSphere resource pool. If left empty, Default: <compute cluster>/Resources

*** DATASTORE**

GFFA-MgmtHX1-Datastore-1

vSphere datastore

*** NETWORK**

gffa-mgmt-401

vSphere network

CCS VM TENANT IMAGE

Recommended

Cisco Kubernetes 1.16.3 CCS VM Tenant Image. If not selected, the installer will attempt to upload an included template.






*** CLUSTER FOLDER**


Mgmt-VMs

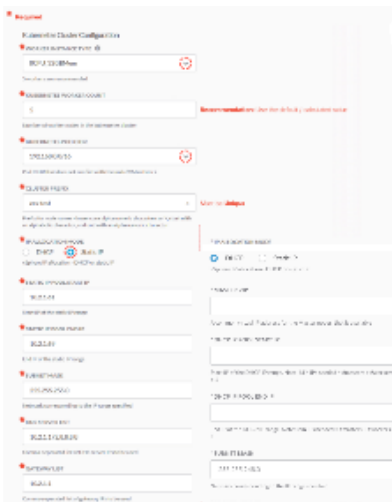
vSphere folder which will contain the kubernetes cluster nodes

4. Scroll down to the next step of the installer and select/input the necessary values for *Kubernetes Cluster Configuration*.

Kubernetes Cluster Configuration			
Field	Input	Condition	Notes
Worker Instance Type	Select the Instance Type with the right amount of CPU / Memory	Required	See Prepare Infrastructure > <i>Resource Requirements for CloudCenter Suite Modules</i> for additional context. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>Recommendation: <i>8CPU_32GBMem</i> will deploy the least number of nodes.</p> </div>
Kubernetes Worker Count	The number of nodes is automatically calculated based on the selection made for <i>Worker Instance Type</i>	Required	See Prepare Infrastructure > <i>Resource Requirements for CloudCenter Suite Modules</i> for additional context. Users can opt to increase or decrease the number of nodes deployed during installation. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px; background-color: #fff9c4;"> <p>! The IP address requirements will change depending on the number of Worker Nodes selected during installation. For example:</p> <ul style="list-style-type: none"> • If the instance type is <i>8CPU_32GBMem</i>, then 5 workers are created and the total static IPs required for this environment are 7 IPs (4 worker VMs, and 3 primary servers). • If the instance type is <i>8CPU_24GBMem</i> memory, then 5 workers are created and the total static IPs required for this environment are 8 IPs (5 worker VMs, and 3 primary servers). • If the instance type is <i>8CPU_16GBMem</i>, then 7 workers are created and the total static IPs required for this environment are 9 IPs (6 worker VMs, and 3 primary servers). • If the instance type is <i>4CPU_16GBMem</i>, then 9 workers are created, so the static IPs required for this environment are 11 IPs (8 worker VMs, and 3 primary servers). </div>

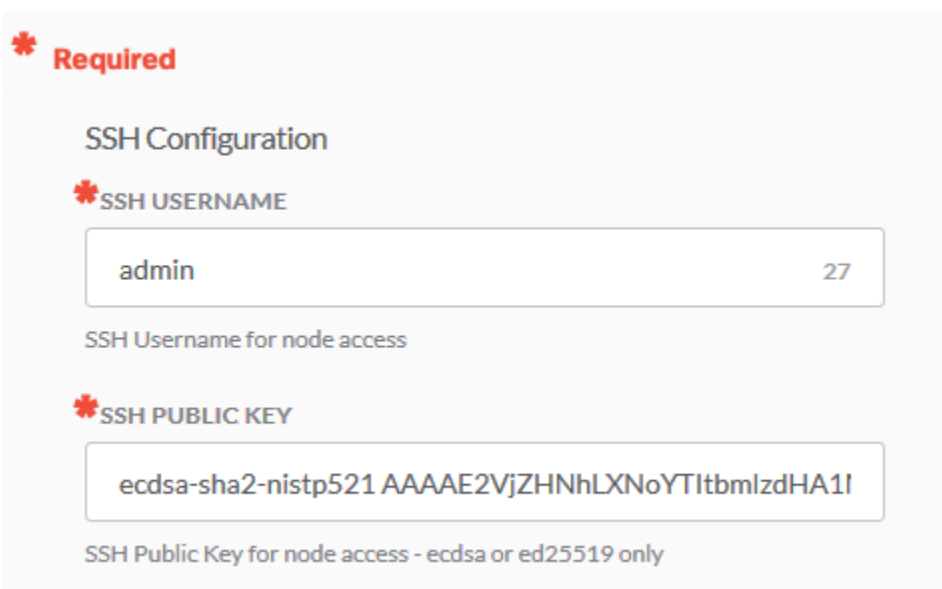
Kubernetes Pod CIDR	Select the IP address block for internal networking between the pods running on each of the nodes	Required	<p>This address space is INTERNAL, and is not routable outside of the Kubernetes Cluster.</p> <div style="border: 1px solid #ffc107; padding: 5px; margin-top: 10px;">  Warning: This address block should NOT conflict with the subnet or IP addresses used for the nodes. </div>
Cluster Prefix	Enter any unique value	Required	<p>Used to identify which VMs / nodes are part of a Kubernetes cluster.</p> <div style="border: 1px solid #ffc107; padding: 5px; margin-top: 10px;">  Value must be lowercase, and must start and end with an alphanumeric character. Input field supports "-" (hyphens) but not "_" (underscores). </div>
IP Allocation Mode	Select either DHCP or Static IP	Required	<div style="border: 1px solid #c6c8ca; padding: 10px; margin-bottom: 10px;">  Note: Besides the assignment of IP addresses, the following selection also determines how Services are exposed outside of the cluster. When DHCP is selected, Kubernetes will use a load balancing service (MetalLB). However, when Static IP is selected, Kubernetes will use NodePorts. </div> <ul style="list-style-type: none"> • DHCP: During the boot process, IP addresses will be allocated via DHCP server. <ul style="list-style-type: none"> • <i>Master VIP</i> - Common virtual IP address shared by the primary server nodes. Users can access the CloudCenter Suite login with this address. • CCS v5.2.3 required users to only provide only a master VIP to create a cluster, but • CCS v5.3.0 now requires users to provide all the network information like static IP case to create a cluster & its resources. DHCP_Pool_Start_IP, DHCP_Pool_End_IP, Subnet mask, Gateway IP & DNS Addresses. <div style="border: 1px solid #dc3545; padding: 5px; margin-top: 10px;">  The IP address for the Master VIP must be unique, and not available to DHCP. </div> • Static: During the boot process, IP addresses will be allocated from a user-defined pool. <div style="border: 1px solid #dc3545; padding: 10px; margin-top: 10px;">  Important: Please note the following requirements when allocating a block of IP addresses - <ul style="list-style-type: none"> • The block of IP addresses must cover the number of <i>Nodes (Workers + primary servers)</i> and <i>(4) additional services</i>. However, we recommend users define larger pools (50% more) to allow for future scalability. <ul style="list-style-type: none"> • e.g. If the instance type is <i>8CPU_32GBMem</i>, 7 IP addresses are required for the nodes and 4 IP addresses are required for the additional services. Therefore total minimum required is 11 IP addresses. • The block of IP addresses for the user-defined pool must be unique. Verify network reachability before proceeding - the installation will fail without complete connectivity. </div> <ul style="list-style-type: none"> • <i>Static IP Pool Start IP</i> - The first IP address in the pool. • <i>Static IP Pool End IP</i> - The last IP address in the pool. • <i>Subnet Mask</i> - The subnet mask of the address pool. • <i>DNS Server List</i> - The available DNS servers in the environment. • <i>Gateway List</i> - The subnet's "Default Gateway".

 The following screenshot is an example. Selections and values may differ between different installation environments.



Scroll down to the next step of the installer and input the necessary values for *SSH Configuration*. This configuration will be used to allow key-based authentication with the worker and primary server nodes via SSH.

SSH Configuration			
Field	Input	Condition	Notes
SSH Username	Enter valid username	Recommended	<p>This is a user-assigned field to identify the user for SSH access into worker(s)/primary server(s).</p> <div style="border: 1px solid red; padding: 5px; margin-top: 10px;"> The username must NOT be <i>root</i> or <i>cloud-user</i>. </div>
SSH Public Key	Enter valid SSH key	Recommended	<p>The encryption formats supported are ecdsa and ed25519.</p> <div style="border: 1px solid green; padding: 5px; margin-top: 10px;"> For additional information - including instructions on how to generate a SSH key - please consider visiting SSH Documentation. </div> <div style="border: 1px solid yellow; padding: 5px; margin-top: 10px;"> The Suite Installer does not require the SSH key to have a comment. However, any comments MUST be in the following format: <username>@<hostname> </div>



5. (Optional) Scroll down to the next step of the installer and input the necessary values for *NTP Configuration*.



Note: The NTP values should match the ESXi NTP configuration as outlined above in **Step 4** of *Prepare/Verify the Installation Environment and Infrastructure - PART 1*.

NTP Configuration

NTP SERVERS

10.2.1.1

Comma-separated list of NTP servers - hostname or IP Address. It is highly recommended to set NTP servers or pools to prevent timing issues between kubernetes nodes.

NTP POOLS

0.us.pool.ntp.org,1.us.pool.ntp.org,2.us.pool.ntp.org

Comma-separated list of NTP pools - hostname or IP Address. It is highly recommended to set NTP servers or pools to prevent timing issues between kubernetes nodes.

- (Optional) Scroll down to the next step of the installer and select/input the necessary values for *Proxy Configuration*. This configuration will define the Docker proxy settings on each worker/primary server node. When attempting to reach the internet, the nodes will use these settings for internet connectivity - this is particularly important during installation.



Reminder: Please review **Step 1** of *Prepare/Verify the Installation Environment and Infrastructure - PART 1* for additional information on the proxy configuration, including a list of required domains.



In Suite Admin 5.2.x, updating proxy configurations must be manually completed on each node. This process is not "hitless" and will require a restart of the VM.

Proxy Configuration			
Field	Input	Conditional	Notes
HTTP Proxy	Enter the IP address and port of the HTTP proxy server	N/A	For proxies requiring Username / Password, select Yes for <i>Proxy Requires User Authentication</i> . Click Validate to ensure the configuration is correct.
HTTPS Proxy	Enter the IP address and port of the HTTPS or HTTP proxy server	N/A	For proxies requiring Username / Password, select Yes for <i>Proxy Requires User Authentication</i> . Click Validate to ensure the configuration is correct. The HTTP proxy value is allowed in HTTPS proxy field. However, HTTP traffic will utilize secure channel (SSL) to connect to internet.
Bypass Proxy Settings	Enter the IP addresses or URLs of the domains you want to bypass the proxy	N/A	This configuration will define the Docker proxy settings on each worker/primary server node. The <i>Bypass Proxy Settings</i> field should be used to define which IP addresses and domains should NOT use the proxy to reach the internet. Example: localhost,10.100.96.168,*.test.example.com,.example2.com,10.1.0.0/16,127.0.0.1

Proxy Configuration

HTTP PROXY

Http proxy host with port

PROXY REQUIRES USER AUTHENTICATION

NO

HTTPS PROXY

Note: Allows HTTP or HTTPS proxy

Https proxy host with port - Proxy that supports either http or https protocol is valid ex: https://proxy.xyz:<port> or http://proxy.xyz:<port>

PROXY REQUIRES USER AUTHENTICATION

YES ← Indicates the proxy configuration requires UN/PW

* USERNAME

* PASSWORD

← Validate to ensure successful login to proxy

BYPASS PROXY SETTINGS

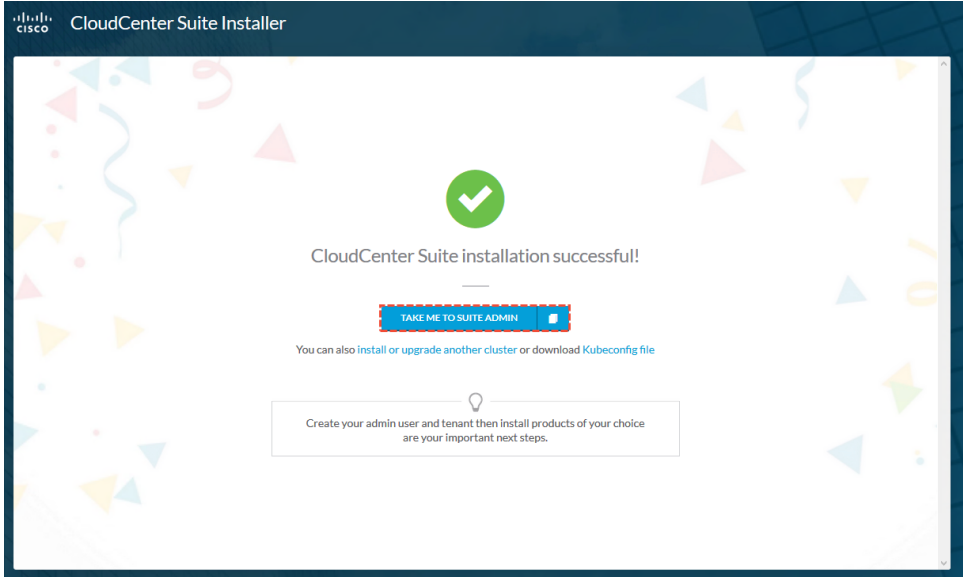
localhost,10.100.96.168,*test.example.com,example2.com,
10.1.0.0/16,127.0.0.1

These are just examples!

Provide values in a comma-separated list without spaces.

7. Once completed, click **Install** to proceed. The installation will take approximately 30-60 minutes depending on the installation environment. Click on **Take Me To Suite Admin** to continue the setup and installation.

Recommendation: Do not forget to download the **kubeconfig file** - this can be used to manage the Kubernetes nodes.



Folder Names in the VMware Environment

When you deploy CloudCenter Suite 5 in a VMware environment, you can specify a top-level folder to place the K8s nodes. The name of that folder cannot be changed after the installation. If the folder name has been changed, installation of new modules will fail. If you rename the folder name to the original name, the installation will succeed.

OpenStack Installation

OpenStack Installation

- [OpenStack Nuances](#)
- [Module Details](#)
- [Installation Process](#)

OpenStack Nuances

Verify the following OpenStack nuances:

- OpenStack newton release with at least the following service versions:
 - Cinder v2
 - Keystone v3
 - OpenStack Nova v2
 - OpenStack Networking v2
 - OpenStack Glance v2
- Ensure to add Port 6443 to the default security group as the security group created for the cluster is not automatically assigned to the load balancer created for the cluster.
- The tenant and project requirements for OpenStack Cloud are identified in the following table.

Model	Quota	Description
For all cases	2 (primary server group, worker group)	Server Groups
	Number of workers + number of primary servers	Server Group Members
	3 (API load balancers)	Load Balancers
	6 (2 for each load balancer)	Health Monitors
	6 (2 for each load balancer)	Pools
	6(2 for each load balancer)	Listeners
	3 (1 for the cluster VMs, 2 for the Kubernetes load balancer services)	Security Groups
	18	Security Group Rules
	See Prepare Infrastructure for additional details	Volume GB
	Number of workers + number of primary servers +3 for each load balancer	Ports
	Number of workers + number of primary servers	Instances
	16 GB (recommended for each worker and each primary server)	RAM
32 (recommended for each workers and each primary server)	vCPUs	
Tenant network	Floating IPs = 3	1 for each load balancer
	Networks = 1	For the tenant network
	Subnet = 1	For the tenant network
	Router = 1	For the tenant network to public network connection
Provider network	Number of workers + number of primary servers + 3 load balancers	Free IPs in the provider network

- **Network Time Protocol (NTP) must be configured this is important as the CloudCenter Suite installation can fail, if NTP is not configured or if it is wrongly configured.**



If you setup CloudCenter Suite in offline mode, you must provide valid NTP server details before you save your configuration.

Module Details

Additionally, refer to your module documentation for module-specific dependencies as identified in the following table:

Module	Documentation
Workload Manager	Cloud Overview
Action Orchestrator	Add Cloud Account
Cost Optimizer	Cloud Overview

Installation Process

To install the CloudCenter Suite on a new OpenStack cluster, perform the following procedure.

1. Verify that you have prepared your environment as listed in the *OpenStack Nuances* section above.
2. Navigate to the Suite Installer Dashboard.
3. Click **New Cluster**.
4. Click the OpenStack card.
5. To connect using OpenStack cloud credentials, enter the OpenStack Placement Property details identified in the following table.

OpenStack Placement Properties	Description
OpenStack Authentication URL	The OpenStack authentication service URL.
OpenStack Region	The OpenStack cloud region.
OpenStack Domain Name	The OpenStack account domain name.
OpenStack Project	The OpenStack project name.
OpenStack Username	The OpenStack account username.
OpenStack Password	The OpenStack account password.
OpenStack CA Certificate	The CA certificate that is required to verify an OpenStack HTTPS URL. This field is mandatory using a HTTPS URL and is not required if using a HTTP URL.



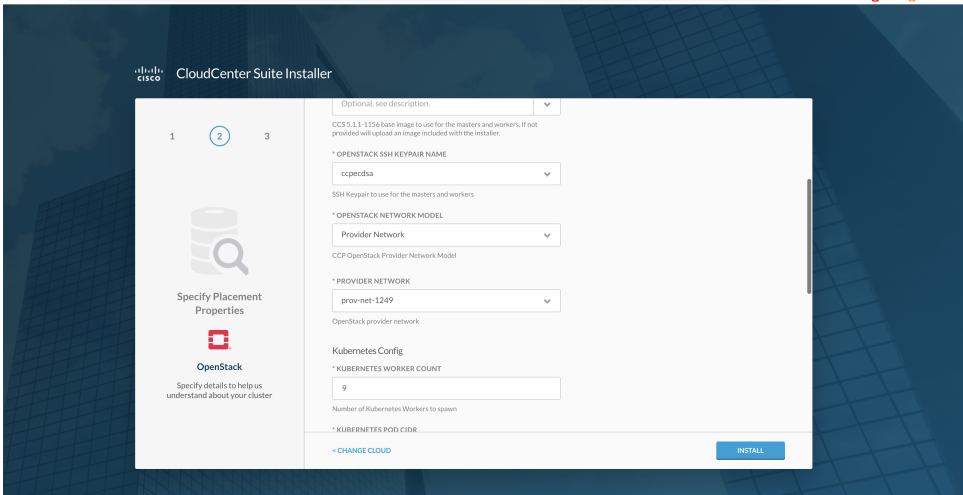
6. Click **Connect**.
7. Once the connection is validated, click **Next**.

To specify the placement properties, enter the following details.




If you setup CloudCenter Suite in offline mode, you must provide valid NTP server details before you save your configuration.

OpenStack Placement Properties	Description
Control Plane Cluster Prefix	Select the OpenStack project to which the Kubernetes cluster is deployed.
OpenStack Details	
OpenStack Flavor UUID	Select one of the existing flavors or VMs. Based on your selection, the recommended number of workers is calculated and displayed in the Kubernetes Worker Count field.

<p>OpenStack Image UUID</p>	<p>Different images will be used for the installer and the cluster launched by the installer. The installer includes a default Kubernetes cluster image (called, <i>CCS-version-Base-Image</i>) with a configurable option to override the use of this default image. The <i>CCS-version-Base-Image</i> image included in the installer is selected if you do not override the setting.</p> <p>To override the <i>CCS-version-Base-Image</i> image used by the Suite installer, be sure to add the applicable image in the OpenStack console and selected the applicable QCOW2 image from the dropdown list in this field.</p> <p>If you use the OVA installer to launch the cluster in an vSphere environment, be sure to override this field and select the applicable QCOW2 <i>CCS-version-Base-Image</i>.</p> <div style="border: 1px solid #ffc107; padding: 5px; margin-top: 10px;"> <p> If you install the CloudCenter Suite using any image other than <i>CCS-version-Base-Image</i>, the installation will fail.</p> </div>
<p>OpenStack SSH Keypair Name</p>	<p>Only SSH keys of type ssh-ed25519 or ecdsa-sha2-nistp256 are supported.</p> <div style="border: 1px solid #ffc107; padding: 5px; margin-top: 10px;"> <p> You must have at least one existing SSH-key in the selected OpenStack environment to begin the installation.</p> </div>
<p>OpenStack Network Model</p>	<p>The functional networking model for OpenStack. See https://docs.openstack.org/security-guide/networking/architecture.html for additional context.</p>
<p>Provider Network or Tenant Network</p>	<p>Provider Network Created by the OpenStack administrator on behalf of tenants and can be dedicated to a particular tenant, shared by a subset of tenants, or shared by all tenants. Refer to https://docs.openstack.org/liberty/networking-guide/intro-os-networking-overview.html for additional details.</p> 

Tenant Network Created by tenants for use by their instances and cannot be shared (based upon default policy settings). Refer to <https://docs.openstack.org/liberty/networking-guide/intro-os-networking-overview.html> for additional details.



Kubernetes Configuration

Kubernetes Worker Count	This field is auto-populated with the recommended number of worker VMs. While you can change the recommended number, be sure to verify that the worker count is adequate to accommodate the modules that you want to install. See Prepare Infrastructure for additional details.
Kubernetes Pod CIDR	Floating IP pool from which IP addresses are assigned to pods. <div style="border: 1px solid green; padding: 5px; margin-top: 10px;"> ✔ Verify that this IP does not conflict with the node/VM IP address. </div>

Proxy Configuration

HTTP Proxy	The hostname or IP address of the proxy host along with the port.
HTTPS Proxy	The hostname or IP address of the secure proxy host along with the port.

NTP Configuration

NTP Servers	A comma-separated list of IP addresses or FQDNs of your NTP server(s) to be used to sync VM clocks.
NTP Pools	A comma-separated list of IP addresses or FQDNs of your NTP cluster(s) to be used to sync VM clocks.

8. Click **Install**. The installation progress is visible on screen.
9. Once successful, you see the following message.

CloudCenter Suite installation successful!

10. You have the following options at this point:

- a. Click **Take Me To Suite Admin** to launch and set up the **Suite Admin**.
- b. Click **Install Another Cluster** to start another installation and go back to the homepage (Installer Dashboard).
- c. Download **Kubeconfig file** to connect to the launched cluster using the **kubect** tool.
- d. After the installation is complete, use the following command to SSH into the workers/primary servers as **ubuntu** and use the private SSH key of the public key (provided when you configured the Placement Properties details above).

✔ Ensure that Port 22 is open on the primary server/worker node so you can provide communication security via Security Groups/Firewall rules for OpenStack environments.

```
#Sample command to SSH into a worker/primary server  
ssh -i <private key> ubuntu@<primary server/Worker IP>
```

11. Be sure to switch off the installer VM. You can reuse this VM for any other purpose, for example, as an Offline Repository or to upgrade the Kubernetes cluster or to upgrade the tenant image on the nodes.

Existing Cluster Installation

Install the CloudCenter Suite on an Existing Kubernetes Cluster

- [Overview](#)
- [Restrictions](#)
- [Prerequisites](#)
- [Procedure](#)

Overview

Once you access the Suite Installer Dashboard (see [Prepare Infrastructure](#)), you can choose to install the Suite Admin on an existing cluster.

Restrictions

Before proceeding with section, adhere to the following restrictions:

- **AWS:** The CloudCenter Suite does not currently support a Suite Admin installation on an existing AWS cluster.
- **Permission:** Admin-level permissions for the cluster are mandatory for a user to install the Suite Admin on an existing cluster.

Prerequisites

Verify that the cluster adheres to the following requirements:

- **Kubernetes Version:** The existing Kubernetes cluster must be of Version v1.18.x or and later.
- **Kubernetes Add Ons:** Install Cert-manager version v1.0.2 (required) using the following command:

```
kubectl apply -f https://github.com/jetstack/cert-manager/releases/download/v1.0.2/cert-manager.yaml
```

Also, see <https://cert-manager.io/docs/>

- **Instance Type:** The instance type for GKE should be n1-standard-8 or higher. Verify that it is large enough to accommodate the installation of Suite Admin and other CloudCenter Suite modules.
- **Basic Authentication:** When creating the GKE cluster, go to **Security** and check the box to **Enable Basic Authentication**.
- **Storage Class:** The default storageClass must be configured.
- **Kubeconfig:** The kubeconfig user must have cluster-admin permission in the kubeconfig namespace.
 - If the cluster does not support Load Balancer.
 - GCP: You must remove auth provider and use the admin user password.
- **RBAC** - Must be enabled.
- **Pod Priority:** Define the PriorityClass for suite-high/suite-medium/suite-low.
 - Refer to <https://kubernetes.io/docs/concepts/configuration/pod-priority-preemption/> for details.
 - The commands to define PriorityClass are listed in the following code block.

```

# create pod priority class: suite-high/suite-medium/suite-low

##### begin create pod priority

cat <<EOF | kubectl apply -f -

apiVersion: scheduling.k8s.io/v1beta1

kind: PriorityClass

metadata:

  name: suite-high

value: 1000000

globalDefault: false

description: "High priority"

---

apiVersion: scheduling.k8s.io/v1beta1

kind: PriorityClass

metadata:

  name: suite-medium

value: 10000

globalDefault: false

description: "Medium priority"

---

apiVersion: scheduling.k8s.io/v1beta1

kind: PriorityClass

metadata:

  name: suite-low

value: 100

globalDefault: false

description: "Low priority"

EOF

##### end create pod priority

```

- GKE clusters with static version or regional location type are supported.
- Use the 'Container-Optimized OS with Docker (cos) (default)' as the container run time for the cluster.
- **Azure AKS clusters in private networks with and without *advanced network configurations*:**
 - For clusters with advanced network configurations with no private network follow the previously mentioned **existing cluster** installation scenario.
 - For clusters in private network with or without advanced network configurations enabled as the Kubernetes API server endpoint is in a private network with no public IP address, but there are multiple ways a network connection between the AKS cluster and installer VM(and therefore successful installation) can be established you will need to use an installer VM that has access to the AKS cluster's Azure Virtual Network (VNet). Use one of the following options:
 - (Easiest) Start the installer VM in the same private network (vnet) where the Kubernetes cluster is so that installer can connect to the cluster.

- Use the Installer VM in a separate network and set up Virtual network peering.
- Use an Express Route or VPN connection.
- Refer to <https://docs.microsoft.com/en-us/azure/aks/private-clusters#options-for-connecting-to-the-private-cluster> for additional details.

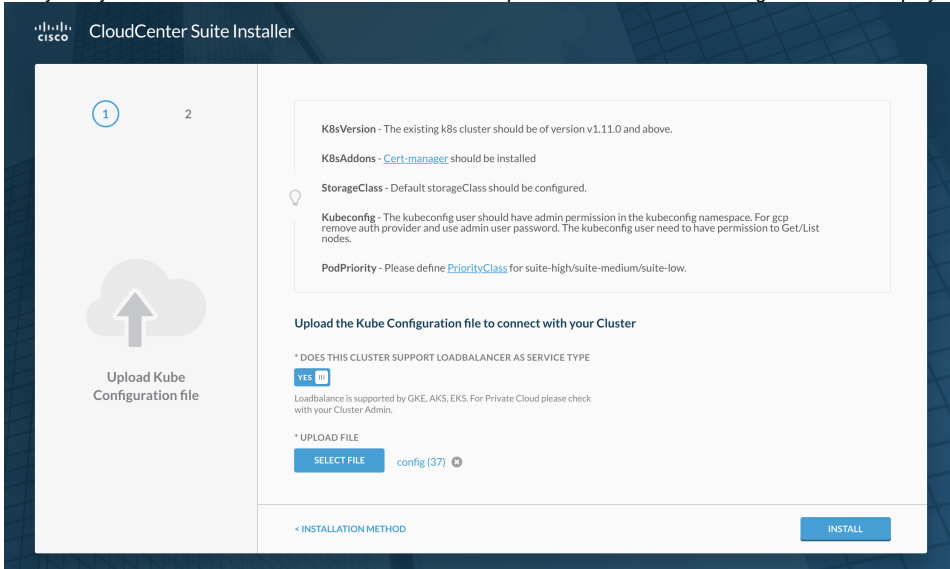
Procedure

To install the CloudCenter Suite on an existing cluster, perform the following procedure.

1. Navigate to the Suite Installer Dashboard.
2. Click **Existing Cluster** to get started as displayed in the following screenshot.



3. Verify that you have met the items identified in the *Prerequisites* section. The following screenshot displays these items as well.



4. Identify if your cluster supports **load balancer as the service type** accordingly, turn this toggle
 - a. **YES** Toggle ON if supported (public clouds generally support load balancers)
 - b. **NO** Toggle OFF if not supported (private clouds generally do not support load balancers)
5. Upload the Kubeconfig file.

Click **Install**. The installation progress is visible on screen. Once successful, you see the following message .

```
CloudCenter Suite installation successful!
```

6. You have the following options at this point:
 - a. Click **Take Me To Suite Admin** to launch and set up the [Suite Admin](#).
 - b. Click **Install Another Cluster** to start another installation on the same cluster.

You have now installed the Suite Admin on an existing cluster.

Upgrade Kubernetes Cluster

Upgrade Kubernetes Cluster

Access the Suite Installer Dashboard (see [Prepare Infrastructure](#)) to install a new cluster and launch nodes for the new Kubernetes cluster

- [Upgrade Approach](#)
- [OpenStack Upgrade](#)
- [VMware vSphere Upgrade](#)

Upgrade Approach

Upgrade Approach

- [Upgrade Approach](#)
 - [Overview](#)
 - [Restrictions](#)
 - [Prerequisites](#)
- [Kubernetes Cluster Upgrade](#)
 - [Process](#)
 - [Upgrading from SA 5.2.3 to 5.2.4](#)

Overview

This section provides details on restrictions, prerequisites, and the process to upgrade the Kubernetes cluster. During this upgrade, the software upgrades the cluster and migrates the pods to new worker instances.



If you restart any worker node, be sure to wait for approximately 10 minutes before logging into the CloudCenter Suite this timeline is determined by the pods taking about 10 minutes to startup.



For private cloud, the CloudCenter Suite 5.2 Installer cannot perform a Kubernetes upgrade on a cluster that was installed using a previous release (for example, any 5.x release). Instead, you should use the backup and restore functionality and restore it on a freshly created CloudCenter Suite 5.2.0 cluster and then perform the Kubernetes version upgrade.

This is the same process that was followed for any other CloudCenter Suite 5.0 or 5.1 releases.

Restrictions

Before proceeding with an upgrade, adhere to the following restrictions:

- **Usage:** To upgrade the Kubernetes cluster to a new version, you can do so from CloudCenter Suite 5.1.0 and later releases.
 - You cannot use the CloudCenter Suite 5.2 upgrader to upgrade a CloudCenter Suite 5.1 or 5.0 cluster. You can only use the CloudCenter Suite 5.2 upgrader effective CloudCenter Suite 5.2.1 to upgrade to a later release.
 - As an upgrader is not available to upgrade from CloudCenter Suite 5.2 to CloudCenter Suite 5.2, you must use the [Backup and Restore](#) procedure to upgrade to a CloudCenter Suite 5.2 cluster.
 - Even if you update the Suite Admin to Suite Admin 5.2, the underlying cluster will not have the capability to be upgraded as it is still using CloudCenter Suite 5.2.
 - **Public Clouds:**
 - Take a backup and then restore on to a new existing cluster with supported kubernetes version for the cloud.
 - **Private Clouds:**
 - By upgrading the cluster, you are performing a rolling upgrade on each base image in the cluster.
 - A rolling upgrade may or may not include a change in the Kubernetes version it may merely apply an OS patch or address vulnerabilities depending on the image version that you use.
 - The installer includes a default Kubernetes cluster image (called, *CCS-version-Base-Image*). The VM Template contains a list of tenant images with a *CCS-version-Base-Image* name format. If you want to upgrade to a version other than the default version provided by the installer, then upload that *CCS-version-Base-Image* under the root folder, so that it will display in this dropdown list. You can use this option to upgrade the cluster across private clouds.
- **Suite Admin-level Permissions:** Suite Admin-level permissions are mandatory for a user to upgrade the cluster.
- **New Clusters Only:** You can upgrade a cluster that is created (from the Suite Installer) using the **New cluster** option.



If you created your cluster by clicking the **Existing cluster** option (using the KubeConfig file), then you cannot upgrade this cluster using the process provided in this section.

Prerequisites

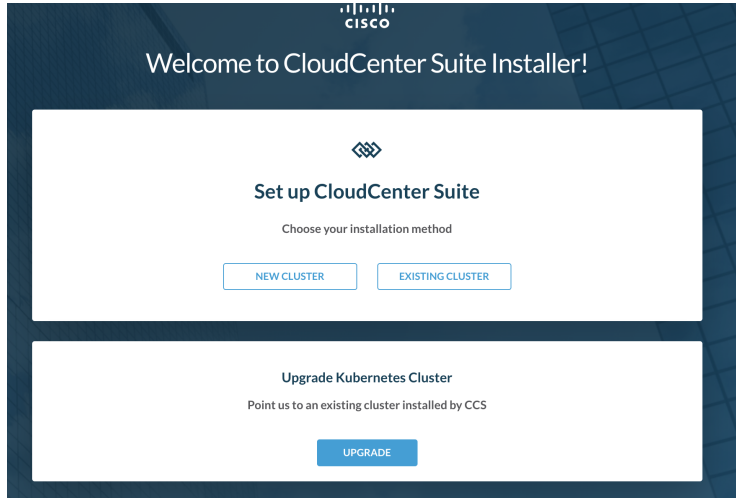
Verify that the cluster adheres to the following requirements:

- **Backup Environment:** Back up your environment before initiating the upgrade. See [Backup Approach](#) for additional details.
- **Schedule Downtime:** Schedule a suitable downtime during off-peak hours to minimize the impact to your users and or customers. Communicate the downtime as the CloudCenter Suite will not be accessible during the upgrade.
- **Verify Kubernetes Version:** Verify that the existing Kubernetes cluster is Version v1.16.3 and above.

Kubernetes Cluster Upgrade

Customers running CCS v5.2.4 can now upgrade their kubernetes cluster using v5.3.0 installer. However they need to follow a few required steps to start upgrading the kubernetes version to 1.18.12.

1. **IMPORTANT:** Follow documented steps to take a backup of the cluster before proceeding with the kubernetes upgrade.
2. Upgrade Suite Admin version to v5.3.0 using the UI. This is a required step to continue the upgrade, if Suite Admin is not upgraded users will see this error message on the installer: "
Cluster is NOT running Common-Framework Suite Admin v5.3.x. Please upgrade the Suite Admin chart to v5.3.x before upgrading kubernetes."
3. **(Only on vSphere)** This release addresses the security issue of encrypting ETCD secrets in Kubernetes. To do this users need to run the documented workaround script to enable ETCD Encryption of secrets in the kubernetes clusters. New cluster installed using v5.3.0 installer has ETCD encryption enabled by default. The contents of the `etcdEncrypt_master_1.sh` script is shown below. **The script is mandatory, not optional for the upgrade to work.**
IMPORTANT: DO NOT run the workaround script TWICE on any of the masters, this can bring the cluster since the script interacts with kubernetes apiserver of the cluster.
4. Use v5.3.0 installer to upgrade the kubernetes cluster & cert-manager
5. After upgrading the cluster, users can use the "Take me to Suiteadmin" link to go back to the Suite Admin UI. Please note that this address might change for DHCP based clusters.
6. In the upgraded cluster, if you want to upgrade the Workload-manager, Cost-optimizer or Action Orchestrator chart and see this error on suite admin UI:
"failed to wait for command /bin/helm, err: exit status 1, msg: Error: [unable to recognize "" : no matches for kind "Certificate" in version "certmanager.k8s.io/v1alpha1", unable to recognize "" : no matches for kind "Issuer" in version "certmanager.k8s.io/v1alpha1"]"
7. Please follow workaround steps mentioned in [Required Post-Kubernetes Upgrade Configuration Tasks](#)



`etcdEncrypt_master_1.sh` from step 3:

Important: File name of the script in the first master must be `etcdEncrypt_master_1.sh`

```
#!/bin/bash

# if encryption is not enabled, enable it.
# Because etcd disk is shared between old and new master node,
# we have to turn on encryption on the old master, and encrypt
# all secrets first. After this, when new master node
# is booted, its own etcd will have encrypted data and kube api
# process will then use encryption key to decrypt secrets stored
# in etcd.
function check_if_running_sudo {
  if [ "$EUID" -ne 0 ]
  then echo "This script needs to run as sudo. Please retry with:"
  echo "sudo bash $0"
  exit
  fi
}

function check_if_vsphere_cloud {
  if kubectl get cm k8s-mgmt.cloudaccount --kubeconfig=/home/${users}/.kube/config -n cisco ; then
    echo "continue with the upgrade..."
  else
```

```

    echo "Workaround is not required for this cloud. Please continue with the upgrade from installer UI."
    exit
fi
kubect1 get cm k8s-mgmt.cloudaccount --kubeconfig=/home/$(users)/.kube/config -n cisco -o jsonpath="{.data.
data}" | base64 -d > k8s-mgmt.cloudaccount.json
sleep 60
CLOUD_TYPE=$(cat k8s-mgmt.cloudaccount.json | jq -r '.cloudType')
if [[ $CLOUD_TYPE == *"vsphere"* ]]; then
echo "Vshere cloud, continue with the workaround."
else
echo "$CLOUD_TYPE cloud, workaround is not required for this cloud. Please continue with the upgrade from
installer UI"
exit
fi
}

function check_for_encryption_file {
#Function to check if this file is executed only Once, if first time, create a lock
FILE=/etc/kubernetes/pki/etcd/encryption.conf
if [ -f "$FILE" ]; then
echo -e "-----"
echo "$FILE already exists.
Looks like workaround is already performed on this master node.
Please do not run the script multiple times on same master. Continue with other master nodes."
echo -e "-----"
exit
fi
}

function backup_resources {
mkdir -p backup && cd backup || exit
echo -e "Backing up all the certificates to backup folder $PWD /backup"
kubect1 get -o yaml --all-namespaces issuer,clusterissuer,certificates,secrets --kubeconfig=/home/$(users)/.
kube/config > cert-manager-backup.yaml
kubect1 get cm k8s-mgmt.cloudaccount --kubeconfig=/home/$(users)/.kube/config -n cisco -o jsonpath="{.data.
data}" | base64 -d > k8s-mgmt.cloudaccount.json
echo -e "Backing up all the secrets to resources_backup.yaml"
kubect1 get -o yaml --all-namespaces secrets --kubeconfig=/home/$(users)/.kube/config > resources_backup.yaml
for n in $(kubect1 get -n cisco -o=name pvc,certificate,configmap,serviceaccount,secret,ingress,service,
deployment,statefulset,hpa,job,cronjob)
do
mkdir -p $(dirname "$n")
kubect1 get -n cisco --kubeconfig=/home/$(users)/.kube/config -o=yaml --export "$n" > "$n".yaml 2> /dev
/null
done
cd ..
echo -e "-----"
}

ETCD_ENCRYPTION_SECRET=$(head -c 32 /dev/urandom | base64)

function generate_encryption_conf_file {
#!/bin/bash

OUTPUT_FILE=/etc/kubernetes/pki/etcd/encryption.conf

echo -e "\nCreating ETCD_ENCRYPTION_SECRET and ETCD_ENCRYPTION_KEY"

ETCD_ENCRYPTION_SECRET=$(head -c 32 /dev/urandom | base64)
ETCD_ENCRYPTION_KEY=$(echo ccp-key-$(echo "$ETCD_ENCRYPTION_SECRET" | base64 -d | sha256sum | cut -c1-10))

echo "Writing content to file $OUTPUT_FILE"

cat <<EOF > $OUTPUT_FILE
apiVersion: apiserver.config.k8s.io/v1
kind: EncryptionConfiguration
resources:
- resources:
- secrets
providers:
- aescbc:

```

```

        keys:
        - name: $ETCD_ENCRYPTION_KEY
          secret: $ETCD_ENCRYPTION_SECRET
    - identity: {}
EOF
}

check_if_running_sudo
check_if_vsphere_cloud
check_for_encryption_file
backup_resources
echo -e "\nCreating etcd Encryption configuration file for all the master nodes"
generate_encryption_conf_file
echo -e "\nSaved etcd Encryption configuration file at $OUTPUT_FILE"
cp $OUTPUT_FILE encryption.conf
echo -e "\nSaved etcd Encryption configuration file's copy at encription.conf\n\n"

#-----
#-----

function check_if_running_sudo {
if [ "$EUID" -ne 0 ]
then echo "This script needs to run as sudo. Please retry with:"
echo "sudo bash $0"
exit
fi
}
function sleep_function {
echo "Sleeping for $1 minutes..."
while true;do echo -n .;sleep 1;done &
sleep "$1" # or do something else here
kill $!; trap 'kill $!' SIGTERM
echo "Done"
echo -e "\n\n-----"
}
function check_if_vsphere_cloud {
if kubectl get cm k8s-mgmt.cloudaccount --kubeconfig=/home/${users}/.kube/config -n cisco ; then
echo "continue with the upgrade..."
else
echo "Workaround is not required for this cloud. Please continue with the upgrade from installer UI."
exit
fi
kubectl get cm k8s-mgmt.cloudaccount --kubeconfig=/home/${users}/.kube/config -n cisco -o jsonpath="{.data.data}" | base64 -d > k8s-mgmt.cloudaccount.json
sleep 60
CLOUD_TYPE=$(cat k8s-mgmt.cloudaccount.json | jq -r '.cloudType')
if [[ $CLOUD_TYPE == *"vsphere"* ]]; then
echo "Vshere cloud, continue with the workaround."
else
echo "$CLOUD_TYPE cloud, workaround is not required for this cloud. Please continue with the upgrade from installer UI"
exit
fi
}
function check_for_encryption_file {
#Function to check if this file is executed only Once, if first time, create a lock
FILE=/etc/kubernetes/pki/etcd/encryption.conf
if [ -f "$FILE" ]; then
echo -e "-----"
echo "$FILE already exists.
Looks like workaround is already performed on this master node.
Please do not run the script multiple times on same master. Continue with other master nodes."
echo -e "-----"
exit
fi
}
check_if_running_sudo
check_if_vsphere_cloud

```

```

OUTPUT_FILE=/etc/kubernetes/pki/etcd/encryption.conf

function copy_encryption_conf_file {
    #!/bin/bash

    OUTPUT_FILE=/etc/kubernetes/pki/etcd/encryption.conf
    echo "Copying script to $OUTPUT_FILE"
    cat <<EOF > $OUTPUT_FILE
$(sed -n -e '/^ETCD_ENCRYPTION_FILE_START$/,/^ETCD_ENCRYPTION_FILE_END$/ { /^ETCD_ENCRYPTION_FILE_START$/g;
/^ETCD_ENCRYPTION_FILE_END$/d; p; }' "$0")
EOF
    cp $OUTPUT_FILE encryption.conf
    echo -e "\nSaved etcd Encryption configuration file's copy at encription.conf\n\n"
}

# File content which updates kube-apiserver.yaml and waits for kubserver to start
function check_kube_api {
    API_PID=$(ps -ef | awk '/kube-ap[i]server/{print $2}')
    echo -e "-----"
    echo "check_kube_api to verify that kube-apiserver is restarted: $API_PID ..."
    while [ $API_PID = $1 ]; do
        echo "wait for kube api server to exist.."
        sleep_function 2m
        API_PID=$(ps -ef | awk '/kube-ap[i]server/{print $2}')
    done

    while ! ps -ef | grep kube-ap[i]server; do
        echo "wait for kube api server to restart.."
        sleep_function 2m
    done

    echo -e "\n=====
echo -e "Waiting. Please DO NOT INTERRUPT."
sleep 300
echo -e "Successfully updated kubernetes manifests"
echo -e "\n=====
}

function update_kube_apiserver_manifest {
    API_PID=$(ps -ef | awk '/kube-ap[i]server/{print $2}')

    echo -e "Adding ETCD encryption resources to kube-apiserver, it might take some time. Please DO NOT
INTERRUPT."

    API_CONF="/etc/kubernetes/manifests/kube-apiserver.yaml"
    ENC_CONF="/etc/kubernetes/pki/etcd/encryption.conf"

    sudo python - <<EOF
import yaml
name = 'k8s-encryption'

with open("$API_CONF", "r") as f:
    data = yaml.safe_load(f)
    data['spec']['volumes'].append({'hostPath': {'path': "$ENC_CONF", 'type': ''}, 'name': name})
    data['spec']['containers'][0]['command'].append("--encryption-provider-config=$ENC_CONF")
    data['spec']['containers'][0]['volumeMounts'].append({'mountPath': "$ENC_CONF", 'name': name})

with open("$API_CONF", "w") as f:
    f.write(yaml.dump(data, default_flow_style=False, indent=2))
EOF

    sleep_function 2m
}

function patch_kubeconfig_secret {

    KUBECONFIG_SECRET_NAME=$(kubectl get secret -n ccp --kubeconfig=/home/${users}/.kube/config | grep kubeconfig

```

```

| awk '/-/{print $1}')
echo "\nCreating patch file for kubeconfig secret $KUBECONFIG_SECRET_NAME"
echo "{\"data\":{\"etcdEncryptionKey\": \"${ETCD_ENCRYPTION_SECRET}\"}}\" > patch_kubeconfig_secret.json
kubectl -n ccp patch secret "$KUBECONFIG_SECRET_NAME" --kubeconfig=/home/$(users)/.kube/config --patch "$(cat
patch_kubeconfig_secret.json)"
echo -e "\nSuccessfully patched kubeconfig secret, sleeping for 2 mins, please DO NOT INTERRUPT.\n\n"
sleep_function 2m
kubectl get secrets --all-namespaces -o json --kubeconfig=/home/$(users)/.kube/config | kubectl replace -f -
--kubeconfig=/home/$(users)/.kube/config
echo -e "\nUpdating the secret, please DO NOT INTERRUPT."
sleep_function 5m

echo -e "\n\n=====
echo -e "=====
echo -e "\nWorkaround completed."
echo -e "\nRestart all 3 Master nodes one by one."
echo -e "\nGive significant time for all services to become RUNNING before proceeding with next master node.
(Recommended ~5 mins)"
echo -e "\n\n=====
echo -e "=====
}

#-----
#-----

API_PID=$(ps -ef | awk '/kube-ap[i]server/{print $2}')
echo -e "\n\nPrinting encryption file, please copy it to other masters and execute\n"
echo ": ' " > etcdEncrypt_master_2.sh
echo ": ' " > etcdEncrypt_master_3.sh
echo "ETCD_ENCRYPTION_FILE_START" >> etcdEncrypt_master_2.sh
echo "ETCD_ENCRYPTION_FILE_START" >> etcdEncrypt_master_3.sh
cat $OUTPUT_FILE >> etcdEncrypt_master_2.sh
cat $OUTPUT_FILE >> etcdEncrypt_master_3.sh
echo "ETCD_ENCRYPTION_FILE_END" >> etcdEncrypt_master_2.sh
echo "ETCD_ENCRYPTION_FILE_END" >> etcdEncrypt_master_3.sh
echo " ' " >> etcdEncrypt_master_2.sh
echo " ' " >> etcdEncrypt_master_3.sh

sed -n '107,245p' etcdEncrypt_master_1.sh >> etcdEncrypt_master_2.sh
sed -n '107,245p' etcdEncrypt_master_1.sh >> etcdEncrypt_master_3.sh
echo "check_for_encryption_file" >> etcdEncrypt_master_2.sh
echo "check_for_encryption_file" >> etcdEncrypt_master_3.sh
echo "copy_encryption_conf_file" >> etcdEncrypt_master_2.sh
echo "copy_encryption_conf_file" >> etcdEncrypt_master_3.sh
echo "update_kube_apiserver_manifest" >> etcdEncrypt_master_2.sh
echo "update_kube_apiserver_manifest" >> etcdEncrypt_master_3.sh
echo -e "export -f check_kube_api \ntimeout 300s bash -c check_kube_api "\$API_PID" " >> etcdEncrypt_master_2.sh
echo -e "export -f check_kube_api \ntimeout 300s bash -c check_kube_api "\$API_PID" " >> etcdEncrypt_master_3.sh
echo -e "ETCD_ENCRYPTION_SECRET="\$ETCD_ENCRYPTION_SECRET">>etcdEncrypt_master_3.sh
echo "patch_kubeconfig_secret" >> etcdEncrypt_master_3.sh

update_kube_apiserver_manifest
export -f check_kube_api
timeout 300s bash -c check_kube_api "$API_PID"

echo -e "\nWorkaround completed on Master-1, please continue with Master-2 and Master-3"
echo -e "\n\n=====
echo -e "\nCopy etcdEncrypt_master_2.sh to second master and run sudo bash etcdEncrypt_master_2.sh"
echo -e "\nCopy etcdEncrypt_master_3.sh to third master and run sudo bash etcdEncrypt_master_3.sh"
echo -e "=====

```

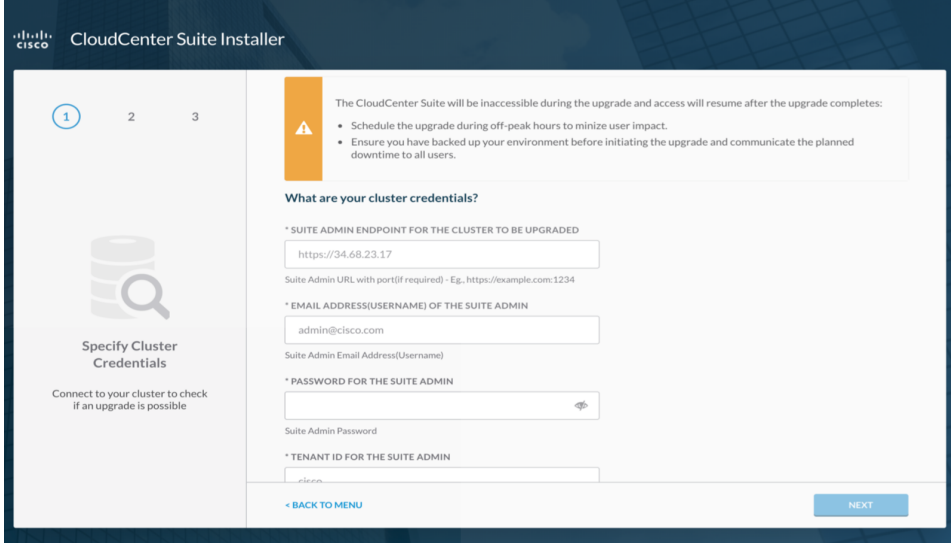
Process

This is the generic process to upgrade a Kubernetes cluster for a cloud that is supported by the CloudCenter Suite.


1. Navigate to the Suite Installer Dashboard (see [Prepare Infrastructure](#)).



2. Click **Upgrade** in the *Upgrade Kubernetes Cluster* section to specify the credentials for your cluster as displayed in the following screenshot.



3. Enter the Suite AdminURL (or DNS), username, password, and Tenant ID for the admin account.

 See the individual cloud upgrade pages for additional notes and nuances.

CloudCenter Suite Installer

1 2 3

Specify Cluster Credentials

Connect to your cluster to check if an upgrade is possible

* SUITE ADMIN ENDPOINT FOR THE CLUSTER TO BE UPGRADED

Suite Admin URL with port(if required) - Eg. https://example.com:1234

* EMAIL ADDRESS(USERNAME) OF THE SUITE ADMIN

Suite Admin Email Address(Username)

* PASSWORD FOR THE SUITE ADMIN

Suite Admin Password

* TENANT ID FOR THE SUITE ADMIN

Suite Admin Tenant ID

* IS THIS AN AMAZON EKS CLUSTER?
 NO
Toggle to provide credentials of Amazon Elastic Kubernetes Service cluster

[← BACK TO MENU](#)

4. Click **Connect** to validate your credentials.
5. At this point, you have multiple scenarios:

- You will be able to click **Next** and select the desired Kubernetes version from the dropdown list for this upgrade. Proceed to Step 8.
- If an upgrade is not available for your cluster as displayed in the following screenshot, some possible reasons are:
 - An upgrade is not currently available as the cluster is already at the latest available version of Kubernetes.

CloudCenter Suite Installer

1 2 3

Specify Cluster Credentials

Connect to your cluster to check if an upgrade is possible

* SUITE ADMIN ENDPOINT FOR THE CLUSTER TO BE UPGRADED

Suite Admin URL with port(if required) - Eg. https://example.com:1234

* EMAIL ADDRESS(USERNAME) OF THE SUITE ADMIN

Suite Admin Email Address(Username)

* PASSWORD FOR THE SUITE ADMIN

Suite Admin Password

* TENANT ID FOR THE SUITE ADMIN

Suite Admin Tenant ID

* IS THIS AN AMAZON EKS CLUSTER?
 NO
Toggle to provide credentials of Amazon Elastic Kubernetes Service cluster

✔ Connected

[← BACK TO MENU](#) NO UPGRADE AVAILABLE FOR YOUR CLUSTER...

No upgrade available for your cluster. You can upgrade another cluster by changing the URL.

- You may have provided the wrong cluster credentials (in this case, you will not see the *Connected* status update when you try to connect). If so, enter the right credentials and try again.

- Once Connected, you see the cloud type and other information on the left side of the screen as visible in the following screenshot (sample of a GKE environment):
- If an upgrade is available, select the **Desired K8s version** for the upgrade.
- Click **Upgrade** to upgrade the Kubernetes cluster as well as the master and worker nodes once the upgrade is complete. A progress bar with relevant status messages is displayed.



An upgrade operation can take more than one hour depending on the number of nodes to be upgraded and cloud response time.

- At this point, you can:
 - Download the latest logs to track the upgrade process.
 - Wait for cluster to finish upgrading.
- The installation progress and success is visible on the screen.



See the individual cloud upgrade pages for which of these options are available and for additional notes and nuances.

- You have the following options at this point depending on your cloud environment:
 - Click **Take Me To Suite Admin** to launch and set up the [Suite Admin](#).
 - Click **Install Another Cluster** to start another installation on the same cluster.
 - Download the **Kubeconfig file**.
 - Download the **SSH private key**.
 - Re-purpose the installer server.
- Login to CloudCenter Suite using valid credentials and verify that your information is preserved and that the cluster was upgraded.

Upgrading from SA 5.2.3 to 5.2.4

If you are upgrading from 5.2.3 or older to 5.2.4, perform the following procedure:

- Run an SQL command in the suite-postgresql database.
- Log in to common-framework-suite-postgresql-0

```
kubectlexec -it common-framework-suite-postgresql-0 -n cisco bash
```

- Run the following bash command.

```
PGDATABASE=suite-samlpsql -c "UPDATE public.saml_infra_config SET cert_source='cert_manager' WHERE id=1;"
```

- Optionally regenerate the SSO certificate.

```
kubectldelete secret suite-saml-sso-tls-n cisco
```

5. Restart the suite-samlpod.

```
kubectldelete pod suite-saml-pod -n cisco
```

6. If you use SSO, reconfigure SSO.

OpenStack Upgrade

OpenStackUpgrade

- [Overview](#)
- [OpenStack Nuances](#)
- [Module Details](#)
- [Upgrade Process](#)

Overview

See [Upgrade Approach](#) for details on permissions and prerequisites.

OpenStack Nuances

Verify the following OpenStack nuances:

- OpenStack newton release with at least the following service versions:
 - Cinder v2
 - Keystone v3
 - OpenStack Nova v2
 - OpenStack Networking v2
 - OpenStack Glance v2
- Ensure to add Port 6443 to the default security group as the security group created for the cluster is not automatically assigned to the load balancer created for the cluster.
- The tenant and project requirements for OpenStack Cloud are identified in the following table.

Model	Quota	Description
For all cases	2 (primary server group, worker group)	Server Groups
	Number of workers + number of primary servers	Server Group Members
	3 (API load balancers)	Load Balancers
	6 (2 for each load balancer)	Health Monitors
	6 (2 for each load balancer)	Pools
	6(2 for each load balancer)	Listeners
	3 (1 for the cluster VMs, 2 for the Kubernetes load balancer services)	Security Groups
	18	Security Group Rules
	See Prepare Infrastructure for additional details	Volume GB
	Number of workers + number of primary servers +3 for each load balancer	Ports
	Number of workers + number of primary servers	Instances
	16 GB (recommended for each worker and each primary server)	RAM
32 (recommended for each workers and each primary server)	vCPUs	
Tenant network	Floating IPs = 3	1 for each load balancer
	Networks = 1	For the tenant network
	Subnet = 1	For the tenant network
	Router = 1	For the tenant network to public network connection
Provider network	Number of workers + number of primary servers + 3 load balancers	Free IPs in the provider network

- **Network Time Protocol (NTP) must be configured this is important as the CloudCenter Suite installation can fail, if NTP is not configured or if it is wrongly configured.**



If you setup CloudCenter Suite in offline mode, you must provide valid NTP server details before you save your configuration.

Module Details

Additionally, refer to your module documentation for module-specific dependencies as identified in the following table:

Module	Documentation
Workload Manager	Cloud Overview
Action Orchestrator	Add Cloud Account
Cost Optimizer	Cloud Overview

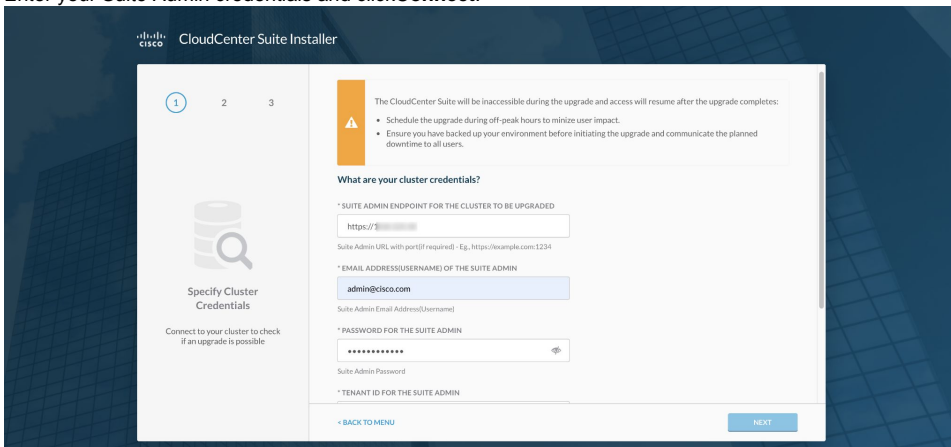
Upgrade Process

To upgrade the cluster for an OpenStack Kubernetes environment, perform the following procedure.

1. Verify that you have prepared your environment as listed in the *OpenStack Nuances* section above.
2. Navigate to the Suite Installer Dashboard.



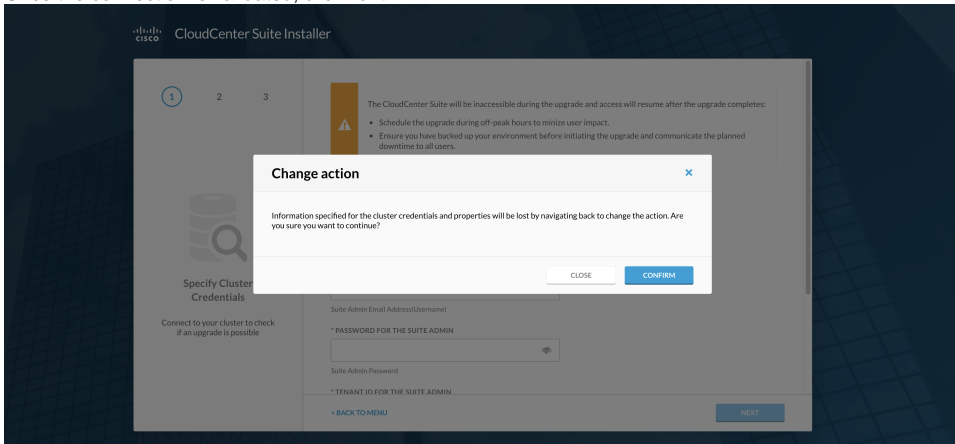
3. Enter your Suite Admin credentials and click **Connect**.




OpenStack Details	Description
Suite Admin Endpoint for the Cluster to be Upgraded	The DNS address or IP address of the vCenter server where you launch the Suite Admin.
Email Address (Username) of the Suite Admin	The email address of Suite Admin(the Initial Administrator) who setup the Suite Admin.
Password for the Suite Admin	The password for the Suite Admin(the Initial Administrator) who setup the Suite Admin.
Tenant ID for the Suite Admin	The Tenant ID for the Suite Admin(the Initial Administrator) who setup the Suite Admin.
Is This an Amazon EKS Cluster	Toggle the switch (default = No). If it is, provide the Access Key and Secret Key details.

The CloudCenter Suite validates the OpenStack credentials to ensure that the cluster is available to this user.

- Once the connection is validated, click **Next**.




- Click **Confirm** to continue with the changes.




6. When Connected, you see the cloud type and other information on the left side of the screen. Enter the information in the Upgrade settings fields.

Cisco CloudCenter Suite Installer

1 2 3



Cluster upgrade settings



OpenStack

Specify the settings required to upgrade your cluster

Upgrade settings

CURRENT KUBERNETES VERSION

1.13.5

Current version of kubernetes installed on the existing cluster

NEW KUBERNETES VERSION

1.13.5

New version of kubernetes available as part of upgrade



NEW CISCO KUBERNETES 1.13.5 OPENSTACK IMAGE

Optional, see description.

Cisco Kubernetes 1.13.5 image to use for the masters and workers. If not provided will upload an image included with the installer.

[< CHANGE CLUSTER CREDENTIALS](#) [UPGRADE](#)


7. Click **Upgrade** to upgrade the Kubernetes cluster as well as the primary server and worker nodes once the upgrade is complete. A progress bar with relevant status messages is displayed as visible in the following screenshots.



8. At this point, you can:

- a. Download the latest logs to track the upgrade process.
- b. Wait for cluster to finish upgrading.

9. Once successful, you see the success message.



You have the following options at this point:

- a. Click **Take Me To Suite Admin** to launch and set up the **Suite Admin**.
 - b. Click **Install Another Cluster** to start another installation and go back to the homepage (Installer Dashboard).
 - c. Download **Kubeconfig file** to connect to the launched cluster using the **kubect** tool.
10. After the installation is complete, use the following command to SSH into the workers/primary servers as **cloud-user** and use the private SSH key or the public key (provided when you configured the Placement Properties details above).

```
#Sample command to SSH into a worker/primary server  
ssh -I <private key> cloud-user@<primary server/worker IP>
```

11. Login to CloudCenter Suite using valid credentials and verify that your information is preserved and that the cluster was upgraded.

You have now upgraded the cluster on the OpenStack cloud. Verify your Suite Admin and tenant data.

VMware vSphere Upgrade

VMware vSphere Upgrade

- [Overview](#)
- [Upgrade Process](#)

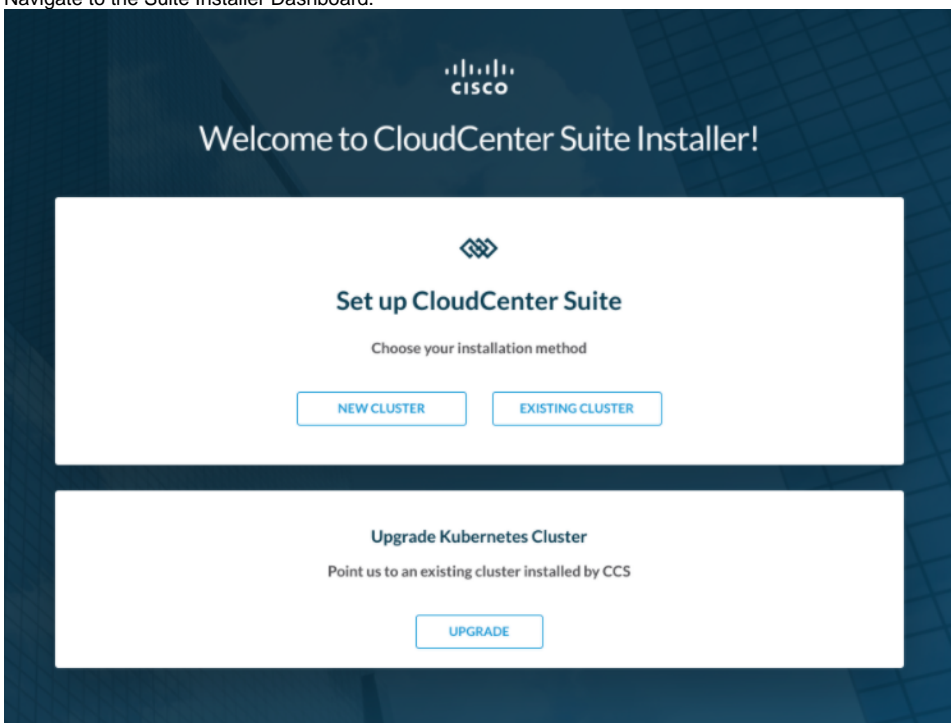
Overview

See [Upgrade Approach](#) for details on permissions and prerequisites.

Upgrade Process

To install the CloudCenter Suite on a new vSphere cluster, perform the following procedure.

1. Verify that you have prepared your environment as listed in the *VMware Nuances* section above.
2. Navigate to the Suite Installer Dashboard.



3. Enter your Suite Admin credentials and click **Connect**.

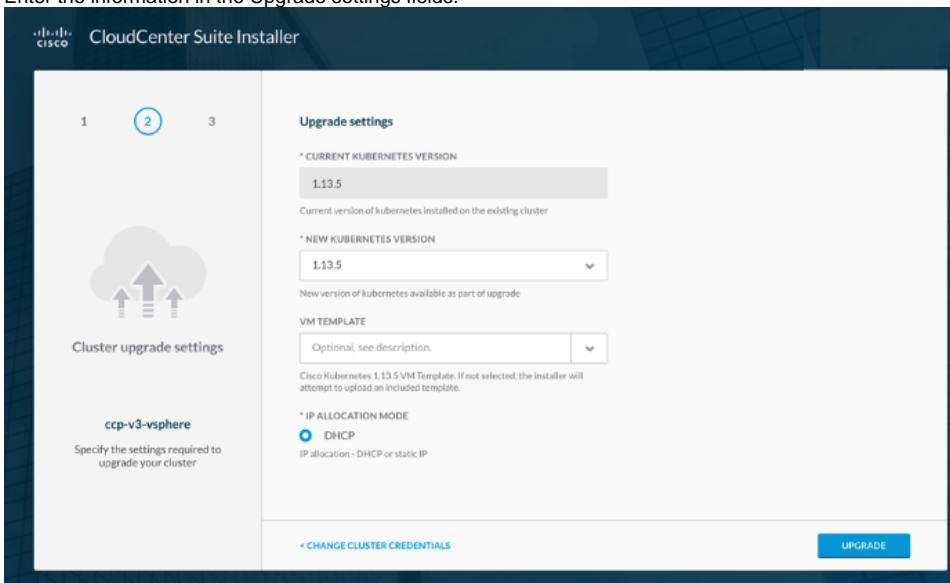
vSphere Details	Description
Suite Admin Endpoint for the Cluster to be Upgraded	The DNS address or IP address of the vCenter server where you launch the Suite Admin.
Email Address (Username) of the Suite Admin	The email address of Suite Admin (the Initial Administrator) who setup the Suite Admin.
Password for the Suite Admin	The password for the Suite Admin (the Initial Administrator) who setup the Suite Admin.
Tenant ID for the Suite Admin	The Tenant ID for the Suite Admin (the Initial Administrator) who setup the Suite Admin.
Is This an Amazon EKS Cluster	Toggle the switch (default is No). If it is, provide the Access Key and Secret Key details.

The CloudCenter Suite validates the vSphere credentials to ensure that the cluster is available to this user.

4. Once the connection is validated, click **Next**.

Once Connected, you see the cloud type and other information on the left side off the screen

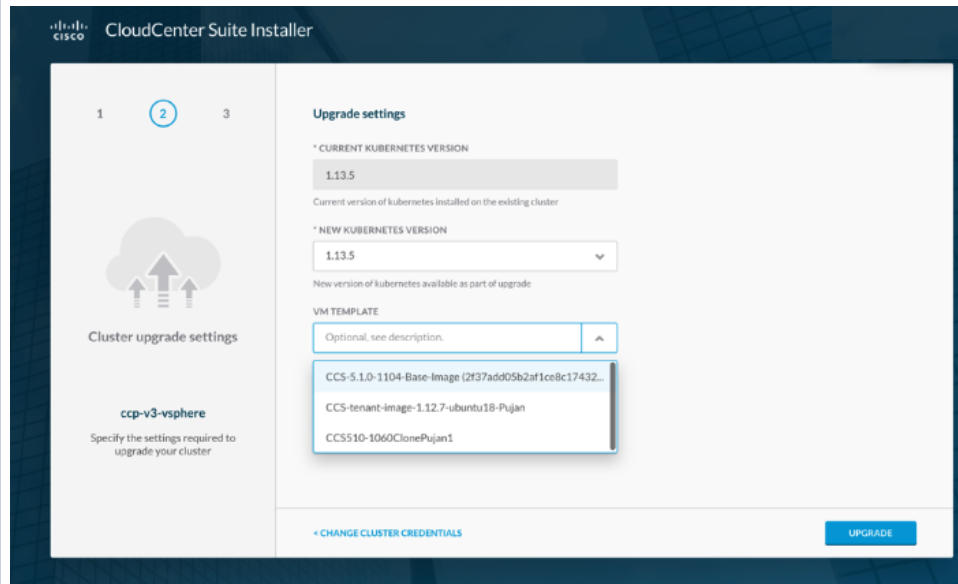
5. Enter the information in the Upgrade settings fields.



Upgrade Settings Field	Description
Current Kubernetes Version	The current version for your Kubernetes setup is pre-populated in this field.
New Kubernetes Version	If an upgrade is available, it is listed in this dropdown list. Select the Desired K8s version for the upgrade.

VM Template

Different images will be used for the installer and the cluster launched by the installer as visible in the following screenshot.



The installer includes a default Kubernetes cluster image (called, *CCS-version-Base-Image*). The VM Template contains a list of tenant images with a *CCS-version-Base-Image* name format. If you want to upgrade to a version other than the default version provided by the installer, then upload that *CCS-version-Base-Image* under the root folder, so that it will display in this dropdown list.

The *CCS-version-Base-Image* image included in the installer is selected if you do not override the setting.

To override the *CCS-version-Base-Image* image used by the Suite installer, be sure to add the applicable image in the vSphere console and selected the applicable **OVA** from the dropdown list in this field.

If you use the **OVA** installer to launch the cluster in an OpenStack environment, be sure to override this field and select the applicable **QCOW2** *CCS-version-Base-Image*.



If you install the CloudCenter Suite using any image other than *CCS-version-Base-Image*, the installation will fail.

IP Allocation Mode

This switch allows you to select the mode. Currently, only DHCP is supported.


- **DHCP:** This strategy allows the IP to be allocated by the DHCP server to the instance on server boot up.
 - **Master VIP:** The IP address for the **Take Me to Suite Admin** link. Users can determine the IP address that should have the primary server role for the **Take Me to Suite Admin** link.




This should be a unique IP and should not be assigned to any other resource.

- **Static IP:** This strategy allows the customer to provide the IP address. As this IP address may or may not be available to the server (based on the availability), you must perform adequate checks to ensure IP availability before using this strategy.
 - **Static IP Pool Start IP:** The first IP address of the static IP range. If you need to scale up nodes after setting up the Suite Admin, then you must ensure a wider range.
 - **Static IP Pool End IP:** The last IP address for the static IP range.
 - **Subnet Mask:** The netmask corresponding to the specified IP range.
 - **DNS Server List:** The comma separated list of DNS server IP addresses.
 - **Gateway List:** The comma separated list of Gateway server IP addresses.

6. Click **Upgrade** to upgrade the Kubernetes cluster as well as the primary server and worker nodes once the upgrade is complete. A progress bar with relevant status messages is displayed as visible in the following screenshot.



7. At this point, you can:
- Download the latest logs to track the upgrade process.
 - Wait for cluster to finish upgrading.
8. Once successful, you see the success message as visible in the following screenshot.



9. You have the following options at this point:
- Click **Take Me To Suite Admin** to launch and set up the **Suite Admin**.
 - Click **Install or Upgrade Another Cluster** to start another installation and go back to the homepage (Installer Dashboard).
 - Download **KubeConfig file** to connect to the launched cluster using the **kubect** tool.
 - After the installation is complete, use the following command to SSH into the workers/primary servers as **cloud-user** and use the private SSH key or the public key (provided when you configured the Placement Properties details above).

```
#Sample command to SSH into a worker/primary server
ssh -I <private key> cloud-user@<primary server/worker IP>
```

10. Be sure to switch off the installer VM. You can reuse this VM for any other purpose, for example, as an Offline Repository.

You have now upgraded the cluster on the VMware cloud. Verify your Suite Admin and tenant data.

Air Gap Installation

End-to-End Air Gap Installation in an Isolated Environment

- [Overview](#)
- [Limitations](#)
- [Prerequisites to Configure an Air Gap Setup](#)
- [Workflow](#)
- [Download the Offline Appliance and Import into vSphere](#)
- [Configure the Offline Repository](#)
- [Configure the Installer to use the Offline Repository](#)
- [Post Configuration Verification](#)
 - [For the CloudCenter Suite Installer](#)
 - [For the Launched Kubernetes Cluster](#)
 - [Using Config Map](#)
- [When Is the Offline Repository Used?](#)
- [Alternate Path of Airgap Certificates](#)

Overview

The term *air gap* refers to security measures implemented for computers, computer systems, or networks requiring airtight security without the risk of compromise. It ensures total isolation of the system from other less secure networks.

An Air Gap installation in the CloudCenter Suite context refers to the ability to support an installation of the CloudCenter Suite in environments that do not have an internet connection (equivalent of an isolated network). While **Air Gap Installation** refers to the feature, **Offline Repository** refers to the delivery mechanism for the Air Gap Installation feature.

Until 5.1, CloudCenter Suite had the concept of offline repository and installations using offline repository but that offline repository was accessed using proxy server settings that were accessed through the CloudCenter Suite cluster and is no longer available in CloudCenter Suite 5.2. To use the Air Gap solution, you must use the offline repository appliance to create a dedicated repository server which is introduced with CloudCenter Suite 5.2.0.



You cannot re-purpose or reuse the installer server.

Effective CloudCenter Suite 5.2, the CloudCenter Suite installer exchanges certificates and host information with the offline repository as soon as the installer is launched, it connects to the offline repository VM (equivalent of an isolated network).

After the cluster is launched, you can use the same offline appliance at any point and install modules. When a newer CloudCenter Suite version becomes available, the corresponding new offline appliance will also be available you can use the new appliance and upgrade to the latest version of all CloudCenter Suite modules.

Limitations

Be aware of the following limitations for the air gap feature:

- The offline repository appliance that is available in CloudCenter Suite 5.2 does not have a UI.
- This feature is only available for VMware environments.
- To upgrade from Action Orchestrator 5.1.4 to Action Orchestrator 5.2 in offline mode, follow this procedure.
 1. Verify that you have already installed [ArangoDB](#) and [NPM](#) in the device that you will be backing up.
 2. [Backup](#) up your Action Orchestrator 5.1.4 setup.
 3. [Uninstall](#) Action Orchestrator 5.1.4 from this device.
 4. [Upgrade Offline Repository](#) over to a new repository which contains Action Orchestrator 5.2
 5. [Install](#) Action Orchestrator 5.2.
 6. [Restore](#) the backed up data.

Prerequisites to Configure an Air Gap Setup

Verify these prerequisites before setting up an Air Gap installation environment:

- You must get a valid Certificate Authority to sign the certificate and a private key pair for the DNS name.
- The offline repository must be accessible from the Kubernetes cluster through the domain name.

Workflow

The following process identifies the high-level process of the change between previous releases and the new CloudCenter Suite 5.2.0 solution:

1. To deploy an offline/Air Gap appliance containing all the Docker images and Helm charts hosted in a local registry backed by a web server, the CloudCenter Suite uses harbor see <https://goharbor.io/> for additional details.
2. To configure the offline appliance and upload user defined certificates or generate self signed certificates:

- a. User defined certs, if using FQDN, your DNS should be able to resolve within the network or else the IP address of the offline repository should be part of cert as an alt alias. Also you must provide the CA cert for generated certificates. Self signed certificates can use the IP or FQDN of the offline appliance. Along with the certificates, you must also change the admin password. The out-of-box password is **Cisco123**.
3. To install the CloudCenter Suite using an offline appliance, you must turn **ON** the Air Gap setup option in the Installer page (Select **VMware**, enter your credentials of vSphere, for this option to display).

Download the Offline Appliance and Import into vSphere

1. Download the offline appliance (suite_offline OVA) from software.cisco.com.
2. Login into vSphere as an administrator or with an user with the following permissions.
 - a. The installation process requires a vSphere User with specific Permissions. For users who do not want to use the default administrator, use the following steps to create a new Role and User for the installation.
 - b. In vSphere, login into vSphere as an administrator user. Navigate to **Home > Administration > Roles** and create a Role by providing the following privileges to this role:



All listed permissions are required to proceed with this installation. Missing even one role will lead to unpredictable consequences.

- Datastore.Allocate space
 - Datastore.Browse datastore
 - Datastore.Low level file operations
 - Datastore.Remove file
 - Folder. Create folder
 - Global.Manage Custom Attributes
 - Global.Set custom attribute
 - Network.Assign network
 - Resource.Apply recommendation
 - Resource.ApplyvApp to resource pool
 - Resource.Apply virtual machine to resource pool
 - Storage views. View
 - Tasks.Create task
 - Tasks.Update task
 - Virtual machine (check all the permissions under this privilege)
 - vApp.Import
 - vApp.Power off
 - vApp.Power on
 - vApp.Suspend
 - vApp.vApp application configuration
 - vApp.vApp instance configuration
 - vApp.vAppmanagedBy configuration
 - vApp.vApp resource configurationIn
- c. Navigate to **Home > Administration > User and Groups**. Click on the + icon and create a new user. Remember the username and password - these will be used in subsequent steps.
 - d. Click on **Global Permissions**. Click on the + icon to open *Global Permission Root - Add Permission*. Click on **Add** to map the previously created user to the Role created in Step 1 - make sure to click **Propagate to children**.
3. Click on **VM and Templates**, and then select the vSphere Datacenter where the Installer needs to be uploaded. Right-click and select **Deploy OVF Template ...**
 4. In the *Deploy OVF Template* wizard, select **Local File** and open the previously downloaded OVA from your computer's file browser. Click **Nextto** proceed.
 5. For the *Select name and location* step of the wizard, select a folder directly underneath the Datacenter. **Do NOT select a sub-folder**. Click **Nextto** proceed.
 6. For the *Select resource* step of the wizard, select an ESX Host from the Cluster. Click **Nextto** proceed.
 7. For the *Select storage* step of the wizard, select an Datastore with necessary permissions as outlined above in *Prepare the vSphere Infrastructure (Prerequisites)*. Click **Nextto** proceed.



Recommendation: Select **Thin** for the *Virtual Disk Format*.

8. For the *Select clone options* of the wizard, select the checkbox for each of the following options. Click **Nextto** proceed.

- **(Optional)** *Customize the operating system*



Note: This selection is only required for environments with **OUT** access to DHCP. The requirement and "workaround" is outlined above in *Prepare the vSphere Infrastructure (Prerequisites)*.

- *Customize this virtual machine's hardware (Experimental)*
- *Power on virtual machine after creation*

Select the Customization Spec created during *Prepare the vSphere Infrastructure (Prerequisites)*. This Customization Spec was created to assign a Static IP to the CloudCenter Suite Installer. Click **Nextto** proceed.

9. For the *Customize hardware* step of the wizard, select the appropriate network for *Network adapter 1*. Click **Next** to proceed.
10. For the *Customize template* step of the wizard, use the following table to complete the form:

Field	Description	Condition
Unique ID	This value must be unique within the vSphere networking domain. This field will be used to generate the hostname.	Required
SSH Private Key	This value will be used to allow key-based authentication with the Installer VM via SSH. <div style="border: 1px solid green; padding: 5px; margin: 5px 0;"> When creating a VM, you provided the public key, here you need to provide the private key of the public key that you used to install the VM. </div>	Recommended
Hostname	This value must be unique within the vSphere networking domain. This field will be used to generate the hostname.	Required

11. Click **Next** and then **Finish** to proceed. The OVA will start uploading - this will take approximately 5-10 minutes.

Recommendation: Once the OVA is finished uploading, it is recommended to create a VM Template from the uploaded installer image. This template can be used in future installations. Right-click on the OVA and select **Clone > Clone to Template**.

This completes the import of the CloudCenter Suite Installer into vSphere.

Configure the Offline Repository

To configure the offline repository, follow this procedure.

1. SSH into the offline repository using one of two methods.

The offline repository has the same user details as the CloudCenter Suite installer VM.

- a. Method 1: Using self-signed certificate.

```
sudo config-airgap-repo -i <ip address> -s
```

- b. Method 2: Using customized certificates.

```
sudo config-airgap-repo -c /tmp/certs/airgap-setup.cisco.com.crt -k airgap-setup.cisco.com.key -r ca.crt -i <ip address> # for user provided certificates
```

2. Verify that Harbor and its associated services are up and running and that the health of the system is successful as displayed in the following screenshot. This may take up to 20 seconds.

```
sudo docker ps # Verify the services are up.
```

```
sudo docker ps
CONTAINER ID        IMAGE                                     COMMAND                  CREATED
7f5393c333ac1     goharbor/harbor-jobservice:v1.9.4      "/harbor/harbor_jobs..." 2 minutes ago
8b2303339f129     goharbor/nginx-photon:v1.9.4          "nginx -g 'daemon of..." 2 minutes ago
65496cf5d0cd      goharbor/harbor-core:v1.9.4           "/harbor/harbor_core"    2 minutes ago
9e761da661f9      goharbor/redis-photon:v1.9.4          "redis-server /etc/r..." 2 minutes ago
cf7f33e1ee29     goharbor/chartmuseum-photon:v0.9.0-v1.9.4 "/docker-entrypoint..." 2 minutes ago
a7708ba3302c     goharbor/harbor-registryctl:v1.9.4    "/harbor/start.sh"       2 minutes ago
c161a9df5635     goharbor/harbor-portal:v1.9.4         "nginx -g 'daemon of..." 2 minutes ago
a93698e1fa6d     goharbor/registry-photon:v2.7.1-patch-2819-2553-v1.9.4 "/entrypoint.sh /etc..." 2 minutes ago
53e121d5951e     goharbor/harbor-db:v1.9.4             "/docker-entrypoint..." 2 minutes ago
44969e1a2168     goharbor/harbor-log:v1.9.4            "/bin/sh -c /usr/loc..." 2 minutes ago
```


3. Enter the password for this user:

```
Cisco123


# Be sure to change this default password.
```

4. Change the admin password using the following command.


```
sudo change-repo-password <oldpassword> <newpassword> # First time users use 'Cisco123' as the bootstrap password.
```

 Note down this **admin** password as you will need it in the later in this procedure!


5. Verify if the Harbor console is accessible via `https://<IP address>:8443`. Use **admin** as the username along with the newly updated password.




6. Click **Configuration > System Settings** and download the *Registry Root Certificate*.



- Once the CA certificate is downloaded, add it to your local keychain/truststore depending on the OS and verify that you can pull the Docker images from this appliance. You can view sample pull commands by navigating to any project > repository.



- To test, pull the Helm charts, add the offline repository as the Helm repository using the CA file downloaded along with the credentials.

 Use the **admin** password that you changed in Step 4 above.

For example:

```
helm repo add --username admin --password <YourNewAdminPassword> --ca-file ~/Downloads/ca-helm.crt
airgap https://10.11.84.50:8443/chartrepo

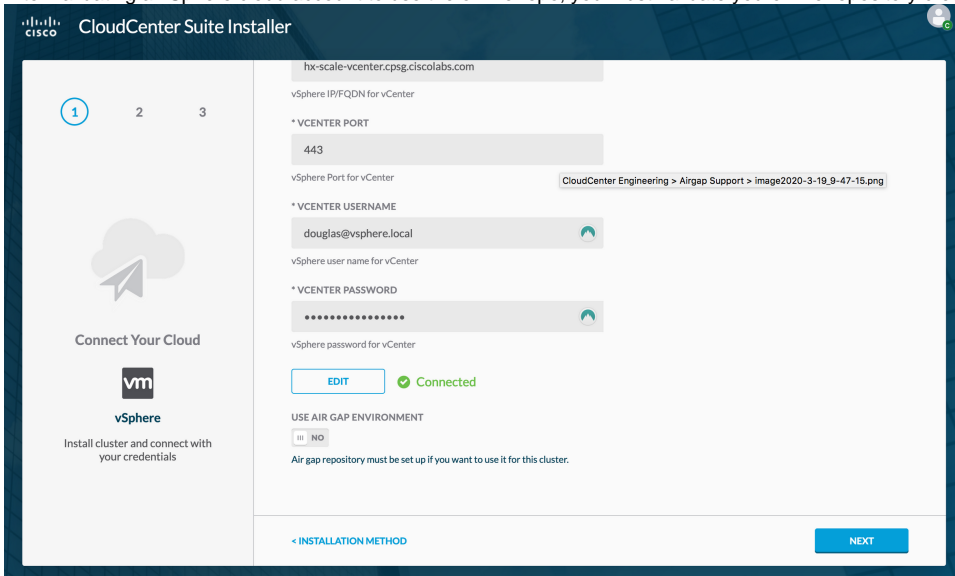
helm repo update
Hang tight while we grab the latest from your chart repositories...
...Successfully got an update from the "airgap" chart repository
Update Complete. Happy Helming!

helm search common-framework
NAME                                CHART VERSION    APP VERSION      DESCRIPTION
airgap/library/common-framework    5.2.0-16798     1.0              Common framework
multicloud suite
```

Configure the Installer to use the Offline Repository

To configure the Installer to use the Air Gap environment, follow this procedure.

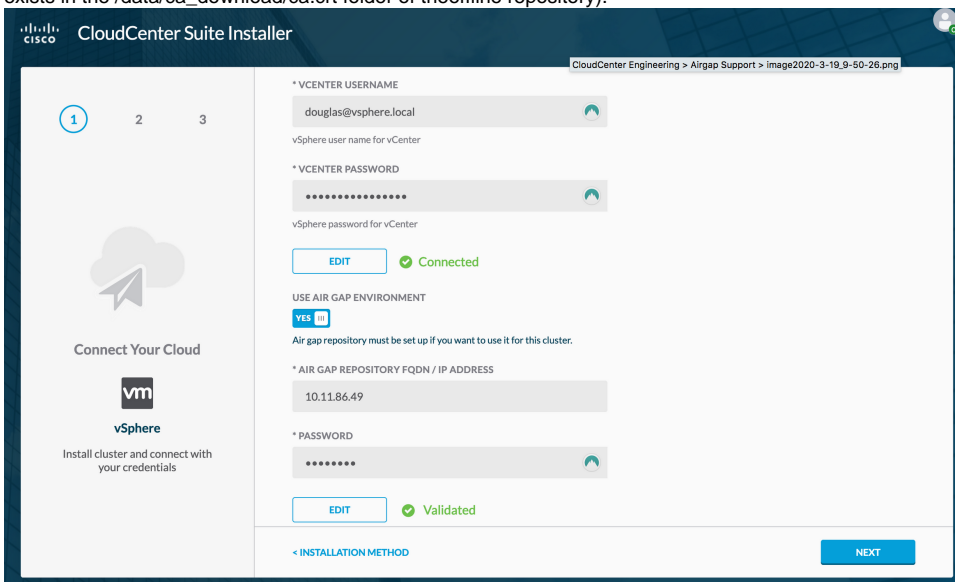
1. After validating a vSphere cloud account to use the offline repo, you must validate your offline repository credentials.




2. The following fields are required to validate your offline repository credentials:

- Offline Repository FQDN / IP Address (without port)
- Offline Repository password

3. After validating the vSphere cloud account, toggle the **Use Air Gap Environment** switch to **Yes** and provide the domain name, and password for the offline repository in the applicable fields. The installer fetches the certificate from the offline repository (as long as the required CA certificate exists in the /data/ca_download/ca.crt folder of the offline repository).




4. To continue the installation, click **Next** and continue with installation as usual. If you prefer to change back to a non-Air Gap setup, toggle the **Use Air Gap Environment** switch to **No** and click **Next**.

 You cannot change from an Air Gap to a non-Air Gap mode or vice-versa after moving away from this screen of the installation process. You must return to the first screen and restart this process if you choose to change at a later point.

Post Configuration Verification

This section identifies the verification process for each step in this process.

 These steps are only required for troubleshooting purposes if the installation fails at any point.

For the CloudCenter Suite Installer

To verify the CloudCenter Suite installation, follow this process.

1. SSH into the installer using the private key.
2. Check if the offline repo configurations and the certificates are stored in the file system by running the following commands.

installer verification steps

```
$ sudo -i
$ cat /home/cloud-user/.installer/k8s-mgmt/store/k8s-mgmt.airgap-repo-config
```

3. You should see the CA certificated displayed as follows (similar):

```
{
  "cacertificate": "-----BEGIN CERTIFICATE-----
\nMIIFFrzCCA5egAwIBAgIJAJ6s98GvsCp0MA0GCSqGSIb3DQEEDQUAMG4xCzAJBgNV\nnBAYTA
...
HxE\nkETWdSfgVVC3A1ZDYLNQWYKGM7i49jV7P0kOVjWuRQNFey/KZSXGvixqJjUtQzo\nn9y3RvH8Oi3CzKsr2AyMlcWT
/eMpB3qAMaJjBxXyngZJevVr12NJOuMG8jA jv104e\n3eap8
/MRptR9NpDvTVuVcOafW151ysnQmaOH7+N3dVHPPQ+AEKcj0Ck781GXNI1N\nnx6wymfE43cIPyvdHOxcpIOkQsMx0dx5RagMBAAGjUDBO
MB0GA1UdDgQWBRR54c3g\nDWHfulMZ2fVb2C/CXkQohTafBgNVHSMEGDAWgBR54c3gDWHfulMZ2fVb2C
/CXkQo\nnhTAMBgNVHRMERTADAQH/MA0GCSqGSIb3DQEEDQUAA4ICAQCaDMYnveDld41L8T4y\n\nlxJ8f7CiZDBGpML
/yI1Hjkl1tsh9BvyvhvAFjhjXzQphSxz15hzaJxnviPeCK/usI\nnq+cael71GFTet5xLfIU/fLq3/AxrvTeZCMz
/tSYU1shNUs4EiJKEBNTSLCjU1349\nnipz100fnCoByYORwFp7bQ3pHBTYZDUNI+VmuPL
/D50zqCB0vF0AC2uJhQAbSS9Xg\nni6bYfG9uYEMcHwTafL8fcw7YoSiaOL8wHGDUvNHP7m726BtH9D3rV0cV69a475EF\n\nndkOvzJcMq
/Zq0iOogGfe3K3Pof2Q74/6WW+j00ChKuD3NPVotVZh/INRLItH7ZGB\n/n/2NXSVrZ6S9mmuohzyD0xFCAXiPXnwjMC9Lo0
/4ah992eeKBXx1xw4+Ykid6Yjtg\nndQLang6J93ozKb4YJhlwmT8I+yad7RyHg8+4UTodlXxdqJXFZ2fSSGF9mbz
/ZD17\nnZ3IuWXUIe+nvsczBzw8yJlg3buJZlxbI7fDKCSwwXEUUO/IO7eoZa60kwCtCaY09\n\nnw5pnYuxJx8
/hRPLhbq8SQmU3dm7VhCPMMjLoTYo32dh6gpLi7Q6HKJVjii2iU9mA\n\nnio2rR1czsoG1qf3VGrgJ9nGfEk4RSag0Jp642JdUokWstPR1G
lp16ang5+EBTK9e\nnEiWISD6XT8lps04b5m2DZ79kig==\n-----END CERTIFICATE-----\n",
  "domainName": "10.11.4.0",
  "password": "*****",
  "username": "admin"
}
```

For the Launched Kubernetes Cluster

To verify the Kubernetes cluster, follow this process.

1. Check if the certificate is present on all the master nodes using the following script.

k8s-cluster verification steps

```
awk -v cmd='openssl x509 -noout -subject' '
/BEGIN/{close(cmd)};{print | cmd}' < /etc/ssl/certs/ca-certificates.crt
```

2. Verify if the offline repository certificate/CN is present in the certificate list generated by the script.



You can upgrade modules to a later version (when available) using the same offline repository certificates if you save the certificate details.

After installing CloudCenter Suite using the Offline Repository, be sure to take a backup of the certificates if the certificates were generated using SSL. This backup is for future reference so you can reuse the same certificates and configure the Offline Repository for a later version, when available.

3. Check if the offline repository configurations and certificates from the installer are migrated and stored in configmap.

```
$ kubectl get configmaps k8s-mgmt.offline-repo -n cisco -o yaml
```

Using Config Map

You can also verify the launched Kubernetes cluster using configmap as displayed in the following code block

```

kubect1 describe cm k8s-mgmt.airgap-repo-config -n cisco
Result:
Name:      k8s-mgmt.airgap-repo-config
Namespace: cisco
Labels:    <none>
Annotations: <none>

Data
====
data:
----
eyJjYWNlcnRpZmljYXRlIjoitFMwdExTMUNSVWRKVGlCRFJWSlVTVVpKUTBGVVJTMHRMUzB0Q2sxs1NVWnlla05EUVRWbFowRjNTVUpCWjBsS1FW
QnpUMUEwYUZwWE1USjFUVUV3UjBOVGNVZFRtV016UkZGRlFrUlJWVUZOUnpSNFEzcEJTa0puVGxZS1FrRlpWRUZZVmxST1VYTjNRMUZaUkZaU1Vv
bEVrVXBfVZSR1VrMUJPRWRCTVZWR1FuZDNTV1V5Um5WSlJYQjJZekpWZUVScVFVMUNaMDVXUWtGdlRRCeNwVTV3WxPKT2RrMVJOSGRFUVZsRVZs
RlJURVJCVMtSaFdFNXFZbnBGWmsxQ01FZEJNV1ZGUVhkM1YxbPhiSGxhTWtaM1RGaE9iR1JjVm5kTWJVNXdDbU15VG5aTWJGVNTJZbFJCWlVaM01I
bE5SRUY2VFZSbmVFMtZUVEJOVkJZaaFJuY3dlazFFUvVhTlZGbdRUWHBOTUUXVZtRk5SelI0UTNwQ1NrSm5UbFlLUWtGw1ZFRnNwbFJOVhOM1Ex
RlpSRlpSVVVsRVFVcEUVV1JGVWsxQk9FZEJNV1ZGUW5kM1NWVX1SblZKUlhCM11
...
VTNwWFRTOVFlR2RHYjFGMFpEQk1RV3R4Ww10b09GTTBtekJIY1hsRlJXZzNwWGRWVudSaFREWTFaQXBxV1RwcFFYQk1krWh2ZVRWbE16W11jRtVR
WjFsQ1JHwklkalpEzDJ4UFExSxpURTA1VW1ZeWFFWnNUWghzY0U1b1lWwM51RW93T0hSTE1rUjNNSFpTQ2t0d1QySmFhbTA0VjJwbkwxWjVNM1ZL
YUc5a2VWQ1pNRgt3T1haWU5FaHhSMUJIT0hwTlowcG5TRUZrZW0xYWMzQ1lNRfYyTTBweE9UWndjbmRPZG1jS2VXNXZNVFJ4YTJseVZFY3lWRWxQ
V0Uxck5tTjJiVkZ5ZEdwTE1tcDJjRFZDTXpJmU5IRjVZbt1RVFdoYVRHMVlRMghwYw5jM1ZuQjVpVE15UTNwWVJBcEdSbEJoV1h0aVrtUkZnWgcz
ZVhCTVpWVkl1VWhWVDA1b2R6UXhTek5EY2twRGNWUkRiWHA1WkdrM1RveGpSU3N2V1c1c1pqZHNjUz10V2pSRWF5dElDbkplUWt3ck1EQ1NRA2xI
TUVwM2EwWjJSelpmZG10eViYsm5QVDBLTFMwdExTMUZUa1FnUBTWU1ZfbEdTVU5CVkVvVdExTMHRMUW89Iiwib2ZmbGluZV9yZXBvX2FkZHZJlc3Mi
OiIxMC4xMS44Ni40OSIsIm9mZmxpbmVfcmlvYXNzd29yZCI6Ikpnc2NvMTIzIiwib2ZmbGluZV9yZXBvX3BvcnQiOiI4NDQzIiwib2ZmbGlu
ZV9yZXBvX3VzZXJlIjoieWRtaW4ifQ==
Events:  <none>

```

When Is the Offline Repository Used?

Once you complete the CloudCenter Suite installation and see the **Take Me to the Suite Admin** screen, the CloudCenter Suite pulls the information from the offline repository. This transition works seamlessly as it does in situations where you have internet connectivity!

Alternate Path of Airgap Certificates

Before beginning the below steps to upgrade the offline repository, you need to copy the following files under the /tmp/certs folder to the new offline repository.

- airgap-setup.cisco.com.crt
- airgap-setup.cisco.com.key
- ca.crt

If the /certs folder is deleted from the /tmp folder, the above files can be found under the /data folder with the different names as below:

- /tmp/certs/airgap-setup.cisco.com.key -> /data/certs/harbor.key
- /tmp/certs/airgap-setup.cisco.com/crt -> /data/certs/harbor.crt
- /tmp/ca.crt -> /data/ca_download/ca.crt

To upgrade the offline repository, follow this procedure.

1. Note the IP address and take back up of certificates of your current Air Gap environment.
2. Power off your current offline repository.
3. Create a new offline repository using the information provided in the Air Gap Installation section.
4. SSH into the offline repository using customized certificates.

Upgrade Offline Repository

Upgrade Offline Repository for an Air Gap Setup

- [Overview](#)
- [Restrictions](#)
- [Prerequisites](#)
- [Upgrade the Offline Repository](#)
- [Verify that the Offline Repository Is Correctly Upgraded](#)
- [Alternate Path of Airgap Certificates](#)

Overview

This section provides details on restrictions, prerequisites, and the process to upgrade the offline repository in an Air Gap environment. During this upgrade, the software upgrades create a new repository using the **SAME** certificates and IP address.

Restrictions

Before proceeding with an upgrade, adhere to the following restrictions:

- **Usage:** To upgrade the Air Gap environment to a new version, you can only use the CloudCenter Suite 5.3.0 upgrader to upgrade to a later release.
- **Suite Admin-level Permissions:** Suite Admin-level permissions are mandatory for a user to upgrade the cluster.

Prerequisites

Verify that the cluster adheres to the following requirements:

- **Backup Environment:** Back up your environment before initiating the upgrade. See [Backup Approach](#) for additional details.
- **Schedule Downtime:** Schedule a suitable downtime during off-peak hours to minimize the impact to your users and or customers. Communicate the downtime as the CloudCenter Suite will not be accessible during the upgrade.
- **Action Orchestrator environments:** The back up process for Action Orchestrator environments is different than from other CloudCenter Suite modules. See [Migrating Database to Install Action Orchestrator 5.2.0](#) to ensure that this processes has already been addressed.

Upgrade the Offline Repository

To upgrade the offline repository, follow this procedure.

1. Note the IP address and take back up of certificates of your current Air Gap environment.
2. Power off your current offline repository.
3. Create a new offline repository using the information provided in the [Air Gap Installation](#) section.
4. SSH into the offline repository. sing customized certificates.



The offline repository has the same user details as the CloudCenter Suite installer VM.

Be sure to use the SAME IP address and certificates for your current air gap environment that you noted down in Step 1 above.

```
sudo config-airgap-repo -c /tmp/certs/airgap-setup.cisco.com.crt -k airgap-setup.cisco.com.key -r ca.crt  
-i <ip address> # for user provided certificates
```

5. Verify that Harbor and its associated services are up and running and that the health of the system is successful as displayed in the following screenshot. This may take up to 20 seconds.

```
sudo docker ps # Verify the services are up.
```

```

sudo docker ps
CONTAINER ID        IMAGE                                     COMMAND                                     CREATED
7f5393e33ae1      goharbor/harbor-jobservice:v1.9.4      "/harbor/harbor_jobs..."             2 minutes ago
8b230339f129      goharbor/nginx-photon:v1.9.4          "nginx -g 'daemon of..."             2 minutes ago
65496cf5d0cd      goharbor/harbor-core:v1.9.4           "/harbor/harbor_core..."             2 minutes ago
9e761da661f9      goharbor/redis-photon:v1.9.4          "redis-server /etc/r..."             2 minutes ago
cf7f33e1ee29      goharbor/chartmuseum-photon:v0.9.0-v1.9.4 "/docker-entrypoint..."             2 minutes ago
a7708ba3302c      goharbor/harbor-registryctl:v1.9.4    "/harbor/start.sh"                    2 minutes ago
c161a9df5635      goharbor/harbor-portal:v1.9.4         "nginx -g 'daemon of..."             2 minutes ago
a93698e1fa6d      goharbor/registry-photon:v2.7.1-patch-2819-2553-v1.9.4 "/entrypoint.sh /etc..."             2 minutes ago
53e121d5951e      goharbor/harbor-db:v1.9.4            "/docker-entrypoint..."             2 minutes ago
44969e1a2168      goharbor/harbor-log:v1.9.4           "/bin/sh -c /usr/loc..."             2 minutes ago

```

6. Enter the password for this user:

```

Cisco123

# Be sure to change this default password.


```

7. Change the **admin** password using the following command.

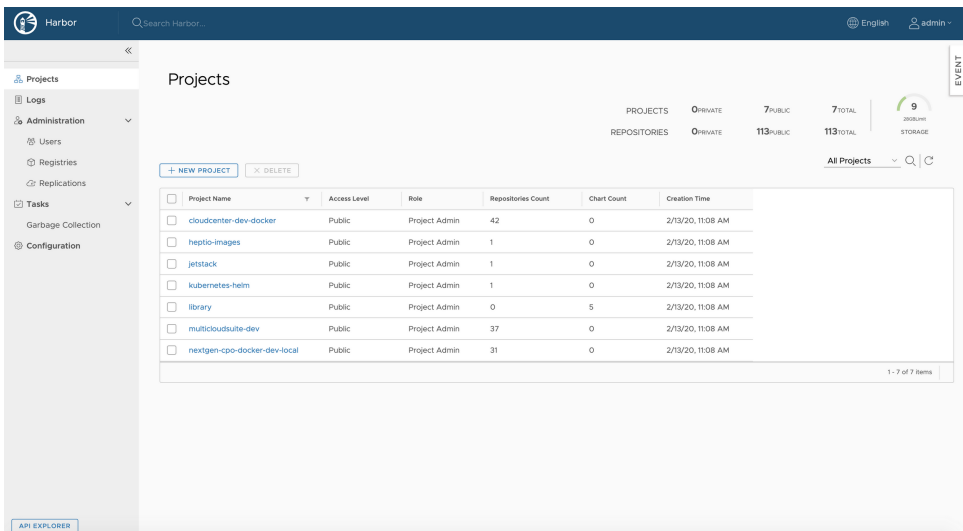
```

sudo change-repo-password <oldpassword> <newpassword> # First time users use 'Cisco123' as the bootstrap password.


```

 Note down this **admin** password as you will need it in the later in this procedure!

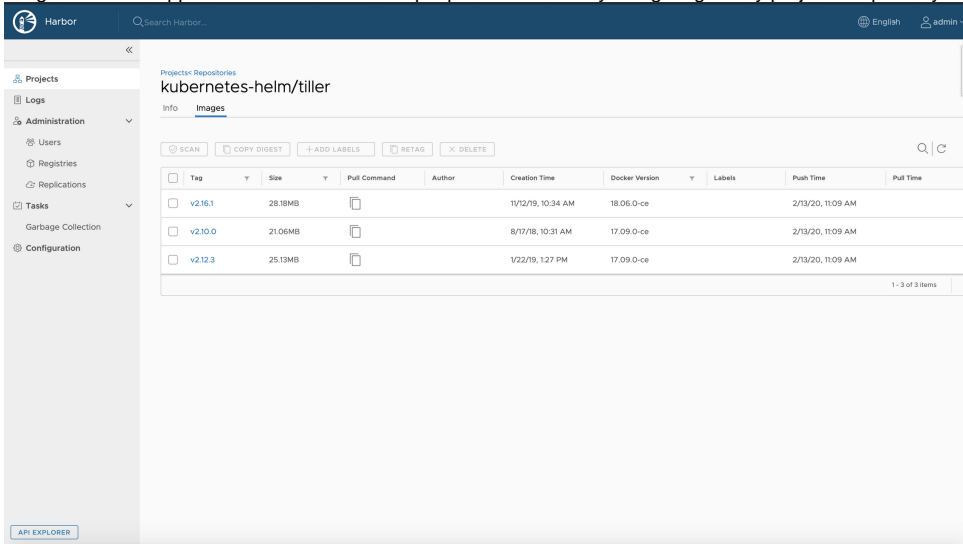
8. Verify if the Harbor console is accessible via `https://<IP address>:8443`. Use **admin** as the username along with the newly updated password.



9. Click **Configuration > System Settings** and download the *Registry Root Certificate*.



10. Once the CA certificate is downloaded, add it to your local keychain/truststore depending on the OS and verify that you can pull the Docker images from this appliance. You can view sample pull commands by navigating to any project > repository.



11. To test, pull the Helm charts, add the offline repository as the Helm repository using the CA file downloaded along with the credentials.

Use the **admin** password that you changed in Step 7 above.

For example:

```

helm repo add --username admin --password <YourNewAdminPassword> --ca-file ~/Downloads/ca-helm.crt
airgap https://10.11.84.50:8443/chartrepo

helm repo update
Hang tight while we grab the latest from your chart repositories...
...Successfully got an update from the "airgap" chart repository
Update Complete. Happy Helming!

helm search common-framework
NAME                                CHART VERSION  APP VERSION  DESCRIPTION
airgap/library/common-framework 5.2.0-16798    1.0          Common framework multicloud suite

```

Verify that the Offline Repository Is Correctly Upgraded

Once you complete the upgrade and see the **Take Me to the Suite Admin** screen, the CloudCenter Suite pulls the information from the offline repository. This transition works seamlessly as it does in situations where you have internet connectivity!

- If the **Repo Proxy Connectivity** icon is displayed in **green AFTER** you log into the Suite Admin, then you have set up the Offline Repository as listed in this section.
- If the **Repo Proxy Connectivity** icon is displayed in **red**, then the repo configuration has an issue perhaps an incorrect IP address or/and incorrect certificates. In this case:
 - Recheck your offline repository certificates and verify if they are applied correctly.
 - If nothing else works, repeat the procedure provided above.

Alternate Path of Airgap Certificates

Before beginning the below steps to upgrade the offline repository, you need to copy the following files under the /tmp/certs folder to the new offline repository.

- airgap-setup.cisco.com.crt
- airgap-setup.cisco.com.key
- ca.crt

If the /certs folder is deleted from the /tmp folder, the above files can be found under the /data folder with the different names as below:

- /tmp/certs/airgap-setup.cisco.com.key -> /data/certs/harbor.key
- /tmp/certs/airgap-setup.cisco.com/crt -> /data/certs/harbor.crt
- /tmp/ca.crt -> /data/ca_download/ca.crt

To upgrade the offline repository, follow this procedure.

1. Note the IP address and take back up of certificates of your current Air Gap environment.
2. Power off your current offline repository.
3. Create a new offline repository using the information provided in the Air Gap Installation section.
4. SSH into the offline repository using customized certificates.

Backup and Restore

Backup and Restore

- [Public Cloud](#)
 - [Backup Approach](#)
 - [Restore Approach](#)
 - [Restore without Proxy](#)
 - [Restore with Proxy](#)
- [Private Cloud](#)

Public Cloud

Public Cloud with Internet Access

- Backup Approach
- Restore Approach
 - Restore without Proxy
 - Restore with Proxy

Backup Approach


Backup Approach

- [Overview](#)
- [Limitations](#)
- [What Data Is Backed Up?](#)
- [Requirements](#)
- [Process](#)
- [Actions after Configuring the Backup](#)


Overview

You may sometimes need to backup your CloudCenter Suite setup so you have the option to recover the data when required. When you have a cluster running, it can go into a bad state for a number of reasons (resource shortage, application unavailability, infrastructure changes, undependable state and so forth). In these cases, backing up the data allows you to recover data when required.

If you are backing up data in the previous release clusters (for example, 5.2.3 clusters), update all module charts to the current version.

 The backup/restore feature is only available on *new* %ccs clusters installed using CloudCenter Suite installers and *not on* existing Kubernetes clusters.


Limitations

 For isolated, air gap, environments, that do not have internet access, or to back up to a local system, a manual backup procedure is available see [Private Cloud](#) for additional details.


Before proceeding with a backup, adhere to the following limitations:

- **Supported Clouds:** You can backup data to one of the following locations:
 - Google Cloud Storage (use the procedure below)
 - AWS S3 (use the procedure below)
- **Switching between Clouds and Cloud Accounts:**
 - While editing the storage location in the CloudCenter Suite, if you switch to a new cloud type or cloud account within the same cloud type, be aware that backups in the previously configured storage location will no longer be accessible from the CloudCenter Suite.
 - The backup files from the previously configured storage location will continue to be available via your cloud console.
- **Restoring to a Different Cluster:**
 - This feature is only supported for clusters launched by the %ccs installer.
 - You cannot backup from and restore to the same cluster you **can only** backup to one cluster and restore to a different cluster.
 - The backed up cluster and the target restore cluster should both be on the same cloud.
 - The backup taken on private clouds after running the `pod_vol_restic_scan.py` script skips backup of elasticsearch-master and elasticsearch-data pods. When this backup is restored on a different cluster you will not see logs in Kibana.
- **User Credentials:**
 - The credentials are specific to your service account in the cloud and only the user with those credentials can configure and initiate the backup.
 - If you change the credentials you will see a warning message to indicate that you cannot access previous backups.

What Data Is Backed Up?

 The CloudCenter Suite does NOT provide a granular option to backup Kubernetes resources or application-specific databases. Additionally, you CANNOT take volume snapshots.

The CloudCenter Suite uses the *latest* cloud/cloud account and bucket configurations to retrieve the list of existing backups, displayed in the table in the **Admin > Backup** page (under the Data Recovery section in the Suite Admin UI).

 If you update the existing configuration for any reason, users cannot manage the backups from the earlier cloud/cloud account and bucket configuration.

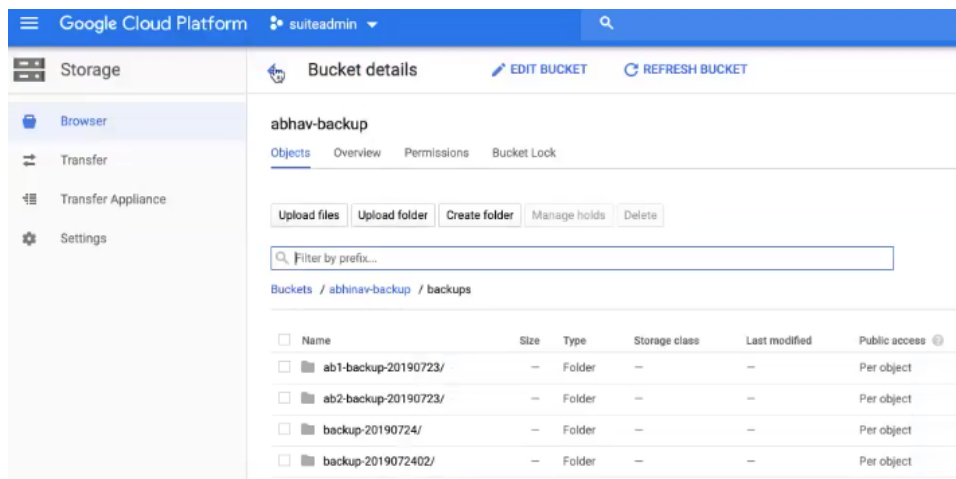
The backup action backs up the ENTIRE *cisco* namespace.

- **Backed Up:**
 - Any data under the Cisco (*cisco*) name space.
 - This includes users, groups, and roles for all modules.
 - This also includes but is not restricted to the Kubernetes resources with associated application data, pod data, secrets, PersistentVolumeClaim (PVC) data, PersistentVolume (PV) data, and other relevant data associated with these sub-systems
- **Not Backed Up:** Any data that is not under the Cisco (*cisco*) name space.
 - The backup taken on private clouds after running the `pod_vol_restic_scan.py` script skips backup of elasticsearch-master and elasticsearch-data pods.
 - Action Orchestrator Nuances:
 - The backup and restore procedures do not back up Action Orchestrator-specific data like workflows, targets, and so forth.
 - This type of Action Orchestrator-specific data is stored in arangoDB and requires arangodump and arangorestore to backup and restore the data.
 - To backup the data (Without internet access or proxy), the Arangodump should occur *before* you install the new Action Orchestrator version. See for additional details on [Private Cloud > Action Orchestrator-Specific Post-Restore Procedure](#) for additional details.
 - Action Orchestrator Backup Requirements:
 1. Backup the Action Orchestrator database using the arangodump tool. Uninstall %aofrom the CCS cluster.
 2. Backup Suite Admin, workload manager, and Cost Optimizer, Ringo, and Veler.
 - Action Orchestrator Restore Requirements:
 1. Restore Suite Admin, workload manager, and % using Veler. Reinstall Action Orchestrator.
 2. Restore the Action Orchestrator database using arangorestore tool

Requirements

Before proceeding with a backup, adhere to the following limitations:

- **General:** When configuring a backup for the first time, verify that the storage bucket is empty before scheduling any backups.
- **GCP:**
 - Configure a Storage Bucket with the required permissions: The following screenshot displays a sample storage bucket in a GCP environment:



- The cloud account used to configure the backup must have an empty **storage.bucket.list**.
- The bucket must have its ACL set to **storage.objects(create,delete,get,list)**.
- **AWS:**
 - The storage bucket in your AWS S3 environment must be empty with the applicable ACL permission.
 - The IAM user permissions define the user privilege on the S3 bucket as listed in the following screenshot:



In the following code block, the bucket name is defined as **velero-cisco** this is just an example! Be sure to change this value to reflect the name of your own bucket!

```


{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeRegions",
        "ec2:DescribeVolumes",
        "ec2:DescribeSnapshots",
        "ec2:CreateTags",
        "ec2:CreateVolume",
        "ec2:CreateSnapshot",
        "ec2>DeleteSnapshot"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:DeleteObject",
        "s3:PutObject",
        "s3:AbortMultipartUpload",
        "s3:ListMultipartUploadParts"
      ],
      "Resource": [
        "arn:aws:s3:::velero-cisco/*"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::velero-cisco"
      ]
    },
    {
      "Effect": "Allow",
      "Action": "s3:ListAllMyBuckets",
      "Resource": [
        "arn:aws:s3:::*"
      ]
    }
  ]
}

```

Process

To backup the CloudCenter Suite data, follow this procedure.

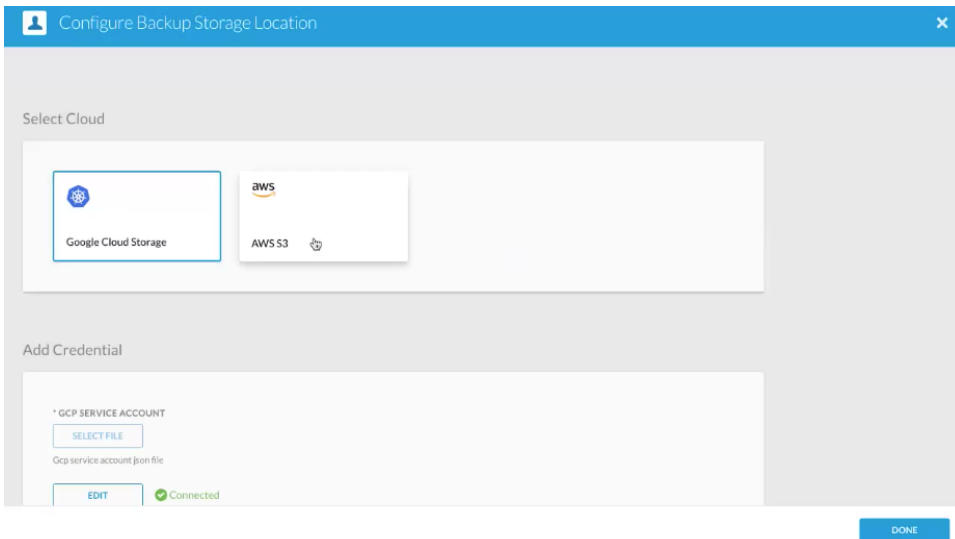
1. Navigate to the [Suite Admin Dashboard](#).
2. Click **Admin > Backup** (under the Data Recovery section) to access the Backup page as displayed in the following screenshot.



3. Click the **cog** icon in the Backup page (as displayed in the following screenshot) to configure a new backup storage location.



4. Select the required cloud in the Configure a Backup Storage Location page as displayed in the following screenshot.




5. Depending on the selected cloud, the Add Credential section differs:

- GCP:


a. Select the file containing the credentials is displayed in the following screenshot.



b. Select the Storage bucket as displayed in the following screenshot.




- c. Click **Done** to save the backup configuration as displayed in the following screenshot.




- AWS S3:

- a. Select the file containing the credentials as displayed in the following screenshot.




- b. Select the **Storage bucket** as displayed in the following screenshot.




- c. Click **Done** to save the backup configuration as displayed in the following screenshot.

Select Cloud



Google Cloud Storage



AWS S3

Add Credential

* AWS ACCESS KEY ID

AWS Access Key ID

* AWS SECRET ACCESS KEY

AWS Secret Access Key

* AWS REGION

AWS Region (optional)

AWS S3 ENDPOINT

AWS S3 Endpoint (optional)

✔ Connected

6. Once configured, click **Backup** in the Backup page to initiate the data backup. Until you initiate the first backup, this page will be empty. Once you have initiated one or more backups, they are automatically listed in this page as visible in the following screenshot.

CISCO SUITE ADMIN Welcome, Admin

Backup

NAME	CREATED DATE	CREATED BY	LOCATION	ACTIONS
ab1-backup-20190723	2 days ago	Admin Cllgrtech	gcp > abnav-backup	<input type="button" value="v"/>
ab2-backup-20190723	2 days ago	Admin Cllgrtech	gcp > abnav-backup	<input type="button" value="v"/>

7. In the Backup Name popup, assign a unique name (by default, the current date is listed) for this backup task and click **OK** as displayed in the following screenshot.

Backup Name


* BACKUP NAME

You have now backed up the CloudCenter Suite data to a cloud of choice.

Actions after Configuring the Backup

Once you have configured one or more backup settings in the Backup page, you may see the following actions in the Actions column.

- **Delete:** You can delete the configured backup as visible in the following screenshot:



- **Cancel:** You will only see the **Cancel** option when you are in the process of backing up a storage location. After you create the location, the only option you will see is **Delete**.

Back to: [Public Cloud](#)

Restore Approach

Restore Approach

- [Restore without Proxy](#)
- [Restore with Proxy](#)

Back to:[Public Cloud](#)

Restore without Proxy

Restore without Proxy

- [Overview](#)
- [Limitations](#)
- [Requirements](#)
 - [1. Launch the Target Cluster](#)
 - [2. Download the KubeConfig Files](#)
 - [3. Download Velero](#)
 - [4. Download JQ](#)
 - [5. Pre-Restore Procedure](#)
 - [6. Restore Procedure](#)
 - [7. Post-Restore Procedure](#)
- [workload manager-Specific Post-Restore Procedure Cloud Remote](#)
 - [a. Understand the Restore Context](#)
 - [b. Retrieve the Port Numbers from the NEW Restored Cluster](#)
 - [c. Retrieve the IP Address of the NEW Restored Cluster](#)
 - [d. Change the IP Address and Port Numbers for the NEW Restored Cluster](#)
 - [e. Perform the Pre-Migrate Activities](#)
 - [f. Migrate Deployments from the OLD Cluster to the NEW Cluster](#)

Overview

To restore data, the CloudCenter Suite requires that you launch a new cluster.



The backup/restore feature is only available on CloudCenter Suite clusters installed using CloudCenter Suite installers and not on existing Kubernetes clusters.

Limitations

If you configured the old cluster using a DNS, be sure to update the new IP address (from the restored cluster) that is mapped to the DNS entry. Once you update the DNS entry of your new cluster, these services will continue to work as designed.

Additionally, be aware that you may need to update the DNS for the [Base URL Configuration](#) and [SSO Setup](#) (both ADFS and SP).



Reconfiguration of Base URL and SSO are only applicable for backup & restore functions IF the source cluster is created using the CloudCenter Suite **5.0.x installer** and the destination cluster is freshly created using the CloudCenter Suite **5.1.1 installer**.

Requirements

Before proceeding with a restore, adhere to the following limitations:

- The Velero tool must be installed. Velero Version 1.5.3 - refer to <https://velero.io/docs/v1.5> for details.
- Launch a new cluster to restore the data.
- You will need to execute multiple scripts as part of these procedures. Make sure to use the 755 permission to execute each script mentioned in this section.

1. Launch the Target Cluster

To launch CloudCenter Suite on a new target cluster and access the Suite Admin UI for this cluster.

1. Navigate to the [Suite Admin Dashboard](#) for the new cluster.
2. Configure the identical backup configuration that you configured in your old cluster. See [Backup Approach > Process](#) additional details. When you provide the credentials, the new cluster automatically connects to the cloud storage location.



This step is REQUIRED to initiate the connection and fetch the backup(s).

3. Wait for a few minutes (at least 5 Mins, maybe more) for the Velero service in the new cluster to be synced up with the cloud storage location. At this point return to your local command window (shell console or terminal window) to perform the remaining steps in this process.



If both your clusters are accessible from your local machine, the scripts used in the following steps can be executed as designed.

If either one of your clusters uses proxy access or if you cannot recover/download the KubeConfig file from your old cluster, follow the instructions provided in the [Restore with Proxy](#) section.

2. Download the KubeConfig Files

You must download the KubeConfig file from the Suite Admin Kubernetes cluster management page for your source and target clusters to your local machine via a local command window (shell console or terminal window):

- From the source cluster, download the KubeConfig file and name it **KUBECONFIG_OLD**.
- From the target cluster, download the KubeConfig file and name it **KUBECONFIG_NEW**.

See [Kubernetes Cluster Management](#) for additional details on accessing the KubeConfig file as displayed in the following screenshot.

The screenshot shows the Suite Admin interface for a Kubernetes cluster. A red box highlights the 'Download KubeConfig File' button. Below the cluster information, there is a table of Virtual Machines.

NAME	IP ADDRESS	STATUS	CPU	MEMORY (GB)	RUNTIME
ab1473-8ca90e44-4e3f-4f2d-adf6-52a18809bf9a-mg-1-master-0	10.10.98.247	Up	2	16.82	1h
ab1473-8ca90e44-4e3f-4f2d-adf6-52a18809bf9a-mg-1-master-1	10.10.97.97	Up	2	16.82	1h
ab1473-8ca90e44-4e3f-4f2d-adf6-52a18809bf9a-mg-1-master-2	10.10.99.7	Up	2	16.82	1h

3. Download Velero

The restore process requires Velero and must be performed on a local command window (shell console or terminal window).

To download Velero, use one of the following options:

- OSX option:

```
$ cd <VELERO_DIRECTORY>
$ curl -L -O https://github.com/vmware-tanzu/velero/releases/download/v1.5.3/velero-v1.5.3-darwin-amd64.tar.gz
$ tar -xvf velero-v1.5.3-darwin-amd64.tar.gz
```

- CentOS Option:

```
$ mkdir -p /velero-test && cd /velero-test
$ curl -LO https://github.com/vmware-tanzu/velero/releases/download/v1.5.3/velero-v1.5.3-linux-amd64.tar.gz
$ tar -xvf velero-v1.5.3-linux-amd64.tar.gz && rm -rf velero-v1.5.3-linux-amd64.tar.gz
$ cp /velero-test/velero /usr/local/bin/
```

After you download Velero, export the KubeConfig file of the target (restore) cluster using the downloaded file:

```
export KUBECONFIG=<KUBECONFIG_PATH>
```

4. Download JQ

The restore process requires that you install JQ on your machine. Refer to <https://stedolan.github.io/jq/download> for additional details.

```
# To install jq on MacOS
$ brew install jq

# To install jq on Debian and Ubuntu
$ sudo apt-get install jq

# To install jq on CentOS
$ sudo yum install epel-release -y
$ sudo yum install jq -y
$ sudo jq --version
```

5. Pre-Restore Procedure

The pre-restore script creates the storageclass, if it does not exist on destination cluster, and saves the nginx-ingress-controller YAML file as well as the config maps for the following Suite Admin services:

- The suite-k8 service
- The suite-prod service

To execute the pre-restore script, run the **pre-restore.sh** script with the provided parameters:

```
# Command to execute the bashscript
$ ./pre-restore.sh 5101 </pathTo/oldCluster/kube_config> </pathTo/targetCluster/kube_config>

#Note: For 5.2.0 or later release, continue to provide the 5101 value.
#</pathTo/oldCluster/kube_config> is the path to the OLD KubeConfig file downloaded in Step 2.
#</pathTo/targetCluster/kube_config> is the path to the NEW KubeConfig file downloaded in Step 2.
{{}}
```



Make sure that the backup folder does not exist at `~/backup` on the device in which you are execute these scripts. If `~/backup` exists, delete it using the following command:

```
rm -rf ~/backup
```

The following code block includes the **pre-restore.sh** script:

```
#!/bin/bash
INSTALLER_VERSION_OLD=$1
KUBECONFIG_OLD=$2
KUBECONFIG_NEW=$3

declare INSTALLER_STORAGECLASS
INSTALLER_STORAGECLASS["500"]="thin"
INSTALLER_STORAGECLASS["501"]="thin"
INSTALLER_STORAGECLASS["502"]="thin"
INSTALLER_STORAGECLASS["51"]="standard"
INSTALLER_STORAGECLASS["510"]="standard"

if [[ ( ($KUBECONFIG_OLD == "" && $INSTALLER_VERSION_OLD == "") || $KUBECONFIG_NEW == "" ) ]]; then
    echo "Missing Paths for kubeconfigs"
    echo "Quitting"
    exit 0
else
    export KUBECONFIG_SAVED=$KUBECONFIG
    export KUBECONFIG=$HOME/.kube/config

    mkdir $HOME/backup
    cp $HOME/.kube/config $HOME/backup/saved_config

    if [[ $KUBECONFIG_OLD != "" ]]; then
```

```

        # Fetching the storage class name for the old(backup) cluster and storing it in variable
STORAGECLASS_NAME_OLD
        cp $KUBECONFIG_OLD $HOME/.kube/config
        STORAGECLASS_NAME_OLD=$(kubectl get storageclass -o json | jq '.items[0].metadata.name' | sed -e 's/^"
//' -e 's/"$//') # Extracting the storage class name from the json file of old cluster
        echo "Creating storage class "${STORAGECLASS_NAME_OLD} "in the target cluster."

    else
        echo "Creating storage class "${INSTALLER_STORAGECLASS[$INSTALLER_VERSION_OLD]} "in the target cluster."
        STORAGECLASS_NAME_OLD=${INSTALLER_STORAGECLASS[$INSTALLER_VERSION_OLD]}
    fi

    # Creating a storage class with the name STORAGECLASS_NAME_OLD in the target(restore) cluster
    cp $KUBECONFIG_NEW $HOME/.kube/config
    kubectl get storageclass -o json | jq --arg inpl $STORAGECLASS_NAME_OLD '.items[0].metadata.name=$inpl' >
$HOME/backup/storageclass.json
    cat $HOME/backup/storageclass.json | kubectl create -f -

    #setting the old storage class as "not default"
    if [[ $STORAGECLASS_NAME_OLD != "standard" ]]; then
        kubectl annotate --overwrite storageclass $STORAGECLASS_NAME_OLD storageclass.beta.kubernetes.io/is-
default-class='false' -n cisco
    fi

    #Scripts to backup ingress service spec, k8s, proxy settings, ssh keys and prod-mgmt configmaps on the
target cluster
    mkdir -p $HOME/backup/configmap
    mkdir -p $HOME/backup/service
    mkdir -p $HOME/backup/sshkeys
    mkdir -p $HOME/backup/proxy

    kubectl get svc -n cisco common-framework-nginx-ingress-controller -o json > $HOME/backup/service/ingress.
json

    for cm in $(kubectl get configmaps -n cisco -o custom-columns=:metadata.name --no-headers=true | grep "k8s-
mgmt")
    do
        kubectl get configmap $cm -n cisco -o yaml > $HOME/backup/configmap/$cm
    done

    for cm in $(kubectl get configmaps -n cisco -o custom-columns=:metadata.name --no-headers=true | grep "prod-
mgmt")
    do
        kubectl get configmap $cm -n cisco -o yaml > $HOME/backup/configmap/$cm
    done

    kubectl get configmap suite.key -n cisco -o yaml > $HOME/backup/sshkeys/suite.key
    kubectl get configmap suite.pub -n cisco -o yaml > $HOME/backup/sshkeys/suite.pub

    kubectl get configmap proxy.settings -n cisco -o yaml > $HOME/backup/proxy/proxy.settings

    kubectl set env deployment/common-framework-suite-prod-mgmt --list -n cisco | grep "CLOUD_TYPE" >> $HOME
/backup/proxy/proxy_variables
    kubectl set env deployment/common-framework-suite-prod-mgmt --list -n cisco | grep "HTTP_PROXY" >> $HOME
/backup/proxy/proxy_variables
    kubectl set env deployment/common-framework-suite-prod-mgmt --list -n cisco | grep "HTTPS_PROXY" >> $HOME
/backup/proxy/proxy_variables
    kubectl set env deployment/common-framework-suite-prod-mgmt --list -n cisco | grep "NO_PROXY" >> $HOME
/backup/proxy/proxy_variables

    cp $HOME/backup/saved_config $HOME/.kube/config

    export KUBECONFIG=$KUBECONFIG_SAVED

fi

echo 'Successful!'

```


6. Restore Procedure

To restore the backed up data to the target cluster, run the following Velero commands from your local machine.

1. List available backups.

```
$ ./<VELERO_DIRECTORY>/velero backup get
```



Verify if the backups are listed BEFORE proceeding to the next step.

2. Make sure the backed up *cisco* namespace does not exist in the target cluster. Be sure to delete the *cisco* name space, if it exists, before you restore.

```
$ kubectl delete ns cisco
```

3. Restore from one of the listed backups.

```
$ ./velero restore create --from-backup <BACKUPNAME>
```

You have now restored the CloudCenter Suite data to the new cluster.

7. Post-Restore Procedure

At this stage, you must restore the config maps for the following Suite Admin services:

- The suite-k8 service
- The suite-prod service

If the new cluster is accessible (from the local device) using theKubeConfig file, execute the following post-restore.sh script.

With Internet Access - The post-restore.sh script

```
#!/bin/bash

KUBECONFIG_NEW=$1

if [[ ( $KUBECONFIG_NEW == "" ) ]]; then
    echo "Missing Paths for kubeconfig"
    echo "Quitting"
    exit 0
else
    export KUBECONFIG_SAVED=$KUBECONFIG
    export KUBECONFIG=$HOME/.kube/config

    cp $HOME/.kube/config $HOME/backup/saved_config
    cp $KUBECONFIG_NEW $HOME/.kube/config

    kubectl delete svc -n cisco common-framework-nginx-ingress-controller
    cat $HOME/backup/service/ingress.json | kubectl create -f -

    for cm in $(ls $HOME/backup/configmap)
    do
        kubectl delete configmap $cm -n cisco
    done

    for cm in $(ls $HOME/backup/configmap)
    do
        cat $HOME/backup/configmap/$cm | kubectl create -f -
    done

    kubectl delete configmap suite.key -n cisco
    kubectl delete configmap suite.pub -n cisco
    kubectl delete configmap proxy.settings -n cisco
    cat $HOME/backup/sshkeys/suite.key | kubectl create -f -
    cat $HOME/backup/sshkeys/suite.pub | kubectl create -f -
    cat $HOME/backup/proxy/proxy.settings | kubectl create -f -

    while IFS= read -r line; do kubectl set env deployment/common-framework-suite-prod-mgmt $line -n cisco;
done < $HOME/backup/proxy/proxy_variables

    cp $HOME/backup/saved_config $HOME/.kube/config
    export KUBECONFIG=$KUBECONFIG_SAVED

    rm -r $HOME/backup/
fi

echo 'Successful!'
```

workload manager-Specific Post-Restore Procedure







This migration procedure only applies to **Running** deployments.

Be sure to verify that you are only migrating deployment in the **Running** state.



The first few steps differ based on your use of private clouds or public clouds. Be sure to use the procedure applicable to your cloud environment.

Cloud Remote Considerations

Scenario	Cloud Remote Configured	Settings	Notes
1	No	No additional settings	Proceed with the steps provided below, other than the note that only applies to Scenario 3. You must repeat this procedure for each region.
2	Yes	1. Cloud endpoint accessible from CloudCenter Suite = No 2. CloudCenter Suite AMQP reachable from worker VMs = No 3. CloudCenter Suite AMQP accessible from cloud = No	You do not need to perform any additional configurations and can skip this section. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  To ensure that the source (old) cluster does not connect to Cloud Remote, click Edit Connectivity in the Regions page and change the settings to Yes for all <i>three</i> settings. </div>
3		1. Cloud endpoint accessible from CloudCenter Suite = No 2. CloudCenter Suite AMQP reachable from worker VMs = No 3. CloudCenter Suite AMQP accessible from cloud = Yes	Proceed with the steps provided below, INCLUDING the note that is specific to this scenario. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  If you have multiple deployments that use both Scenario 1 and 3, you must perform these additional steps for deployments that use both Scenarios 1 and 3. You must repeat this procedure for each region. </div> You must repeat this procedure for each region.
4		1. Cloud endpoint accessible from CloudCenter Suite = Yes 2. CloudCenter Suite AMQP reachable from worker VMs = No 3. CloudCenter Suite AMQP accessible from cloud = No	You do not need to perform any additional configurations and can skip this section (similar to Scenario 2 above). <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  To ensure that the source (old) cluster does not connect to Cloud Remote, click Edit Connectivity in the Regions page and change the settings to Yes for all <i>three</i> settings. </div>
5		1. Cloud endpoint accessible from CloudCenter Suite = Yes 2. CloudCenter Suite AMQP reachable from worker VMs = No 3. CloudCenter Suite AMQP accessible from cloud = Yes	Proceed with the steps provided below, INCLUDING the note that is specific to this scenario (similar to Scenario 3 above). <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  If you have multiple deployments that use both Scenario 1 and 3, you must perform these additional steps for deployments that use both Scenarios 1 and 3. You must repeat this procedure for each region. </div> You must repeat this procedure for each region.

a. Understand the workload manager Restore Context

If you have installed the workload manager module, you must perform this procedure to update the DNS/IP address for the private cloud resources listed below and displayed in the following image:

- The Worker AMQP IP
- The Guacamole Public IP and Port
- The Guacamole IP Address and Port for Application VMs

Cloud endpoint accessible from CloudCenter Suite	Yes
CloudCenter Suite AMQP reachable from worker VM's	Yes
CloudCenter Suite AMQP accessible from cloud	Yes
Remote AMQP IP	
Worker AMQP IP	10.8.1.140:26642
Guacamole Public IP and Port	10.8.1.140:708
Guacamole IP Address and Port for Application VMs	10.8.1.140:32941
Blade Name	cloudcenter-blade-vmware-1-2033



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

b. Retrieve the Port Numbers from the NEW Restored Cluster

The Kubernetes cluster contains the information that is required to update the workload manager UI. This section provides the commands required to retrieve this information.



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

To retrieve the port numbers from the new cluster for private clouds, follow this procedure.

1. The port numbers for each component will differ.
 - a. Run the following command on the new cluster (login to the KubeConfig of the new cluster) to locate the new port numbers for the **Worker AMQP IP**.

```
kubectl get service -n cisco | grep rabbitmq-ext | awk '{print $5}'  
  
# In the resulting response, locate the port corresponding to Port 443 and use that port number!  
  
443:26642/TCP,15672:8902/TCP
```

- b. Run the following command on the new cluster to retrieve the port number for the **Guacamole Public IP and Port**.

```
kubectl get service -n cisco | grep cloudcenter-guacamole | awk '{print $5}'  
  
# In the resulting response, locate the port corresponding to Port 443 and use that port number  
for the Guacamole port!  
  
8080:2376/TCP,7788:25226/TCP,7789:32941/TCP,443:708/TCP
```

- c. Run the following command on the new cluster to retrieve the port number for the **Guacamole IP Address and Port for Application VMs**.

```
kubectl get service -n cisco | grep cloudcenter-guacamole | awk '{print $5}'  
  
# In the resulting response, locate the port corresponding to Port 7789 and use that port number  
for the Guacamole port!  
  
8080:2376/TCP,7788:25226/TCP,7789:32941/TCP,443:708/TCP
```

c. Retrieve the IP Address of the NEW Restored Cluster

Use the IP address of one of the primary servers of the NEW restored Kubernetes cluster for all the resources where the IP address needs to be replaced.



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

d. Change the IP Address and Port Numbers for the NEW Restored Cluster

The IP addresses and port numbers are not updated automatically in the workload manager UI and you must explicitly update them using this procedure.



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

To configure the IP address and port number in the new cluster, follow this procedure.

1. Access the workload manager module.

2. Navigate to **Clouds > Configure Cloud > Region Connectivity**.

Region Connectivity Running [Edit Connectivity](#)

Cloud endpoint accessible from CloudCenter Suite	Yes
CloudCenter Suite AMQP reachable from worker VM's	Yes
CloudCenter Suite AMQP accessible from cloud	Yes
Remote AMQP IP	
Worker AMQP IP	10.8.1.140:26642
Guacamole Public IP and Port	10.8.1.140:708
Guacamole IP Address and Port for Application VMs	10.8.1.140:32941
Blade Name	cloudcenter-blade-vmware-1-2033

Strategy [Edit Strategy](#)

Strategy Bundle

3. Click **Edit Connectivity** in the Region Connectivity settings.

4. In the Configure Region popup, change the 3 fields mentioned above to ensure that the IP and port details are updated to the NEW restored VM.

Configure Region

IS CLOUD END POINT DIRECTLY ACCESSIBLE?
 YES

SHOULD WORKER VMS DIRECTLY CONNECT WITH CLOUDCENTER SUITE?
 YES

WORKER AMQP IP ADDRESS
10.8.1.140:26642

GUACAMOLE PUBLIC IP AND PORT
10.8.1.140:708

GUACAMOLE IP ADDRESS AND PORT FOR APPLICATION VMS
10.8.1.140:32941

OK

DO NOT MAKE ANY OTHER CONFIGURATION CHANGES!

5. Click **OK** to save your changes.

Saving your changes may not automatically update the information in the Region Connectivity settings. Be sure to refresh the page to see the saved information.

6. You have now updated the DNS/IP/Port for the restored WM for this particular cloud. If you have configured other clouds in this environment, be sure to repeat this procedure for each cloud. Once you complete this procedure for all configured clouds, you can resume new deployment activities using the workload manager.

Only for Scenario 3

i Only required for Scenario 3 in the Workload Manager table above

With [Cloud Remote](#) configured in your old cluster, you must also reconfigure Cloud Remote to communicate with the new cluster by following this procedure.

1. Click **Download Configuration** in the Region Connectivity section as displayed in the following screen shot.

Region Connectivity	Running	Download Configuration	Copy Encryption Key	Edit Connectivity
Cloud endpoint accessible from CloudCenter Suite	No			
CloudCenter Suite AMQP reachable from worker VM's	Yes			
CloudCenter Suite AMQP accessible from cloud	Yes			
Local AMQP IP	192.168.113.240:31364			
Worker AMQP IP and Port	192.168.113.240:31364			
Guacamole Public IP and Port	192.168.113.240:31740			
Guacamole IP Address and Port for Application VMs	192.168.113.240:32065			
Blade Name	cloudcenter-blade-vmware-8-23a5			

2. Click **Copy Encryption Key**.
3. Access the Cloud Remote UI.
4. Apply the downloaded configuration on the Cloud Remote.

e. Perform the Pre-Migrate Activities

Before you migrate the deployment details you need to ensure that you can connect to both clusters and have the required files to perform the migration.

To perform the pre-migrate activities, follow this procedure.

1. Verify that the OLD cluster VMs can reach the NEW cluster. The remaining steps in this procedure are dependent on this connectivity in your environment.
2. Save the contents of the following **actions.json** file using the same name and file extension to your local directory with a file type JSON format.

The actions.json file


```
{ "repositories": [], "actions": { "resource": null, "size": 2, "pageNumber": 0, "totalElements": 2, "totalPages": 1, "actionJaxbs": [ { "id": "57", "resource": null, "name": "AgentReConfig_Linux", "description": "", "actionType": "EXECUTE_COMMAND", "category": "ON_DEMAND", "lastUpdatedTime": "2019-09-19 22:14:54.245", "timeOut": 1200, "enabled": true, "encrypted": false, "explicitShare": false, "showExplicitShareFeature": false, "deleted": false, "systemDefined": false, "bulkOperationSupported": true, "isAvailableToUser": true, "currentlyExecuting": false, "owner": 1, "actionParameters": [ { "paramName": "downloadFromBundle", "paramValue": "true", "customParam": false, "required": true, "useDefault": false, "preference": "VISIBLE_UNLOCKED" }, { "paramName": "bundlePath", "paramValue": "http://10.0.0.3/5.1-release/ccs-bundle-artifacts-5.1.0-20190819/agent.zip", "customParam": false, "required": true, "useDefault": false, "preference": "VISIBLE_UNLOCKED" }, { "paramName": "script", "paramValue": "agent/agentReconfig.sh", "customParam": false, "required": true, "useDefault": false, "preference": "VISIBLE_UNLOCKED" }, { "paramName": "executeOnContainer", "paramValue": "false", "customParam": false, "required": true, "useDefault": false, "preference": "VISIBLE_UNLOCKED" }, { "paramName": "rebootInstance", "paramValue": "false", "customParam": false, "required": true, "useDefault": false, "preference": "VISIBLE_UNLOCKED" }, { "paramName": "refreshInstanceInfo", "paramValue": "false", "customParam": false, "required": true, "useDefault": false, "preference": "VISIBLE_UNLOCKED" } ], "actionResourceMappings": { "type": "VIRTUAL_MACHINE", "actionResourceFilters": { "cloudRegionResource": null, "serviceResource": null, "applicationProfileResource": null, "deploymentResource": null, "vmResource": { "type": "DEPLOYMENT_VM", "appProfiles": [ "all" ], "cloudRegions": [ "all" ], "cloudAccounts": [ "all" ], "services": [ "all" ], "osTypes": [], "cloudFamilyNames": [], "nodeStates": [], "cloudResourceMappings": [], "isEditable": true } }, "actionResourceMappingAncillaries": [], "actionCustomParamSpecs": [ { "paramName": "brokerHost", "displayName": "BrokerHost", "helpText": "Ip Address or Hostname of Rabbit MQ cluster", "type": "string", "valueList": null, "defaultValue": "", "confirmValue": "", "pathSuffixValue": "", "userVisible": true, "userEditable": true, "systemParam": false, "exampleValue": null, "dataUnit": null, "optional": false, "deploymentParam": false, "multiselectSupported": false, "useDefault": true, "valueConstraint": { "minValue": 0, "maxValue": 255, "maxLength": 255, "regex": null, "allowSpaces": true, "sizeValue": 0, "step": 0, "calloutWorkflowName": null }, "scope": null, "webserviceListParams": { "url": "", "protocol": "", "username": "", "password": "", "requestType": null, "contentType": null, "commandParams": null, "requestBody": null, "resultString": null, "secret": null, "tabularTypeData": null, "collectionList": [], "preference": "VISIBLE_UNLOCKED" }, { "paramName": "brokerPort", "displayName": "BrokerPort", "helpText": "RabbitMQ Port number", "type": "string", "valueList": null, "defaultValue": "", "confirmValue": "", "pathSuffixValue": "", "userVisible": true, "userEditable": true, "systemParam": false, "exampleValue": null, "dataUnit": null, "optional": false, "deploymentParam": false, "multiselectSupported": false, "useDefault": true, "valueConstraint": { "minValue": 0, "maxValue": 255, "maxLength": 255, "regex": null, "allowSpaces": true, "sizeValue": 0, "step": 0, "calloutWorkflowName": null }, "scope": null, "webserviceListParams": { "url": "", "protocol": "", "username": "", "password": "", "requestType": null, "contentType": null, "commandParams": null, "requestBody": null, "resultString": null, "secret": null, "tabularTypeData": null, "collectionList": [], "preference": "VISIBLE_UNLOCKED" } ] }
```

```

requestBody":null,"resultString":null},"secret":null,"tabularTypeData":null,"collectionList":[],"
preference":"VISIBLE_UNLOCKED"]}],{"id":"58","resource":null,"name":"AgentReConfig_Win","
description":"","actionType":"EXECUTE_COMMAND","category":"ON_DEMAND","lastUpdatedTime":"2019-09-19 22:
15:02.311","timeOut":1200,"enabled":true,"encrypted":false,"explicitShare":false,"
showExplicitShareFeature":false,"deleted":false,"systemDefined":false,"bulkOperationSupported":true,"
isAvailableToUser":true,"currentlyExecuting":false,"owner":1,"actionParameters":[{"paramName":"
downloadFromBundle","paramValue":"true","customParam":false,"required":true,"useDefault":false,"
preference":"VISIBLE_UNLOCKED"},{"paramName":"bundlePath","paramValue":"http://10.0.0.3/5.1-release/ccs-
bundle-artifacts-5.1.0-20190819/agent.zip","customParam":false,"required":true,"useDefault":false,"
preference":"VISIBLE_UNLOCKED"},{"paramName":"script","paramValue":"agent\\agentReconfig.ps1","
customParam":false,"required":true,"useDefault":false,"preference":"VISIBLE_UNLOCKED"},{"paramName":"
executeOnContainer","paramValue":"false","customParam":false,"required":true,"useDefault":false,"
preference":"VISIBLE_UNLOCKED"},{"paramName":"rebootInstance","paramValue":"false","customParam":false,"
required":true,"useDefault":false,"preference":"VISIBLE_UNLOCKED"},{"paramName":"refreshInstanceInfo","
paramValue":"false","customParam":false,"required":true,"useDefault":false,"preference":"
VISIBLE_UNLOCKED"}],"actionResourceMappings":[{"type":"VIRTUAL_MACHINE","actionResourceFilters":
[{"cloudRegionResource":null,"serviceResource":null,"applicationProfileResource":null,"
deploymentResource":null,"vmResource":{"type":"DEPLOYMENT_VM","appProfiles":["all"],"cloudRegions":
["all"],"cloudAccounts":["all"],"services":["all"],"osTypes":[],"cloudFamilyNames":[],"nodeStates":[],"
cloudResourceMappings":[]},"isEditable":true},{"cloudRegionResource":null,"serviceResource":null,"
applicationProfileResource":null,"deploymentResource":null,"vmResource":{"type":"IMPORTED_VM","
appProfiles":["all"],"cloudRegions":["all"],"cloudAccounts":["all"],"services":["all"],"osTypes":["all"],"
cloudFamilyNames":["all"],"nodeStates":["all"],"cloudResourceMappings":[]},"isEditable":true}]}],"
actionResourceMappingAncillaries":[{"paramName":"brokerHost","displayName":"
BrokerHost","helpText":"Ip Address or Hostname of Rabbit MQ cluster","type":"string","valueList":null,"
defaultValue":"","confirmValue":"","pathSuffixValue":"","userVisible":true,"userEditable":true,"
systemParam":false,"exampleValue":null,"dataUnit":null,"optional":false,"deploymentParam":false,"
multiselectSupported":false,"useDefault":true,"valueConstraint":{"minValue":0,"maxValue":255,"maxLength":
255,"regex":null,"allowSpaces":true,"sizeValue":0,"step":0,"calloutWorkflowName":null},"scope":null,"
webserviceListParams":{"url":"","protocol":"","username":"","password":"","requestType":null,"
contentType":null,"commandParams":null,"requestBody":null,"resultString":null},"secret":null,"
tabularTypeData":null,"collectionList":[]},"preference":"VISIBLE_UNLOCKED"},{"paramName":"brokerPort","
displayName":"BrokerPort","helpText":"RabbitMQ Port number","type":"string","valueList":null,"
defaultValue":"","confirmValue":"","pathSuffixValue":"","userVisible":true,"userEditable":true,"
systemParam":false,"exampleValue":null,"dataUnit":null,"optional":false,"deploymentParam":false,"
multiselectSupported":false,"useDefault":true,"valueConstraint":{"minValue":0,"maxValue":255,"maxLength":
255,"regex":null,"allowSpaces":true,"sizeValue":0,"step":0,"calloutWorkflowName":null},"scope":null,"
webserviceListParams":{"url":"","protocol":"","username":"","password":"","requestType":null,"
contentType":null,"commandParams":null,"requestBody":null,"resultString":null},"secret":null,"
tabularTypeData":null,"collectionList":[]},"preference":"VISIBLE_UNLOCKED"}]}],"
repositoriesMappingRequired":false,"actionTypesCounts":[{"key":"EXECUTE_COMMAND","value":"2"}]}

```

3. Access workload manager in your OLD cluster and navigate to the Actions Library page.
4. Import the actions.json file that you saved in Step 2 above. You should see two files (**AgentReconfig_Linux** and **AgentReconfig_Win**) as displayed in the following screenshot.



5. The files are disabled by default (OFF) enable both files by toggling each switch to ON.
6. Save the following script to a file in your local directory and name it **agentReconfig.sh**. This is the file to use for Linux environments.

The agentReconfig.sh file

```

#!/bin/bash

#Write to system log as well as to terminal
logWrite()
{
    msg=$1
    echo "$(date) ${msg}"
    logger -t "OSMOSIX" "${msg}"
    return 0
}

logWrite "Starting agent migrate..."

env_file="/usr/local/osmosix/etc/userenv"

```

```

if [ -f $env_file ];
then
    logWrite "Source the userenv file..."
    . $env_file
fi

if [ -z $brokerHost ];
then
    logWrite "Broker Host / Rabbit Server Ip not passed as action parameter"
    exit 3;
fi

if [ -z $brokerPort ];
then
    logWrite "Broker Port / Rabbit Server Port not passed as action parameter"
    exit 4
fi

replaceUserDataValue() {
    key=$1
    value=$2

    if [ -z $key ] || [ -z $value ];
    then
        logWrite "Command line arguments missing to update user-data file, key: $key, value:$value"
        return
    fi

    user_data_file="/usr/local/agentlite/etc/user-data"
    if [ -f $user_data_file ];
    then
        json_content=`cat $user_data_file`
        old_value=`echo $json_content | awk -F $key '{print $2}' | awk -F \ " '{print $3}'`
        sed -i 's@"$old_value"@"$value"@g' $user_data_file
    fi
}

}

export AGENT_HOME="/usr/local/agentlite"

logWrite "Updating the user data file"
replaceUserDataValue "brokerClusterAddresses" "$brokerHost:$brokerPort"

logWrite "Updating config.json file"
sed -i '/AmqpAddress/c\      "AmqpAddress": "'"$brokerHost":'$"$brokerPort"'"', ' "$AGENT_HOME/config/config.json"

cd $AGENT_HOME
echo "sleep 10" > execute.sh
echo "/usr/local/agentlite/bin/agent-stop.sh" >> execute.sh
echo "/usr/local/agentlite/bin/agent-start.sh" >> execute.sh
chmod a+x execute.sh
nohup bash execute.sh > /dev/null 2>&1 &

exit 0

```

7. Save the following script to a file in your local directory and name it **agentReconfig.ps1**. This is the file to use for Windows environments.

The agentReconfig.ps1 file

```
param (
    [string]$brokerHost = "$env:brokerHost",
    [string]$brokerPort = "$env:brokerPort"
)

$SERVICE_NAME = "AgentService"
$SYSTEM_DRIVE = (Get-WmiObject Win32_OperatingSystem).SystemDrive
. "$SYSTEM_DRIVE\temp\userenv.ps1"

if ($brokerHost -eq 0 -or $brokerHost -eq $null -or $brokerHost -eq "") {
    echo "Variable brokerHost not available in the env file"
    exit 1
}

if ($brokerPort -eq 0 -or $brokerPort -eq $null -or $brokerPort -eq "") {
    echo "Variable brokerPort not available in the env file"
    exit 2
}

$AGENTGO_PARENT_DIR = "$SYSTEM_DRIVE\opt"

echo "Check if AgentGo Parent directory exists. If not create it: '$AGENTGO_PARENT_DIR'"
if (-not (Test-Path $AGENTGO_PARENT_DIR)) {
    echo "Create $AGENTGO_PARENT_DIR..."
    mkdir $AGENTGO_PARENT_DIR
}
else {
    echo "$AGENTGO_PARENT_DIR already exists."
}

$AGENT_CONFIG="{0}\agentlite\config\config.json" -f $AGENTGO_PARENT_DIR
if (Test-Path $AGENT_CONFIG) {
    echo "Changing the config.json file with the new broker host $env:brokerHost and port $env:
brokerPort"
    $confJson = get-content $AGENT_CONFIG | out-string | convertfrom-json
    $confJson.AmqpAddress = "$($env:brokerHost): $($env:brokerPort)"
    $confJson | ConvertTo-Json | set-content $AGENT_CONFIG
}

$USER_DATA_FILE = "{0}\agentlite\etc\user-data" -f $AGENTGO_PARENT_DIR
if (Test-Path $USER_DATA_FILE) {
    echo "Changing user-data file with new broker host $env:brokerHost and port $env:brokerPort"
    $userDataJson = get-content $USER_DATA_FILE | out-string | convertfrom-json
    $userDataJson.brokerClusterAddresses = "$($env:brokerHost): $($env:brokerPort)"
    $userDataJson | ConvertTo-Json | set-content $USER_DATA_FILE
}

$AGENT_SERVICE_NAME = "AgentService"
echo "Stop-Service $AGENT_SERVICE_NAME" > $AGENTGO_PARENT_DIR\exec.ps1
echo "sleep 10" >> $AGENTGO_PARENT_DIR\exec.ps1
echo "Start-Service $AGENT_SERVICE_NAME" >> $AGENTGO_PARENT_DIR\exec.ps1

echo "Restarting agent"
Start-Process -filepath "powershell" -argumentlist "-executionpolicy bypass -noninteractive -file
`"$AGENTGO_PARENT_DIR\exec.ps1`""

echo "Agent set to restart after config changes"
```


8. Add these two files to a folder called **agent** (just an example) and compress the folder to create **agent.zip** with the same structure displayed here.

agent

agentReconfig.ps1

agentReconfig.sh

- 9. Move the **agent.zip** folder to an HTTP repository in your local environment that is accessible from the OLD and NEW clusters.

 This procedure uses the following URL as an example:
http://10.0.0.3/repo/agent.zip

You have now ensured cluster connectivity and saved the required files for the migration procedure.

f. Migrate Deployments from the OLD Cluster to the NEW Cluster

To migrate the deployment details from the old cluster to the new cluster, follow this procedure.


- 1. Navigate to the workload manager **Actions Libray** page and edit the **AgentReconfig_Linux** action. This procedure continues to use the Linux file
- 2. going forward. Scroll to the **Actions Definition** section and update the URL as displayed in the following screenshot.

Action Definition

* EXECUTE FROM BUNDLE
 YES

* LOCATION * URL
URL http://10.0.0.3/repo/agent.zip




* SCRIPT FROM BUNDLE
agent/agentReconfig.sh

 The URL and Script from Bundle fields in the above screenshot are in accordance with the steps above.

- 3. Scroll to the **Custom Fields** section and change the default value of the **Broker Host** to use the NEW cluster IP.

Custom Fields

If desired add custom fields to the action. They can be made to be user entered or defined here by you, locked and hidden

  BrokerHost 

* DISPLAY NAME
BrokerHost

* PARAMETER NAME
brokerHost

HELP TEXT
Ip Address or Hostname of Rabbit MQ cluster

* TYPE MAX LENGTH
String 255

DEFAULT VALUE

REQUIRED FIELD ?
 YES

4. Scroll down to the **Broker Port** and change the default to use the NEW Worker AMQP IP port (for example, 26642 in Step 8 above).

BrokerPort

* DISPLAY NAME
BrokerPort

* PARAMETER NAME
brokerPort

HELP TEXT
RabbitMQ Port number


* TYPE
String

MAX LENGTH
255

DEFAULT VALUE

REQUIRED FIELD ?
YES

5. Click **Done** to save your default configuration changes in the OLD cluster.
6. Navigate to the **Virtual Machines** page and locate the VM to migrate to the new cluster.
7. Click the **Actions** dropdown and verify if your newly modified actions are visible under the Custom Actions section in the dropdown list as visible in the following screenshot.



8. Click one of the actions and verify that the configured defaults are displayed in the Broker host and Broker port fields as indicated earlier.
9. Click **Submit** to migrate this VM to the new cluster.
10. Verify that the migration is complete by going to the Deployment page in your NEW cluster and the VM is listed as RUNNING (green line).
11. Repeat Steps 6 through 10 for each VM that needs to be migrated to the NEW cluster.

You have now migrated the deployment details from the old cluster to the new cluster

Back to: [Public Cloud](#)

Restore with Proxy

Restore with Proxy

- [Overview](#)
- [Limitations](#)
- [Requirements](#)
 - [1. Launch the Target Cluster](#)
 - [2. Download the KubeConfig Files](#)
 - [3. Download Velero](#)
 - [4. Download JQ](#)
 - [5. Pre-Restore Procedure](#)
 - [6. Restore Procedure](#)
 - [7. Post-Restore Procedure](#)
- [workload manager-Specific Post-Restore Procedure Cloud Remote](#)
 - [a. Understand the Source Restore Context](#)
 - [b. Retrieve the Port Numbers from the NEW Restored Cluster](#)
 - [c. Retrieve the IP Address of the NEW Restored Cluster](#)
 - [d. Change the IP Address and Port Numbers for the NEW Restored Cluster](#)
 - [e. Perform the Pre-Migrate Activities](#)
 - [f. Migrate Deployments from the OLD Cluster to the NEW Cluster](#)

Overview

To restore data, the CloudCenter Suite requires that you launch a new cluster.



The backup/restore feature is only available on CloudCenter Suite clusters installed using CloudCenter Suite installers and not on existing Kubernetes clusters.

Limitations

If you configured the old cluster using a DNS, be sure to update the new IP address (from the restored cluster) that is mapped to the DNS entry. Once you update the DNS entry of your new cluster, these services will continue to work as designed.

Additionally, be aware that you may need to update the DNS for the [Base URL Configuration](#) and [SSO Setup](#) (both ADFS and SP).



Reconfiguration of Base URL and SSO are only applicable for backup & restore functions IF the source cluster is created using the CloudCenter Suite **5.0.x installer** and the destination cluster is freshly created using the CloudCenter Suite **5.1.1 installer**.

Requirements

Before proceeding with a restore, adhere to the following limitations:

- The Velero tool must be installed. Velero Version 1.5.3 - refer to <https://velero.io/docs/v1.5> for details.
- Launch a new cluster to restore the data.
- You will need to execute multiple scripts as part of these procedures. Make sure to use the 755 permission to execute each script mentioned in this section.

1. Launch the Target Cluster

To launch CloudCenter Suite on a new target cluster and access the Suite Admin UI for this cluster.

1. Navigate to the [Suite Admin Dashboard](#) for the new cluster.
2. Configure the identical backup configuration that you configured in your old cluster. See [Backup Approach > Process](#) additional details. When you provide the credentials, the new cluster automatically connects to the cloud storage location.



This step is REQUIRED to initiate the connection and fetch the backup(s).

3. Wait for a few minutes (at least 5 Mins, maybe more) for the Velero service in the new cluster to be synced up with the cloud storage location. At this point return to your local command window (shell console or terminal window) to perform the remaining steps in this process.



If both your clusters are accessible from your local machine, the scripts used in the following steps can be executed as designed.

If either one of your clusters uses proxy access or if you cannot recover/download the KubeConfig file from your old cluster, follow the instructions provided in the [Restore with Proxy](#) section.

2. Download the KubeConfig Files

You must download the KubeConfig file from the Suite Admin Kubernetes cluster management page for your source and target clusters to your local machine via a local command window (shell console or terminal window):

- From the source cluster, download the KubeConfig file and name it **KUBECONFIG_OLD**.
- From the target cluster, download the KubeConfig file and name it **KUBECONFIG_NEW**.

See [Kubernetes Cluster Management](#) for additional details on accessing the KubeConfig file as displayed in the following screenshot.

The screenshot shows the Suite Admin interface for a Kubernetes Cluster. The 'Download KubeConfig File' button is highlighted with a red box. Below the cluster information, there is a table of Virtual Machines with the following data:

NAME	IP ADDRESS	STATUS	CPU	MEMORY (GB)	RUNTIME
ab1473-8ca90e44-4e3f-4f2d-adf6-52a18809bf9a-mg-1-master-0	10.10.98.247	Up	2	16.82	1h
ab1473-8ca90e44-4e3f-4f2d-adf6-52a18809bf9a-mg-1-master-1	10.10.97.97	Up	2	16.82	1h
ab1473-8ca90e44-4e3f-4f2d-adf6-52a18809bf9a-mg-1-master-2	10.10.99.7	Up	2	16.82	1h

3. Download Velero

The restore process requires Velero and must be performed on a local command window (shell console or terminal window).

To download Velero, use one of the following options:

- OSX option:

```
$ cd <VELERO_DIRECTORY>
$ curl -L -O https://github.com/vmware-tanzu/velero/releases/download/v1.5.3/velero-v1.5.3-darwin-amd64.tar.gz
$ tar -xvf velero-v1.5.3-darwin-amd64.tar.gz
```

- CentOS Option:

```
$ mkdir -p /velero-test && cd /velero-test
$ curl -LO https://github.com/vmware-tanzu/velero/releases/download/v1.5.3/velero-v1.5.3-linux-amd64.tar.gz
$ tar -xvf velero-v1.5.3-linux-amd64.tar.gz && rm -rf velero-v1.5.3-linux-amd64.tar.gz
$ cp /velero-test/velero /usr/local/bin/
```

After you download Velero, export the KubeConfig file of the target (restore) cluster using the downloaded file:

```
export KUBECONFIG=<KUBECONFIG_PATH>
```

4. Download JQ

The restore process requires that you install JQ on your machine. Refer to <https://stedolan.github.io/jq/download> for additional details.

```
# To install jq on MacOS
$ brew install jq

# To install jq on Debian and Ubuntu
$ sudo apt-get install jq

# To install jq on CentOS
$ sudo yum install epel-release -y
$ sudo yum install jq -y
$ sudo jq --version
```

5. Pre-Restore Procedure

If either one of your clusters uses proxy access or if you cannot recover/download the KubeConfig file from your old cluster, follow the instructions provided in this section.

1. SSH into one of the VMs in your old cluster and retrieve the storageclass names.



This step is required because of changes in the storageclass name between CloudCenter Suite 5.0.0 and 5.1.0.

```
$ kubectl get storageclass -o json | grep '"name"' | cut -d ':' -f 2 | sed 's/"/\//g' | sed 's/[\,]/ /g'
```

For example:

Example

```
$ kubectl get storageclass -o json | grep '"name"' | cut -d ':' -f 2 | sed 's/"/\//g' | sed 's/[\,]/ /g' "thin"
```

2. SSH into one of the VMs in your new cluster and retrieve the storageclass names:

```
$ kubectl get storageclass -o json | grep '"name"' | cut -d ':' -f 2 | sed 's/"/\//g' | sed 's/[\,]/ /g'
```

For example:

Example

```
$ kubectl get storageclass -o json | grep '"name"' | cut -d ':' -f 2 | sed 's/"/\//g' | sed 's/[\,]/ /g'
"standard"
```

3. Copy the contents of storageclass from the new cluster using the command below: (use the storageclass_name retrieve using the above step). You need to run the following command, copy the output, and save the output to a file called backupStorageclass.yaml.

```
$ kubectl get storageclass <storageclass_name> -o yaml
```

For example:

```

cloud-user@ab21461-fcc43751-1381-4e98-8d45-934bb965edfe-mg-1-primary-0:~$ kubectl get storageclass
standard -o yaml
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  annotations:
    kubectl.kubernetes.io/last-applied-configuration: |
      {"apiVersion":"storage.k8s.io/v1beta1","kind":"StorageClass","metadata":{"annotations":
{"storageclass.beta.kubernetes.io/is-default-class":"true"},"name":"standard"},"parameters":
{"diskformat":"thin"},"provisioner":"kubernetes.io/vsphere-volume"}
      storageclass.beta.kubernetes.io/is-default-class: "true"
  creationTimestamp: "2019-07-31T23:26:57Z"
  name: standard
  resourceVersion: "605"
  selfLink: /apis/storage.k8s.io/v1/storageclasses/standard
  uid: b045d700-b3ea-11e9-9b1d-0050569f28fd
parameters:
  diskformat: thin
  provisioner: kubernetes.io/vsphere-volume
  reclaimPolicy: Delete
  volumeBindingMode: Immediate

```

4. Create a new file backupStorageclass.yaml and paste the contents copied from the previous step.
5. Replace the field **name** in the backupStorageclass.yaml file with the OLD storage_classname from the old cluster from Step 1.

For example:

```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  annotations:
    kubectl.kubernetes.io/last-applied-configuration: |
      {"apiVersion":"storage.k8s.io/v1beta1","kind":"StorageClass","metadata":{"annotations":{"storageclass.beta.kubernetes.io/is-default-class":"t
storageclass.beta.kubernetes.io/is-default-class: "true"
  creationTimestamp: "2019-07-31T23:26:57Z"
→ name: thin
  resourceVersion: "605"
  selfLink: /apis/storage.k8s.io/v1/storageclasses/standard
  uid: b045d700-b3ea-11e9-9b1d-0050569f28fd
parameters:
  diskformat: thin
  provisioner: kubernetes.io/vsphere-volume
  reclaimPolicy: Delete
  volumeBindingMode: Immediate

```

6. Create a new storageclass in the new cluster using the command below

```
$ cat /path/backupStorageclass.yaml | kubectl create -f -
```

7. Create a backup of the Kubernetes config maps of the following services by executing the script provided in this step.
 - The suite-k8 service
 - The suite-prod service
8. Run the command to execute the backup_configmap.sh script

```

#Execute the script as sudo user
$ sudo /path/to/script/backup_configmap.sh.sh

```

The backup_configmap.sh script

backup_configmap.sh

```
#!/bin/bash

#Scripts to backup ssh keys, proxy settings, k8s and prod-mgmt configmaps on the target cluster
mkdir -p $HOME/backup/configmap
mkdir -p $HOME/backup/service
mkdir -p $HOME/backup/sshkeys
mkdir -p $HOME/backup/proxy

kubectl get svc -n cisco common-framework-nginx-ingress-controller -o json > $HOME/backup/service/ingress.json

for cm in $(kubectl get configmaps -n cisco -o custom-columns=:metadata.name --no-headers=true | grep "k8s-mgmt")
do
    kubectl get configmap $cm -n cisco -o yaml > $HOME/backup/configmap/$cm
done

for cm in $(kubectl get configmaps -n cisco -o custom-columns=:metadata.name --no-headers=true | grep "prod-mgmt")
do
    kubectl get configmap $cm -n cisco -o yaml > $HOME/backup/configmap/$cm
done

kubectl get configmap suite.key -n cisco -o yaml > $HOME/backup/sshkeys/suite.key
kubectl get configmap suite.pub -n cisco -o yaml > $HOME/backup/sshkeys/suite.pub

kubectl get configmap proxy.settings -n cisco -o yaml > $HOME/backup/proxy/proxy.settings

kubectl set env deployment/common-framework-suite-prod-mgmt --list -n cisco | grep "CLOUD_TYPE" >> $HOME/backup/proxy/proxy_variables
kubectl set env deployment/common-framework-suite-prod-mgmt --list -n cisco | grep "HTTP_PROXY" >> $HOME/backup/proxy/proxy_variables
kubectl set env deployment/common-framework-suite-prod-mgmt --list -n cisco | grep "HTTPS_PROXY" >> $HOME/backup/proxy/proxy_variables
kubectl set env deployment/common-framework-suite-prod-mgmt --list -n cisco | grep "NO_PROXY" >> $HOME/backup/proxy/proxy_variables

echo 'Successful!'
```

6. Restore Procedure

1. List available backups.



Verify if the backups are listed BEFORE proceeding to the next step.

```
$ ./<VELERO_DIRECTORY>/velero backup get
```

2. Make sure the backed up namespace does not exist in the target cluster (for example, if the *cisco* namespace was backed up it shouldn't be here on the cluster).

```
$ kubectl delete ns cisco
```

3. Restore from one of the listed backups.

```
$ ./velero restore create --from-backup <BACKUPNAME>
```

You have now restored the CloudCenter Suite data to the new cluster.

7. Post-Restore Procedure

At this stage, you must restore the config maps for the following Suite Admin services:

- The suite-k8 service
- The suite-prod service

If the new cluster is NOT accessible (from the local device) using kubeconfig, execute the following script from the remote device after the restore process is complete.

```
#Execute the script as sudo user
$ sudo /path/to/script/post-restore.sh
```

Without Internet Access - The post-restore.sh script

```
#!/bin/bash
kubectl delete svc -n cisco common-framework-nginx-ingress-controller
cat $HOME/backup/service/ingress.json | kubectl create -f -

for cm in $(ls $HOME/backup/configmap)
do
    kubectl delete configmap $cm -n cisco
done

for cm in $(ls $HOME/backup/configmap)
do
    cat $HOME/backup/configmap/$cm | kubectl create -f -
done

kubectl delete configmap suite.key -n cisco
kubectl delete configmap suite.pub -n cisco
kubectl delete configmap proxy.settings -n cisco
cat $HOME/backup/sshkeys/suite.key | kubectl create -f -
cat $HOME/backup/sshkeys/suite.pub | kubectl create -f -
cat $HOME/backup/proxy/proxy.settings | kubectl create -f -

while IFS= read -r line; do kubectl set env deployment/common-framework-suite-prod-mgmt $line -n cisco; done <
$HOME/backup/proxy/proxy_variables

rm -r $HOME/backup/configmap

echo 'Successfull!'
```

You have now restored the Suite Admin data to the new cluster. You can now follow the post-restore procedure specific to workload manager as provided in the next section.

workload manager-Specific Post-Restore Procedure



This migration procedure only applies to **Running** deployments.





Be sure to verify that you are only migrating deployment in the **Running** state.



The first few steps differ based on your use of private clouds or public clouds. Be sure to use the procedure applicable to your cloud environment.

Cloud Remote Considerations

Scenario	Cloud Remote Configured	Settings	Notes
----------	-------------------------	----------	-------

1	No	No additional settings	Proceed with the steps provided below, other than the note that only applies to Scenario 3. You must repeat this procedure for each region.
2	Yes	<ol style="list-style-type: none"> 1. Cloud endpoint accessible from CloudCenter Suite = No 2. CloudCenter Suite AMQP reachable from worker VMs = No 3. CloudCenter Suite AMQP accessible from cloud = No 	<p>You do not need to perform any additional configurations and can skip this section.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p> To ensure that the source (old) cluster does not connect to Cloud Remote, click Edit Connectivity in the Regions page and change the settings to Yes for all <i>three</i> settings.</p> </div>
3		<ol style="list-style-type: none"> 1. Cloud endpoint accessible from CloudCenter Suite = No 2. CloudCenter Suite AMQP reachable from worker VMs = No 3. CloudCenter Suite AMQP accessible from cloud = Yes 	<p>Proceed with the steps provided below, INCLUDING the note that is specific to this scenario.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p> If you have multiple deployments that use both Scenario 1 and 3, you must perform these additional steps for deployments that use both Scenarios 1 and 3.</p> <p>You must repeat this procedure for each region.</p> </div> <p>You must repeat this procedure for each region.</p>
4		<ol style="list-style-type: none"> 1. Cloud endpoint accessible from CloudCenter Suite = Yes 2. CloudCenter Suite AMQP reachable from worker VMs = No 3. CloudCenter Suite AMQP accessible from cloud = No 	<p>You do not need to perform any additional configurations and can skip this section (similar to Scenario 2 above).</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p> To ensure that the source (old) cluster does not connect to Cloud Remote, click Edit Connectivity in the Regions page and change the settings to Yes for all <i>three</i> settings.</p> </div>
5		<ol style="list-style-type: none"> 1. Cloud endpoint accessible from CloudCenter Suite = Yes 2. CloudCenter Suite AMQP reachable from worker VMs = No 3. CloudCenter Suite AMQP accessible from cloud = Yes 	<p>Proceed with the steps provided below, INCLUDING the note that is specific to this scenario (similar to Scenario 3 above).</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p> If you have multiple deployments that use both Scenario 1 and 3, you must perform these additional steps for deployments that use both Scenarios 1 and 3.</p> <p>You must repeat this procedure for each region.</p> </div> <p>You must repeat this procedure for each region.</p>

a. Understand the workload manager Restore Context

If you have installed the workload manager module, you must perform this procedure to update the DNS/IP address for the private cloud resources listed below and displayed in the following image:

- The Worker AMQP IP
- The Guacamole Public IP and Port
- The Guacamole IP Address and Port for Application VMs

Cloud endpoint accessible from CloudCenter Suite	Yes
CloudCenter Suite AMQP reachable from worker VM's	Yes
CloudCenter Suite AMQP accessible from cloud	Yes
Remote AMQP IP	
Worker AMQP IP	10.8.1.140:26642
Guacamole Public IP and Port	10.8.1.140:708
Guacamole IP Address and Port for Application VMs	10.8.1.140:32941
Blade Name	cloudcenter-blade-vmware-1-2033



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

b. Retrieve the Port Numbers from the NEW Restored Cluster

The Kubernetes cluster contains the information that is required to update the workload manager UI. This section provides the commands required to retrieve this information.



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

To retrieve the port numbers from the new cluster for private clouds, follow this procedure.

1. The port numbers for each component will differ.
 - a. Run the following command on the new cluster (login to the KubeConfig of the new cluster) to locate the new port numbers for the **Worker AMQP IP**.

```
kubectl get service -n cisco | grep rabbitmq-ext | awk '{print $5}'  
  
# In the resulting response, locate the port corresponding to Port 443 and use that port number!  
  
443:26642/TCP,15672:8902/TCP
```

- b. Run the following command on the new cluster to retrieve the port number for the **Guacamole Public IP and Port**.

```
kubectl get service -n cisco | grep cloudcenter-guacamole | awk '{print $5}'  
  
# In the resulting response, locate the port corresponding to Port 443 and use that port number  
for the Guacamole port!  
  
8080:2376/TCP,7788:25226/TCP,7789:32941/TCP,443:708/TCP
```

- c. Run the following command on the new cluster to retrieve the port number for the **Guacamole IP Address and Port for Application VMs**.

```
kubectl get service -n cisco | grep cloudcenter-guacamole | awk '{print $5}'  
  
# In the resulting response, locate the port corresponding to Port 7789 and use that port number  
for the Guacamole port!  
  
8080:2376/TCP,7788:25226/TCP,7789:32941/TCP,443:708/TCP
```

c. Retrieve the IP Address of the NEW Restored Cluster

Use the IP address of one of the primary servers of the NEW restored Kubernetes cluster for all the resources where the IP address needs to be replaced.



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

d. Change the IP Address and Port Numbers for the NEW Restored Cluster

The IP addresses and port numbers are not updated automatically in the workload manager UI and you must explicitly update them using this procedure.



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

To configure the IP address and port number in the new cluster, follow this procedure.

1. Access the workload manager module.

2. Navigate to **Clouds > Configure Cloud > Region Connectivity**.

Region Connectivity Running [Edit Connectivity](#)

Cloud endpoint accessible from CloudCenter Suite	Yes
CloudCenter Suite AMQP reachable from worker VM's	Yes
CloudCenter Suite AMQP accessible from cloud	Yes
Remote AMQP IP	
Worker AMQP IP	10.8.1.140:26642
Guacamole Public IP and Port	10.8.1.140:708
Guacamole IP Address and Port for Application VMs	10.8.1.140:32941
Blade Name	cloudcenter-blade-vmware-1-2033

Strategy [Edit Strategy](#)

Strategy Bundle

3. Click **Edit Connectivity** in the Region Connectivity settings.

4. In the Configure Region popup, change the 3 fields mentioned above to ensure that the IP and port details are updated to the NEW restored VM.

Configure Region

IS CLOUD END POINT DIRECTLY ACCESSIBLE?
 YES

SHOULD WORKER VMS DIRECTLY CONNECT WITH CLOUDCENTER SUITE?
 YES

WORKER AMQP IP ADDRESS
10.8.1.140:26642

GUACAMOLE PUBLIC IP AND PORT
10.8.1.140:708

GUACAMOLE IP ADDRESS AND PORT FOR APPLICATION VMS
10.8.1.140:32941

OK

DO NOT MAKE ANY OTHER CONFIGURATION CHANGES!

5. Click **OK** to save your changes.

Saving your changes may not automatically update the information in the Region Connectivity settings. Be sure to refresh the page to see the saved information.

6. You have now updated the DNS/IP/Port for the restored WM for this particular cloud. If you have configured other clouds in this environment, be sure to repeat this procedure for each cloud. Once you complete this procedure for all configured clouds, you can resume new deployment activities using the workload manager.

Only for Scenario 3

i Only required for Scenario 3 in the Workload Manager table above

With [Cloud Remote](#) configured in your old cluster, you must also reconfigure Cloud Remote to communicate with the new cluster by following this procedure.

1. Click **Download Configuration** in the Region Connectivity section as displayed in the following screen shot.

Region Connectivity	Running	Download Configuration	Copy Encryption Key	Edit Connectivity
Cloud endpoint accessible from CloudCenter Suite	No			
CloudCenter Suite AMQP reachable from worker VM's	Yes			
CloudCenter Suite AMQP accessible from cloud	Yes			
Local AMQP IP	192.168.113.240:31364			
Worker AMQP IP and Port	192.168.113.240:31364			
Guacamole Public IP and Port	192.168.113.240:31740			
Guacamole IP Address and Port for Application VMs	192.168.113.240:32065			
Blade Name	cloudcenter-blade-vmware-8-23a5			

2. Click **Copy Encryption Key**.
3. Access the Cloud Remote UI.
4. Apply the downloaded configuration on the Cloud Remote.

e. Perform the Pre-Migrate Activities

Before you migrate the deployment details you need to ensure that you can connect to both clusters and have the required files to perform the migration.

To perform the pre-migrate activities, follow this procedure.

1. Verify that the OLD cluster VMs can reach the NEW cluster. The remaining steps in this procedure are dependent on this connectivity in your environment.
2. Save the contents of the following **actions.json** file using the same name and file extension to your local directory with a file type JSON format.

The actions.json file

```
{
  "repositories": [],
  "actions": {
    "resource": null,
    "size": 2,
    "pageNumber": 0,
    "totalElements": 2,
    "totalPages": 1,
    "actionJaxbs": [
      {
        "id": "57",
        "resource": null,
        "name": "AgentReConfig_Linux",
        "description": "",
        "actionType": "EXECUTE_COMMAND",
        "category": "ON_DEMAND",
        "lastUpdatedTime": "2019-09-19 22:14:54.245",
        "timeOut": 1200,
        "enabled": true,
        "encrypted": false,
        "explicitShare": false,
        "showExplicitShareFeature": false,
        "deleted": false,
        "systemDefined": false,
        "bulkOperationSupported": true,
        "isAvailableToUser": true,
        "currentlyExecuting": false,
        "owner": 1,
        "actionParameters": [
          {
            "paramName": "downloadFromBundle",
            "paramValue": "true",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "bundlePath",
            "paramValue": "http://10.0.0.3/5.1-release/ccs-bundle-artifacts-5.1.0-20190819/agent.zip",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "script",
            "paramValue": "agent/agentReconfig.sh",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "executeOnContainer",
            "paramValue": "false",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "rebootInstance",
            "paramValue": "false",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "refreshInstanceInfo",
            "paramValue": "false",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          }
        ],
        "actionResourceMappings": {
          "type": "VIRTUAL_MACHINE",
          "actionResourceFilters": [
            {
              "cloudRegionResource": null,
              "serviceResource": null,
              "applicationProfileResource": null,
              "deploymentResource": null,
              "vmResource": {
                "type": "DEPLOYMENT_VM",
                "appProfiles": ["all"],
                "cloudRegions": ["all"],
                "cloudAccounts": ["all"],
                "services": ["all"],
                "osTypes": [],
                "cloudFamilyNames": [],
                "nodeStates": [],
                "isEditable": true
              },
              "cloudRegionResource": null,
              "serviceResource": null,
              "applicationProfileResource": null,
              "deploymentResource": null,
              "vmResource": {
                "type": "IMPORTED_VM",
                "appProfiles": [],
                "cloudRegions": ["all"],
                "cloudAccounts": ["all"],
                "services": [],
                "osTypes": ["all"],
                "cloudFamilyNames": [],
                "nodeStates": [],
                "cloudResourceMappings": [],
                "isEditable": true
              }
            ]
          ],
          "actionResourceMappingAncillaries": []
        },
        "actionCustomParamSpecs": [
          {
            "paramName": "brokerHost",
            "displayName": "BrokerHost",
            "helpText": "Ip Address or Hostname of Rabbit MQ cluster",
            "type": "string",
            "valueList": null,
            "defaultValue": "",
            "confirmValue": "",
            "pathSuffixValue": "",
            "userVisible": true,
            "userEditable": true,
            "systemParam": false,
            "exampleValue": null,
            "dataUnit": null,
            "optional": false,
            "deploymentParam": false,
            "multiselectSupported": false,
            "useDefault": true,
            "valueConstraint": {
              "minValue": 0,
              "maxValue": 255,
              "maxLength": 255,
              "regex": null,
              "allowSpaces": true,
              "sizeValue": 0,
              "step": 0,
              "calloutWorkflowName": null,
              "scope": null,
              "webserviceListParams": {
                "url": ""
              },
              "protocol": "",
              "username": "",
              "password": "",
              "requestType": null,
              "contentType": null,
              "commandParams": null,
              "requestBody": null,
              "resultString": null,
              "secret": null,
              "tabularTypeData": null,
              "collectionList": []
            },
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "brokerPort",
            "displayName": "BrokerPort",
            "helpText": "RabbitMQ Port number",
            "type": "string",
            "valueList": null,
            "defaultValue": "",
            "confirmValue": "",
            "pathSuffixValue": "",
            "userVisible": true,
            "userEditable": true,
            "systemParam": false,
            "exampleValue": null,
            "dataUnit": null,
            "optional": false,
            "deploymentParam": false,
            "multiselectSupported": false,
            "useDefault": true,
            "valueConstraint": {
              "minValue": 0,
              "maxValue": 255,
              "maxLength": 255,
              "regex": null,
              "allowSpaces": true,
              "sizeValue": 0,
              "step": 0,
              "calloutWorkflowName": null,
              "scope": null,
              "webserviceListParams": {
                "url": ""
              },
              "protocol": "",
              "username": "",
              "password": "",
              "requestType": null,
              "contentType": null,
              "commandParams": null,


```

```

requestBody":null,"resultString":null},"secret":null,"tabularTypeData":null,"collectionList":[],"
preference":"VISIBLE_UNLOCKED"]}],{"id":"58","resource":null,"name":"AgentReConfig_Win","
description":"","actionType":"EXECUTE_COMMAND","category":"ON_DEMAND","lastUpdatedTime":"2019-09-19 22:
15:02.311","timeOut":1200,"enabled":true,"encrypted":false,"explicitShare":false,"
showExplicitShareFeature":false,"deleted":false,"systemDefined":false,"bulkOperationSupported":true,"
isAvailableToUser":true,"currentlyExecuting":false,"owner":1,"actionParameters":[{"paramName":"
downloadFromBundle","paramValue":"true","customParam":false,"required":true,"useDefault":false,"
preference":"VISIBLE_UNLOCKED"},{"paramName":"bundlePath","paramValue":"http://10.0.0.3/5.1-release/ccs-
bundle-artifacts-5.1.0-20190819/agent.zip","customParam":false,"required":true,"useDefault":false,"
preference":"VISIBLE_UNLOCKED"},{"paramName":"script","paramValue":"agent\\agentReconfig.ps1","
customParam":false,"required":true,"useDefault":false,"preference":"VISIBLE_UNLOCKED"},{"paramName":"
executeOnContainer","paramValue":"false","customParam":false,"required":true,"useDefault":false,"
preference":"VISIBLE_UNLOCKED"},{"paramName":"rebootInstance","paramValue":"false","customParam":false,"
required":true,"useDefault":false,"preference":"VISIBLE_UNLOCKED"},{"paramName":"refreshInstanceInfo","
paramValue":"false","customParam":false,"required":true,"useDefault":false,"preference":"
VISIBLE_UNLOCKED"}],"actionResourceMappings":[{"type":"VIRTUAL_MACHINE","actionResourceFilters":
[{"cloudRegionResource":null,"serviceResource":null,"applicationProfileResource":null,"
deploymentResource":null,"vmResource":{"type":"DEPLOYMENT_VM","appProfiles":["all"],"cloudRegions":
["all"],"cloudAccounts":["all"],"services":["all"],"osTypes":[],"cloudFamilyNames":[],"nodeStates":[],"
cloudResourceMappings":[]},"isEditable":true},{"cloudRegionResource":null,"serviceResource":null,"
applicationProfileResource":null,"deploymentResource":null,"vmResource":{"type":"IMPORTED_VM","
appProfiles":["all"],"cloudRegions":["all"],"cloudAccounts":["all"],"services":["all"],"osTypes":["all"],"
cloudFamilyNames":["all"],"nodeStates":["all"],"cloudResourceMappings":[]},"isEditable":true}]}],"
actionResourceMappingAncillaries":[{"paramName":"brokerHost","displayName":"
BrokerHost","helpText":"Ip Address or Hostname of Rabbit MQ cluster","type":"string","valueList":null,"
defaultValue":"","confirmValue":"","pathSuffixValue":"","userVisible":true,"userEditable":true,"
systemParam":false,"exampleValue":null,"dataUnit":null,"optional":false,"deploymentParam":false,"
multiselectSupported":false,"useDefault":true,"valueConstraint":{"minValue":0,"maxValue":255,"maxLength":
255,"regex":null,"allowSpaces":true,"sizeValue":0,"step":0,"calloutWorkflowName":null},"scope":null,"
webserviceListParams":{"url":"","protocol":"","username":"","password":"","requestType":null,"
contentType":null,"commandParams":null,"requestBody":null,"resultString":null},"secret":null,"
tabularTypeData":null,"collectionList":[]},"preference":"VISIBLE_UNLOCKED"},{"paramName":"brokerPort","
displayName":"BrokerPort","helpText":"RabbitMQ Port number","type":"string","valueList":null,"
defaultValue":"","confirmValue":"","pathSuffixValue":"","userVisible":true,"userEditable":true,"
systemParam":false,"exampleValue":null,"dataUnit":null,"optional":false,"deploymentParam":false,"
multiselectSupported":false,"useDefault":true,"valueConstraint":{"minValue":0,"maxValue":255,"maxLength":
255,"regex":null,"allowSpaces":true,"sizeValue":0,"step":0,"calloutWorkflowName":null},"scope":null,"
webserviceListParams":{"url":"","protocol":"","username":"","password":"","requestType":null,"
contentType":null,"commandParams":null,"requestBody":null,"resultString":null},"secret":null,"
tabularTypeData":null,"collectionList":[]},"preference":"VISIBLE_UNLOCKED"}]}],"
repositoriesMappingRequired":false,"actionTypesCounts":[{"key":"EXECUTE_COMMAND","value":"2"}]}

```

3. Access workload manager in your OLD cluster and navigate to the Actions Library page.
4. Import the actions.json file that you saved in Step 2 above. You should see two files (**AgentReconfig_Linux** and **AgentReconfig_Win**) as displayed in the following screenshot.



5. The files are disabled by default (OFF) enable both files by toggling each switch to ON.
6. Save the following script to a file in your local directory and name it **agentReconfig.sh**. This is the file to use for Linux environments.

The agentReconfig.sh file

```

#!/bin/bash

#Write to system log as well as to terminal
logWrite()
{
    msg=$1
    echo "$(date) ${msg}"
    logger -t "OSMOSIX" "${msg}"
    return 0
}

logWrite "Starting agent migrate..."

env_file="/usr/local/osmosix/etc/userenv"

```

```

if [ -f $env_file ];
then
    logWrite "Source the userenv file..."
    . $env_file
fi

if [ -z $brokerHost ];
then
    logWrite "Broker Host / Rabbit Server Ip not passed as action parameter"
    exit 3;
fi

if [ -z $brokerPort ];
then
    logWrite "Broker Port / Rabbit Server Port not passed as action parameter"
    exit 4
fi

replaceUserDataValue() {
    key=$1
    value=$2

    if [ -z $key ] || [ -z $value ];
    then
        logWrite "Command line arguments missing to update user-data file, key: $key, value:$value"
        return
    fi

    user_data_file="/usr/local/agentlite/etc/user-data"
    if [ -f $user_data_file ];
    then
        json_content=`cat $user_data_file`
        old_value=`echo $json_content | awk -F $key '{print $2}' | awk -F \" '{print $3}'`
        sed -i 's@"$old_value"@"$value"@g' $user_data_file
    fi
}

export AGENT_HOME="/usr/local/agentlite"

logWrite "Updating the user data file"
replaceUserDataValue "brokerClusterAddresses" "$brokerHost:$brokerPort"

logWrite "Updating config.json file"
sed -i '/AmqpAddress/c\    "AmqpAddress": "'"$brokerHost":'$"$brokerPort"'"', '$AGENT_HOME/config/config.json'

cd $AGENT_HOME
echo "sleep 10" > execute.sh
echo "/usr/local/agentlite/bin/agent-stop.sh" >> execute.sh
echo "/usr/local/agentlite/bin/agent-start.sh" >> execute.sh
chmod a+x execute.sh
nohup bash execute.sh > /dev/null 2>&1 &

exit 0

```

7. Save the following script to a file in your local directory and name it **agentReconfig.ps1**. This is the file to use for Windows environments.

The agentReconfig.ps1 file

```
param (
    [string]$brokerHost = "$env:brokerHost",
    [string]$brokerPort = "$env:brokerPort"
)

$SERVICE_NAME = "AgentService"
$SYSTEM_DRIVE = (Get-WmiObject Win32_OperatingSystem).SystemDrive
. "$SYSTEM_DRIVE\temp\userenv.ps1"

if ($brokerHost -eq 0 -or $brokerHost -eq $null -or $brokerHost -eq "") {
    echo "Variable brokerHost not available in the env file"
    exit 1
}

if ($brokerPort -eq 0 -or $brokerPort -eq $null -or $brokerPort -eq "") {
    echo "Variable brokerPort not available in the env file"
    exit 2
}

$AGENTGO_PARENT_DIR = "$SYSTEM_DRIVE\opt"

echo "Check if AgentGo Parent directory exists. If not create it: '$AGENTGO_PARENT_DIR'"
if (-not (Test-Path $AGENTGO_PARENT_DIR)) {
    echo "Create $AGENTGO_PARENT_DIR..."
    mkdir $AGENTGO_PARENT_DIR
}
else {
    echo "$AGENTGO_PARENT_DIR already exists."
}

$AGENT_CONFIG="{0}\agentlite\config\config.json" -f $AGENTGO_PARENT_DIR
if (Test-Path $AGENT_CONFIG) {
    echo "Changing the config.json file with the new broker host $env:brokerHost and port $env:
brokerPort"
    $confJson = get-content $AGENT_CONFIG | out-string | convertfrom-json
    $confJson.AmqpAddress = "$($env:brokerHost): $($env:brokerPort)"
    $confJson | ConvertTo-Json | set-content $AGENT_CONFIG
}

$USER_DATA_FILE = "{0}\agentlite\etc\user-data" -f $AGENTGO_PARENT_DIR
if (Test-Path $USER_DATA_FILE) {
    echo "Changing user-data file with new broker host $env:brokerHost and port $env:brokerPort"
    $userDataJson = get-content $USER_DATA_FILE | out-string | convertfrom-json
    $userDataJson.brokerClusterAddresses = "$($env:brokerHost): $($env:brokerPort)"
    $userDataJson | ConvertTo-Json | set-content $USER_DATA_FILE
}

$AGENT_SERVICE_NAME = "AgentService"
echo "Stop-Service $AGENT_SERVICE_NAME" > $AGENTGO_PARENT_DIR\exec.ps1
echo "sleep 10" >> $AGENTGO_PARENT_DIR\exec.ps1
echo "Start-Service $AGENT_SERVICE_NAME" >> $AGENTGO_PARENT_DIR\exec.ps1

echo "Restarting agent"
Start-Process -filepath "powershell" -argumentlist "-executionpolicy bypass -noninteractive -file
`"$AGENTGO_PARENT_DIR\exec.ps1`""

echo "Agent set to restart after config changes"
```


8. Add these two files to a folder called **agent** (just an example) and compress the folder to create **agent.zip** with the same structure displayed here.

agent

agentReconfig.ps1

agentReconfig.sh

- 9. Move the **agent.zip** folder to an HTTP repository in your local environment that is accessible from the OLD and NEW clusters.

 This procedure uses the following URL as an example:
http://10.0.0.3/repo/agent.zip

You have now ensured cluster connectivity and saved the required files for the migration procedure.

f. Migrate Deployments from the OLD Cluster to the NEW Cluster

To migrate the deployment details from the old cluster to the new cluster, follow this procedure.


- 1. Navigate to the workload manager **Actions Libray** page and edit the **AgentReconfig_Linux** action. This procedure continues to use the Linux file
- 2. going forward. Scroll to the **Actions Definition** section and update the URL as displayed in the following screenshot.

Action Definition

* EXECUTE FROM BUNDLE
 YES

* LOCATION * URL
URL http://10.0.0.3/repo/agent.zip




* SCRIPT FROM BUNDLE
agent/agentReconfig.sh

 The URL and Script from Bundle fields in the above screenshot are in accordance with the steps above.

- 3. Scroll to the **Custom Fields** section and change the default value of the **Broker Host** to use the NEW cluster IP.

Custom Fields

If desired add custom fields to the action. They can be made to be user entered or defined here by you, locked and hidden

  BrokerHost 

* DISPLAY NAME
BrokerHost

* PARAMETER NAME
brokerHost

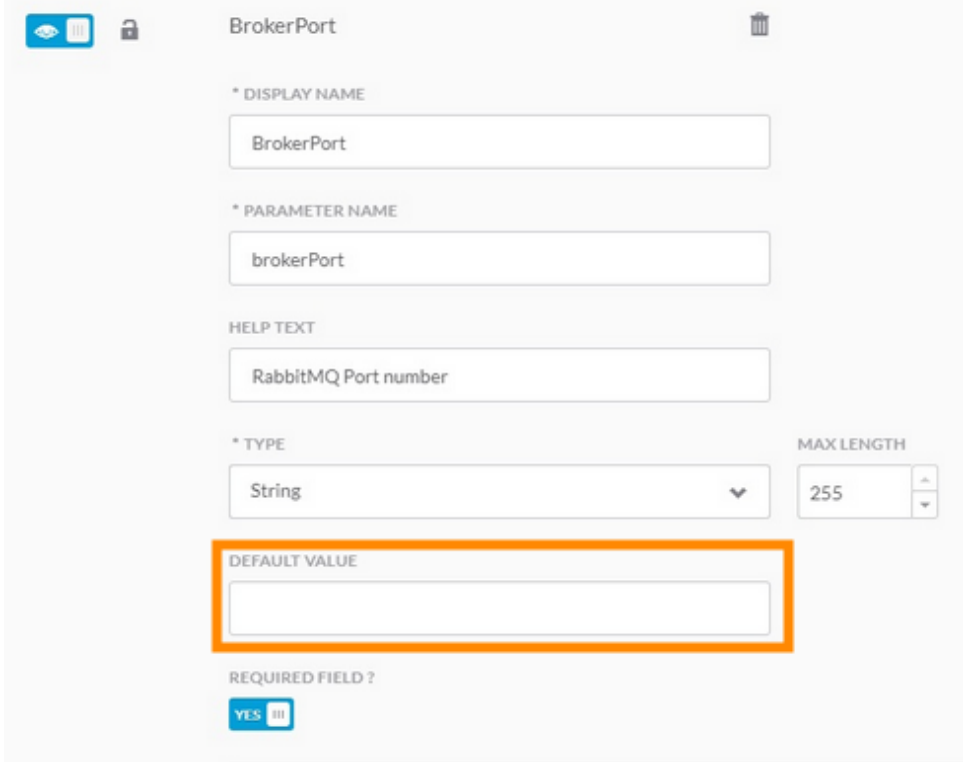
HELP TEXT
Ip Address or Hostname of Rabbit MQ cluster

* TYPE MAX LENGTH
String 255


DEFAULT VALUE

REQUIRED FIELD ?
 YES

- 4. Scroll down to the **Broker Port** and change the default to use the NEW Worker AMQP IP port (for example, 26642 in Step 8 above).



- 5. Click **Done** to save your default configuration changes in the OLD cluster.
- 6. Navigate to the **Virtual Machines** page and locate the VM to migrate to the new cluster.
- 7. Click the **Actions** dropdown and verify if your newly modified actions are visible under the Custom Actions section in the dropdown list as visible in the following screenshot.



- 8. Click one of the actions and verify that the configured defaults are displayed in the Broker host and Broker port fields as indicated earlier.
- 9. Click **Submit** to migrate this VM to the new cluster.
- 10. Verify that the migration is complete by going to the Deployment page in your NEW cluster and the VM is listed as RUNNING (green line).
- 11. Repeat Steps 6 through 10 for each VM that needs to be migrated to the NEW cluster.

You have now migrated the deployment details from the old cluster to the new cluster

Back to: [Public Cloud](#)

Private Cloud

Private Cloud behind Firewalls

- [Overview](#)
- [Minio Server Setup](#)
- [Backup and Restore Process](#)
- [Action Orchestrator-Specific Post-Restore Procedure](#)
- [Post-Restore Procedure](#)
- [Cloud Center Considerations](#)
 - [a. Understand the workload manager Restore Context](#)
 - [b. Retrieve the Port Numbers from the NEW Restored Cluster](#)
 - [c. Retrieve the IP Address of the NEW Restored Cluster](#)
 - [d. Change the IP Address and Port Numbers for the NEW Restored Cluster](#)
 - [e. Perform the Pre-Migrate Activities](#)
 - [f. Migrate Deployments from the OLD Cluster to the NEW Cluster](#)

Overview

You may sometimes need to work in an environment that is completely behind the firewall. This section addresses the backup and restore procedures for those environments.

See [Backup Approach](#) for restrictions and limitations.

Minio Server Setup

You need to set up a Minio server to configure a S3-compatible backup storage location. Refer to <https://min.io/download#/macos> to setup the Minio server.

Once the Minio server is setup, use YOUR Minio server credentials to login to your Minio server.

- Minio server URL
- Minio server username
- Minio server password

To set up a Minio server, use one of the following options:

- Run using Docker:

```
docker run -p 9000:9000 -v /mnt/data:/data minio/minio server /data
```

- Run using Linux binary on any machine:

```
wget https://dl.min.io/server/minio/release/linux-amd64/minio
chmod +x minio
export MINIO_ACCESS_KEY=minio
export MINIO_SECRET_KEY=minio123
./minio server /mnt/data
```

- Run using Windows binary:

```
minio.exe server F:\Data
```

Backup and Restore Process



The script provided as part of this process uses publicly available **Velero 1.5.3** (see <https://velero.io/docs/v1.5/> for details) and **Minio** tools to complete the manual backup and restore process in isolated environments.

To backup and restore the CloudCenter Suite data in an air gap environment, follow this procedure.

1. Create a bucket on the Minio server and provide a meaningful name. This example, uses **velero**. See [Backup Approach](#) for details.
2. Before installing Velero, annotate all the pods in your cluster by using Velero-specific annotations that are provided in the script below.

```
kubectl -n YOUR_POD_NAMESPACE annotate pod/YOUR_POD_NAME backup.velero.io/backup-  
volumes=YOUR_VOLUME_NAME_1,YOUR_VOLUME_NAME_2,...
```

To make the process simpler, here is a utility that does it for you. Be sure to save the following script contents to a file called **pod_vol_restic_scan.py** to your local system.

```
# This utility is used to annotate pods for Velero backups  
import random  
import logging  
import string  
import os  
import time  
import datetime  
from argparse import ArgumentParser  
import sys  
import zipfile  
import shutil  
import subprocess  
import re  
from pprint import pprint as pp  
import yaml  
__copyright__ = "Copyright Cisco Systems"  
__license__ = "Cisco Systems"  
  
def script_run_time(seconds):  
    min, sec = divmod(seconds, 60)  
    hrs, min = divmod(min, 60)  
    timedatastring = "%d:%02d:%02d" % (hrs, min, sec)  
    return timedatastring  
  
def random_char(y):  
    return ''.join(random.choice(string.ascii_letters) for x in range(y))  
  
def border_print(symbol, msg):  
    line = " " + msg + " "  
    totalLength = len(line) + 50  
    logger.info("")  
    logger.info(symbol * totalLength)  
    logger.info(line.center(totalLength, symbol))  
    logger.info(symbol * totalLength)  
    logger.info("")  
  
def setup_custom_logger(name, tcStartTime, fileBaseName, inputName=""):  
    if inputName == "" or inputName == None:  
        st = datetime.datetime.fromtimestamp(  
            tcStartTime).strftime('%Y-%m-%d-%H-%M-%S')  
        filename = fileBaseName + "-" + st + '.log'  
        dirName = "po-scan" + st  
        dirPath = os.path.abspath(os.path.join(  
            os.path.dirname(__file__), '.', dirName))  
        logfilename = os.path.join(dirPath, filename)  
        if not os.path.isdir(dirPath):  
            os.makedirs(dirPath)  
    else:  
        logfilename = inputName  
  
    # print(logfilename)  
    formatter = logging.Formatter(  
        fmt='%(asctime)s %(levelname)-8s %(message)s',  
        datefmt='%Y-%m-%d %H:%M:%S')  
  
    handler = logging.FileHandler(logfilename, mode='w')  
    handler.setFormatter(formatter)
```

```

screen_handler = logging.StreamHandler(stream=sys.stdout)
screen_handler.setFormatter(formatter)
logger = logging.getLogger(name)
logger.setLevel(logging.DEBUG)
logger.addHandler(handler)
logger.addHandler(screen_handler)
return logger, logfilename

def shell_cmd(cmd):
    logger.info("Shell cmd execution >>> '{}'.format(cmd))
    p = subprocess.Popen(
        cmd,
        shell=True,
        stdout=subprocess.PIPE,
        universal_newlines=True)
    output = p.communicate()[0]
    p_status = p.wait()
    return output.split("\n")

def zipdir(path, ziph):
    # ziph is zipfile handle
    for root, dirs, files in os.walk(path):
        for file in files:
            # print(file)
            ziph.write(os.path.join(root, file))

def create_zip():
    st = datetime.datetime.fromtimestamp(
        tcStartTime).strftime('%Y-%m-%d-%H-%M-%S')
    dirName = "ccs-log" + st
    zipFileName = dirName + ".zip"
    zipFilePath = os.path.abspath(os.path.join(os.path.dirname(__file__)))
    logger.info(
        "Generating zip file '{}' at '{}'.format(zipFileName, zipFilePath))
    zipf = zipfile.ZipFile(zipFileName, 'w', zipfile.ZIP_DEFLATED)
    zipdir(dirName, zipf)
    zipf.close()
    shutil.rmtree(dirName)

if __name__ == "__main__":
    fileName = os.path.basename(__file__).split(".")[0]
    tcStartTime = time.time()
    timeStamp = datetime.datetime.fromtimestamp(
        tcStartTime).strftime('%Y%m%d%H%M%S')
    parser = ArgumentParser()
    parser.add_argument(
        "-n", "--namespace", dest="namespace",
        help="Kubernetes Namespace", required=True)
    args = parser.parse_args()
    namespace = args.namespace.strip()
    logger, logFileName = setup_custom_logger(
        "Cloudcenter K8 Debug", tcStartTime, fileName)
    cmd = "kubect1 get pod -n " + namespace + \
        " | grep -v NAME | awk '{print $1}'"
    pod_name_list = shell_cmd(cmd)
    pod_pvc_dict = {}
    pod_vol_dict = {}

    for pod in pod_name_list:
        if pod != "":
            cmd = "kubect1 get pod {} -n {} -o yaml > temp.yaml".format(pod, namespace)
            data = shell_cmd(cmd)
            temp_file = open("temp.yaml", "r")
            with open('temp.yaml', 'r') as temp_file:
                try:
                    file_contents = (yaml.load(temp_file))
                    #print("Pod Name = {}".format(pod.strip()))

```

```

        for vol in file_contents['spec']['volumes']:
            # pp(vol)
            try:
                pvc = vol['persistentVolumeClaim']
                pod_vol_dict[pod.strip()] = vol['name'].strip()
                #print("Vol Name = {}".format(vol['name']))
            except:
                pass
        except yaml.YAMLError as exc:
            logger.error("Error in reading YAML file.")
            logger.error(exc)
        os.remove('temp.yaml')

# pp(pod_vol_dict)
border_print("+", "Applying POD annotations")
for pod in pod_vol_dict.keys():
    if ('elasticsearch-data' in pod) or ('elasticsearch-master' in pod) :
        cmd = "kubectl -n {} annotate pod {} backup.velero.io/backup-volumes-".format(namespace,pod)
        data = shell_cmd(cmd)
        cmd = "kubectl -n {} annotate --overwrite pod {} backup.velero.io/backup-volumes-excludes="
        data = shell_cmd(cmd)
    else:
        cmd = "kubectl -n {} annotate --overwrite pod {} backup.velero.io/backup-volumes={}".format(
namespace,pod,pod_vol_dict[pod])
        data = shell_cmd(cmd)

```

- From where you have saved the `pod_vol_restic_scan.py` script, run the following command **be sure to run this script each time you need a backup!**

```

#Needs Python3
python pod_vol_restic_scan.py -n cisco

```

- Install Velero Version 1.5.3 refer to <https://velero.io/docs/v1.5/> for details.



This is the version used for the client-side CLI commands. You can download from here -<https://github.com/vmware-tanzu/velero/releases/tag/v1.5.3>

- Create a credential file to store your credentials. This example, uses the following URL and credentials *this is only an example!*

Contents of the credentials-minio file

```

[default]
aws_access_key_id = <your Minio username>
aws_secret_access_key = <your Minio password>

```

- On the CloudCenter Suite cluster, you must deploy Velero and configure it with the AWS compatible bucket location, in this example, *Minio*.



Velero and Minio Usage

This process uses Velero to backup the Kubernetes data to a Minio server.

Once you finish this task you can configure the AWS S3 storage provider using the Minio server credentials as specified below. Configuring Minio is similar to configuring an AWS S3 environment, the difference is that you must provide the region and endpoint details when adding the Minio server as AWS S3 storage. You can verify the data from Minio server GUI or command line. The following steps are an example to verify the data from the Minio command line.

Refer to <https://docs.min.io/docs/aws-cli-with-minio.html> for additional details.

- Install Velero manually on the CloudCenter Suite cluster before taking a **backup** of the CloudCenter Suite cluster (assuming kubectl is using kubeconfig of **source/backup** CloudCenter Suite cluster).

- a. Isolated, air gap, environments, that do not have internet access and back up to a local system: Velero images will be pulled from the offline repository.

```
velero install \
  --provider aws \
  --bucket velero \
  --secret-file ./credentials-minio \
  --plugins <offline_repo_url>:8443/velero/velero-plugin-for-aws:v1.1.0 \
  --image <offline_repo_url>:8443/velero/velero:v1.5.3 \
  --use-volume-snapshots=false \
  --backup-location-config region=minio,s3ForcePathStyle="true",s3Url=http://<minio server url>:9000 \
  --use-restic \
  --wait
```

- b. Have internet connectivity and want to back up to a local system: Velero images will be pulled from the online repo.

```
velero install \
  --provider aws \
  --bucket velero \
  --secret-file ./credentials-minio \
  --plugins velero/velero-plugin-for-aws:v1.1.0 \
  --use-volume-snapshots=false \
  --backup-location-config region=minio,s3ForcePathStyle="true",s3Url=http://<minio server url>:9000 \
  --use-restic \
  --wait
```

8. Start a backup using the following command.

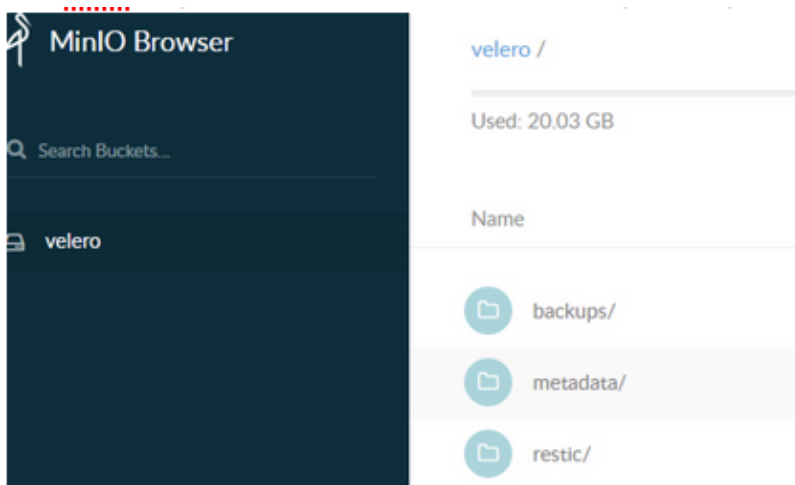
```
velero backup create <Minio backup name> --include-namespaces=cisco --wait
```

Take a backup on source CCS cluster:

- a. Execute `pod_vol_restic_scan.py` to annotate all the pods in your cluster **be sure to run this script each time you need a backup!**
 b. Start a backup using the following command:

```
velero backup create <Minio backup name> --include-namespaces=cisco --wait
```

9. Wait for the backup to complete and watch the logs. Once the backup is complete, the Minio output should look like the information displayed in the following screenshot.



10. Run the Restore Procedure to restore the backup to a different cluster or a fresh cluster.
- Install Velero manually on the CloudCenter Suite cluster before **restoring** the backup on the CloudCenter Suite cluster (assuming kubeclt is using kubeconfig of **destination/restore** CloudCenter Suite cluster).
 - Isolated, air gap, environments, that do not have internet access and back up to a local system: Velero images will be pulled from the offline repository.

```

velero install \
  --provider aws \
  --bucket velero \
  --secret-file ./credentials-minio \
  --plugins <offline_repo_url>:8443/velero/velero-plugin-for-aws:v1.1.0 \
  --image <offline_repo_url>:8443/velero/velero:v1.5.3 \
  --use-volume-snapshots=false \
  --backup-location-config region=minio,s3ForcePathStyle="true",s3Url=http://<minio server url>:9000 \
  --use-restic \
  --wait

```

- ii. Have internet connectivity and want to restore from the local system: Velero images will be pulled from the online repo.

```

velero install \
  --provider aws \
  --bucket velero \
  --secret-file ./credentials-minio \
  --plugins velero/velero-plugin-for-aws:v1.1.0 \
  --use-volume-snapshots=false \
  --backup-location-config region=minio,s3ForcePathStyle="true",s3Url=http://<minio server url>:9000 \
  --use-restic \
  --wait

```

- b. Once the Velero pods are up and running, create the configmap described below to configure the restic to use offline repo for fetching restore-helper image.

```

## Configmap
apiVersion: v1
kind: ConfigMap
metadata:
  name: restic-restore-action-config
  namespace: velero
  labels:
    velero.io/plugin-config: ""
    velero.io/restic: RestoreItemAction
data:
  image: <offline_repo_url>:8443/velero/velero-restic-restore-helper:v1.5.3

```



This step (kubectl create config map) is not applicable if the CloudCenter Suite cluster is online.

```
$ kubectl apply -f /path/to/configmap -n velero
```

- c. Create a backup of the Kubernetes config maps of the following services by executing the script provided on CloudCenter Suite cluster where you are going to perform restore.

- The suite-k8 service
- The suite-prod service

- d. Run the command to execute the backup_configmap.sh script

```

#Execute the script as sudo user
$ sudo /path/to/script/backup_configmap.sh

```

The backup_configmap.sh script

backup_configmap.sh

```
#!/bin/bash

#Scripts to backup ssh keys, proxy settings, k8s and prod-mgmt configmaps on the target cluster
mkdir -p $HOME/backup/configmap
mkdir -p $HOME/backup/service
mkdir -p $HOME/backup/sshkeys
mkdir -p $HOME/backup/proxy

kubectl get svc -n cisco common-framework-nginx-ingress-controller -o json > $HOME/backup/service/ingress.json

for cm in $(kubectl get configmaps -n cisco -o custom-columns=:metadata.name --no-headers=true | grep "k8s-mgmt")
do
    kubectl get configmap $cm -n cisco -o yaml > $HOME/backup/configmap/$cm
done

for cm in $(kubectl get configmaps -n cisco -o custom-columns=:metadata.name --no-headers=true | grep "prod-mgmt")
do
    kubectl get configmap $cm -n cisco -o yaml > $HOME/backup/configmap/$cm
done

kubectl get configmap suite.key -n cisco -o yaml > $HOME/backup/sshkeys/suite.key
kubectl get configmap suite.pub -n cisco -o yaml > $HOME/backup/sshkeys/suite.pub

kubectl get configmap proxy.settings -n cisco -o yaml > $HOME/backup/proxy/proxy.settings

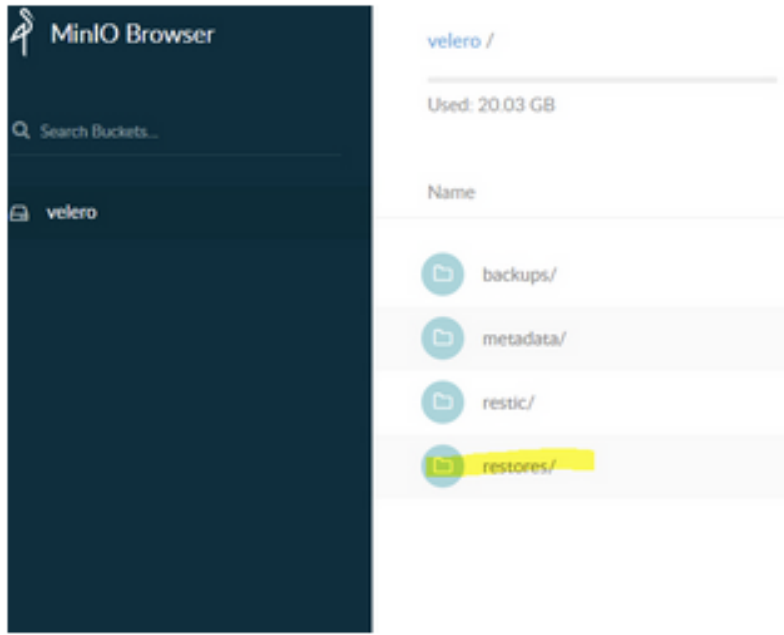
kubectl set env deployment/common-framework-suite-prod-mgmt --list -n cisco | grep "CLOUD_TYPE" >> $HOME/backup/proxy/proxy_variables
kubectl set env deployment/common-framework-suite-prod-mgmt --list -n cisco | grep "HTTP_PROXY" >> $HOME/backup/proxy/proxy_variables
kubectl set env deployment/common-framework-suite-prod-mgmt --list -n cisco | grep "HTTPS_PROXY" >> $HOME/backup/proxy/proxy_variables
kubectl set env deployment/common-framework-suite-prod-mgmt --list -n cisco | grep "NO_PROXY" >> $HOME/backup/proxy/proxy_variables

echo 'Successful!'
```

- e. Start the restore process *after ensuring that the cisco namespace does not exist*.

```
kubectl delete ns cisco
velero restore create --from-backup <Minio backup name>
```

- f. The Minio output should look like the information displayed in the following screenshot you will see an additional restore folder as displayed in the following screenshot



g. At this stage, you must restore the config maps for the following Suite Admin services:

- The suite-k8 service
- The suite-prod service

```
#Execute the script as sudo user
$ sudo /path/to/script/post-restore.sh
```

Without Internet Access - The post-restore.sh script

```
#!/bin/bash
kubect1 delete svc -n cisco common-framework-nginx-ingress-controller
cat $HOME/backup/service/ingress.json | kubect1 create -f -

for cm in $(ls $HOME/backup/configmap)
do
    kubect1 delete configmap $cm -n cisco
done

for cm in $(ls $HOME/backup/configmap)
do
    cat $HOME/backup/configmap/$cm | kubect1 create -f -
done

kubect1 delete configmap suite.key -n cisco
kubect1 delete configmap suite.pub -n cisco
kubect1 delete configmap proxy.settings -n cisco
cat $HOME/backup/sshkeys/suite.key | kubect1 create -f -
cat $HOME/backup/sshkeys/suite.pub | kubect1 create -f -
cat $HOME/backup/proxy/proxy.settings | kubect1 create -f -

while IFS= read -r line; do kubect1 set env deployment/common-framework-suite-prod-mgmt $line -n
cisco; done < $HOME/backup/proxy/proxy_variables

rm -r $HOME/backup/configmap

echo 'Successfull!'
```

You have now restored the Suite Admin data to the new cluster. You can now follow the post-restore procedure specific to workload manager or/and the post-restore procedure specific to Action Orchestrator, as provided in the next section.

Action Orchestrator-Specific Post-Restore Procedure

This section identifies the ArangoDB Backup/Restore Process that is specific to the Action Orchestrator module. If this section is not relevant to your environment, you can skip this section.

1. Ensure the client machine has the ArangoDB client installed. Only the **client** download/install is required. Choose the download appropriate for your operating system: <https://www.arangodb.com/download-major/>
2. After installation, ensure that the tools can be executed:

```
$ arangodump --version
$ arangorestore --version
```

3. Obtain the ArangoDBroot password from the secret.

```
$ kubectl get secrets -n cisco action-orchestrator-pers-arangodb-root-password -o jsonpath={.data.password} | base64 --decode
```

#Output:

```
75e39e601efc0d74d191b53c0a47bca25640acad861b88ff6ae940f172e2c15a
```

4. In a separate terminal window, start a port-forward process to access the arango service from your client.

```
$ kubectl port-forward -n cisco svc/action-orchestrator-pers-arangodb 8529
```

#Output:

```
Forwarding from 127.0.0.1:8529 -> 8529
Forwarding from [::1]:8529 -> 8529
Handling connection for 8529
```

5. Setup environment variables for arangodump/arangorestore commands:

```
export ARANGO_ENDPOINT=http+ssl://localhost:8529
export ARANGO_PWD=75e39e601efc0d74d191b53c0a47bca25640acad861b88ff6ae940f172e2c15a
```

6. Perform the backup:

```
$ arangodump --server.endpoint=$ARANGO_ENDPOINT --server.username=root \
--server.password=$ARANGO_PWD --server.authentication=true \
--all-databases true --threads 8 \
--output-directory $(date "+%Y-%m-%d_%H%M%S")
```

7. Perform the restore.



If the restore is being performed on a separate environment from the backup, ensure that Step 4 has been done in the new client session, and that the variables are appropriate for the new cluster.

8. Ensure the DUMP_FOLDER is replaced with the actual path of the dump.

```
$ arangorestore --server.endpoint=$ARANGO_ENDPOINT --server.username root \
--server.password=$ARANGO_PWD --all-databases true --create-database true \
--replication-factor 3 --threads 4 --overwrite true \
--input-directory {DUMP_FOLDER}
```

7. Log in to arangodb console to verify the cluster is working properly.

workload manager-Specific Post-Restore Procedure



This migration procedure only applies to **Running** deployments.

Be sure to verify that you are only migrating deployment in the **Running** state.



The first few steps differ based on your use of private clouds or public clouds. Be sure to use the procedure applicable to your cloud environment.

Cloud Remote Considerations

Scenario	Cloud Remote Configured	Settings	Notes
1	No	No additional settings	Proceed with the steps provided below, other than the note that only applies to Scenario 3. You must repeat this procedure for each region.
2	Yes	<ol style="list-style-type: none"> 1. Cloud endpoint accessible from CloudCenter Suite = No 2. CloudCenter Suite AMQP reachable from worker VMs = No 3. CloudCenter Suite AMQP accessible from cloud = No 	<p>You do not need to perform any additional configurations and can skip this section.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>i To ensure that the source (old) cluster does not connect to Cloud Remote, click Edit Connectivity in the Regions page and change the settings to Yes for all <i>three</i> settings.</p> </div>
3		<ol style="list-style-type: none"> 1. Cloud endpoint accessible from CloudCenter Suite = No 2. CloudCenter Suite AMQP reachable from worker VMs = No 3. CloudCenter Suite AMQP accessible from cloud = Yes 	<p>Proceed with the steps provided below, INCLUDING the note that is specific to this scenario.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>i If you have multiple deployments that use both Scenario 1 and 3, you must perform these additional steps for deployments that use both Scenarios 1 and 3. You must repeat this procedure for each region.</p> </div> <p>You must repeat this procedure for each region.</p>
4		<ol style="list-style-type: none"> 1. Cloud endpoint accessible from CloudCenter Suite = Yes 2. CloudCenter Suite AMQP reachable from worker VMs = No 3. CloudCenter Suite AMQP accessible from cloud = No 	<p>You do not need to perform any additional configurations and can skip this section (similar to Scenario 2 above).</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>i To ensure that the source (old) cluster does not connect to Cloud Remote, click Edit Connectivity in the Regions page and change the settings to Yes for all <i>three</i> settings.</p> </div>
5		<ol style="list-style-type: none"> 1. Cloud endpoint accessible from CloudCenter Suite = Yes 2. CloudCenter Suite AMQP reachable from worker VMs = No 3. CloudCenter Suite AMQP accessible from cloud = Yes 	<p>Proceed with the steps provided below, INCLUDING the note that is specific to this scenario (similar to Scenario 3 above).</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>i If you have multiple deployments that use both Scenario 1 and 3, you must perform these additional steps for deployments that use both Scenarios 1 and 3. You must repeat this procedure for each region.</p> </div> <p>You must repeat this procedure for each region.</p>

a. Understand the workload manager Restore Context

If you have installed the workload manager module, you must perform this procedure to update the DNS/IP address for the private cloud resources listed below and displayed in the following image:

- The Worker AMQP IP
- The Guacamole Public IP and Port

- The Guacamole IP Address and Port for Application VMs

Cloud endpoint accessible from CloudCenter Suite	Yes
CloudCenter Suite AMQP reachable from worker VM's	Yes
CloudCenter Suite AMQP accessible from cloud	Yes
Remote AMQP IP	
Worker AMQP IP	10.8.1.140:26642
Guacamole Public IP and Port	10.8.1.140:708
Guacamole IP Address and Port for Application VMs	10.8.1.140:32941
Blade Name	cloudcenter-blade-vmware-1-2033



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

b. Retrieve the Port Numbers from the NEW Restored Cluster

The Kubernetes cluster contains the information that is required to update the workload manager UI. This section provides the commands required to retrieve this information.



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

To retrieve the port numbers from the new cluster for private clouds, follow this procedure.

1. The port numbers for each component will differ.
 - a. Run the following command on the new cluster (login to the KubeConfig of the new cluster) to locate the new port numbers for the **Worker AMQP IP**.

```
kubectl get service -n cisco | grep rabbitmq-ext | awk '{print $5}'

# In the resulting response, locate the port corresponding to Port 443 and use that port number!

443:26642/TCP,15672:8902/TCP
```

- b. Run the following command on the new cluster to retrieve the port number for the **Guacamole Public IP and Port**.

```
kubectl get service -n cisco | grep cloudcenter-guacamole | awk '{print $5}'

# In the resulting response, locate the port corresponding to Port 443 and use that port number
for the Guacamole port!

8080:2376/TCP,7788:25226/TCP,7789:32941/TCP,443:708/TCP
```

- c. Run the following command on the new cluster to retrieve the port number for the **Guacamole IP Address and Port for Application VMs**.

```
kubectl get service -n cisco | grep cloudcenter-guacamole | awk '{print $5}'

# In the resulting response, locate the port corresponding to Port 7789 and use that port number
for the Guacamole port!

8080:2376/TCP,7788:25226/TCP,7789:32941/TCP,443:708/TCP
```

c. Retrieve the IP Address of the NEW Restored Cluster

Use the IP address of one of the primary servers of the NEW restored Kubernetes cluster for all the resources where the IP address needs to be replaced.



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

d. Change the IP Address and Port Numbers for the NEW Restored Cluster

The IP addresses and port numbers are not updated automatically in the workload manager UI and you must explicitly update them using this procedure.



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

To configure the IP address and port number in the new cluster, follow this procedure.

1. Access the workload manager module.
2. Navigate to **Clouds > Configure Cloud > Region Connectivity**.

Region Connectivity Running	
Cloud endpoint accessible from CloudCenter Suite	Yes
CloudCenter Suite AMQP reachable from worker VM's	Yes
CloudCenter Suite AMQP accessible from cloud	Yes
Remote AMQP IP	
Worker AMQP IP	10.8.1.140:26642
Guacamole Public IP and Port	10.8.1.140:708
Guacamole IP Address and Port for Application VMs	10.8.1.140:32941
Blade Name	cloudcenter-blade-vmware-1-2033
Strategy	Edit Strategy
Strategy Bundle	

3. Click **Edit Connectivity** in the Region Connectivity settings.
4. In the Configure Region popup, change the 3 fields mentioned above to ensure that the IP and port details are updated to the NEW restored VM.

Configure Region

IS CLOUD END POINT DIRECTLY ACCESSIBLE?
 YES

SHOULD WORKER VMS DIRECTLY CONNECT WITH CLOUDCENTER SUITE?
 YES

WORKER AMQP IP ADDRESS
10.8.1.140:26642

GUACAMOLE PUBLIC IP AND PORT
10.8.1.140:708

GUACAMOLE IP ADDRESS AND PORT FOR APPLICATION VMS
10.8.1.140:32941

OK



DO NOT MAKE ANY OTHER CONFIGURATION CHANGES!

5. Click **OK** to save your changes.



Saving your changes may not automatically update the information in the Region Connectivity settings. Be sure to refresh the page to see the saved information.

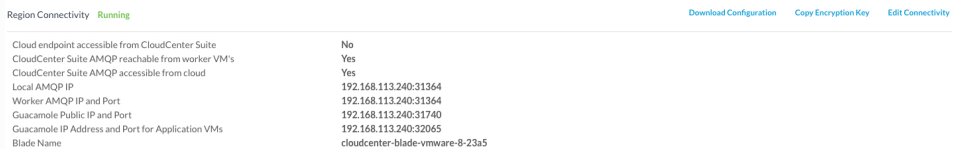
- You have now updated the DNS/IP/Port for the restored WM for this particular cloud. If you have configured other clouds in this environment, be sure to repeat this procedure for each cloud. Once you complete this procedure for all configured clouds, you can resume new deployment activities using the workload manager.

Only for Scenario 3

Only required for Scenario 3 in the Workload Manager table above

With [Cloud Remote](#) configured in your old cluster, you must also reconfigure Cloud Remote to communicate with the new cluster by following this procedure.

- Click **Download Configuration** in the Region Connectivity section as displayed in the following screen shot.



Region Connectivity Running		Download Configuration	Copy Encryption Key	Edit Connectivity
Cloud endpoint accessible from CloudCenter Suite	No			
CloudCenter Suite AMQP reachable from worker VM's	Yes			
CloudCenter Suite AMQP accessible from cloud				
Local AMQP IP	192.168.113.240:31364			
Worker AMQP IP and Port	192.168.113.240:31364			
Guacamole Public IP and Port	192.168.113.240:31740			
Guacamole IP Address and Port for Application VMs	192.168.113.240:32065			
Blade Name	cloudcenter-blade-vmware-8-23a5			

- Click **Copy Encryption Key**.
- Access the Cloud Remote UI.
- Apply the downloaded configuration on the Cloud Remote.

e. Perform the Pre-Migrate Activities

Before you migrate the deployment details you need to ensure that you can connect to both clusters and have the required files to perform the migration.

To perform the pre-migrate activities, follow this procedure.


- Verify that the OLD cluster VMs can reach the NEW cluster. The remaining steps in this procedure are dependent on this connectivity in your environment.
- Save the contents of the following **actions.json** file using the same name and file extension to your local directory with a file type JSON format.

The actions.json file

```
{
  "repositories": [],
  "actions": {
    "resource": null,
    "size": 2,
    "pageNumber": 0,
    "totalElements": 2,
    "totalPages": 1,
    "actionJaxbs": [
      {
        "id": "57",
        "resource": null,
        "name": "AgentReConfig_Linux",
        "description": "",
        "actionType": "EXECUTE_COMMAND",
        "category": "ON_DEMAND",
        "lastUpdatedTime": "2019-09-19 22:14:54.245",
        "timeOut": 1200,
        "enabled": true,
        "encrypted": false,
        "explicitShare": false,
        "showExplicitShareFeature": false,
        "deleted": false,
        "systemDefined": false,
        "bulkOperationSupported": true,
        "isAvailableToUser": true,
        "currentlyExecuting": false,
        "owner": 1,
        "actionParameters": [
          {
            "paramName": "downloadFromBundle",
            "paramValue": "true",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "bundlePath",
            "paramValue": "http://10.0.0.3/5.1-release/ccs-bundle-artifacts-5.1.0-20190819/agent.zip",
            "customParam": false,
            "required": false,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "script",
            "paramValue": "agent/agentReconfig.sh",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "executeOnContainer",
            "paramValue": "false",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "rebootInstance",
            "paramValue": "false",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "refreshInstanceInfo",
            "paramValue": "false",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          }
        ],
        "actionResourceMappings": [
          {
            "type": "VIRTUAL_MACHINE",
            "actionResourceFilters": [
              {
                "cloudRegionResource": null,
                "serviceResource": null,
                "applicationProfileResource": null,
                "deploymentResource": null,
                "vmResource": {
                  "type": "DEPLOYMENT_VM"
                },
                "appProfiles": [
                  "all"
                ],
                "cloudRegions": [
                  "all"
                ],
                "cloudAccounts": [
                  "all"
                ],
                "services": [
                  "all"
                ],
                "osTypes": [],
                "cloudFamilyNames": [],
                "nodeStates": [],
                "cloudResourceMappings": [],
                "isEditable": true
              }
            ],
            "cloudRegionResource": null,
            "serviceResource": null,
            "applicationProfileResource": null,
            "deploymentResource": null,
            "vmResource": {
              "type": "IMPORTED_VM"
            },
            "appProfiles": [],
            "cloudRegions": [
              "all"
            ],
            "cloudAccounts": [
              "all"
            ],
            "services": [],
            "osTypes": [
              "all"
            ],
            "cloudFamilyNames": [],
            "nodeStates": [],
            "cloudResourceMappings": [],
            "isEditable": true
          }
        ],
        "actionResourceMappingAncillaries": [],
        "actionCustomParamSpecs": [
          {
            "paramName": "brokerHost",
            "displayName": "BrokerHost",
            "helpText": "Ip Address or Hostname of Rabbit MQ cluster",
            "type": "string",
            "valueList": null,
            "defaultValue": "",
            "confirmValue": "",
            "pathSuffixValue": "",
            "userVisible": true,
            "userEditable": true,
            "systemParam": false,
            "exampleValue": null,
            "dataUnit": null,
            "optional": false,
            "deploymentParam": false,
            "multiselectSupported": false,
            "useDefault": true,
            "valueConstraint": {
              "minValue": 0,
              "maxValue": 255,
              "maxLength": 255,
              "regex": null,
              "allowSpaces": true,
              "sizeValue": 0,
              "step": 0,
              "calloutWorkflowName": null,
              "scope": null,
              "webserviceListParams": {
                "url": ""
              },
              "protocol": "",
              "username": "",
              "password": "",
              "requestType": null,
              "contentType": null,
              "commandParams": null,
              "requestBody": null,
              "resultString": null,
              "secret": null,
              "tabularTypeData": null,
              "collectionList": [],
              "preference": "VISIBLE_UNLOCKED"
            },
            {
              "paramName": "brokerPort",
              "displayName": "BrokerPort",
              "helpText": ""
            }
          }
        ]
      }
    ]
  }
}
```

```
RabbitMQ Port number", "type": "string", "valueList": null, "defaultValue": "", "confirmValue": "", "
pathSuffixValue": "", "userVisible": true, "userEditable": true, "systemParam": false, "exampleValue": null, "
dataUnit": null, "optional": false, "deploymentParam": false, "multiselectSupported": false, "useDefault": true, "
valueConstraint": {"minValue": 0, "maxValue": 255, "maxLength": 255, "regex": null, "allowSpaces": true, "
sizeValue": 0, "step": 0, "calloutWorkflowName": null}, "scope": null, "webserviceListParams": {"url": "", "
protocol": "", "username": "", "password": "", "requestType": null, "contentType": null, "commandParams": null, "
requestBody": null, "resultString": null}, "secret": null, "tabularTypeData": null, "collectionList": [], "
preference": "VISIBLE_UNLOCKED"}}, {"id": "58", "resource": null, "name": "AgentReConfig_Win", "
description": "", "actionType": "EXECUTE_COMMAND", "category": "ON_DEMAND", "lastUpdatedTime": "2019-09-19 22:
15:02.311", "timeOut": 1200, "enabled": true, "encrypted": false, "explicitShare": false, "
showExplicitShareFeature": false, "deleted": false, "systemDefined": false, "bulkOperationSupported": true, "
isAvailableToUser": true, "currentlyExecuting": false, "owner": 1, "actionParameters": [{"paramName": "
downloadFromBundle", "paramValue": "true", "customParam": false, "required": true, "useDefault": false, "
preference": "VISIBLE_UNLOCKED"}, {"paramName": "bundlePath", "paramValue": "http://10.0.0.3/5.1-release/ccs-
bundle-artifacts-5.1.0-20190819/agent.zip", "customParam": false, "required": true, "useDefault": false, "
preference": "VISIBLE_UNLOCKED"}, {"paramName": "script", "paramValue": "agent\\agentReconfig.ps1", "
customParam": false, "required": true, "useDefault": false, "preference": "VISIBLE_UNLOCKED"}, {"paramName": "
executeOnContainer", "paramValue": "false", "customParam": false, "required": true, "useDefault": false, "
preference": "VISIBLE_UNLOCKED"}, {"paramName": "rebootInstance", "paramValue": "false", "customParam": false, "
required": true, "useDefault": false, "preference": "VISIBLE_UNLOCKED"}, {"paramName": "refreshInstanceInfo", "
paramValue": "false", "customParam": false, "required": true, "useDefault": false, "preference": "
VISIBLE_UNLOCKED"}], "actionResourceMappings": [{"type": "VIRTUAL_MACHINE", "actionResourceFilters":
[{"cloudRegionResource": null, "serviceResource": null, "applicationProfileResource": null, "
deploymentResource": null, "vmResource": {"type": "DEPLOYMENT_VM", "appProfiles": ["all"], "cloudRegions":
["all"], "cloudAccounts": ["all"], "services": ["all"], "osTypes": [], "cloudFamilyNames": [], "nodeStates": [], "
cloudResourceMappings": []}, "isEditable": true}, {"cloudRegionResource": null, "serviceResource": null, "
applicationProfileResource": null, "deploymentResource": null, "vmResource": {"type": "IMPORTED_VM", "
appProfiles": [], "cloudRegions": ["all"], "cloudAccounts": ["all"], "services": [], "osTypes": ["all"], "
cloudFamilyNames": [], "nodeStates": [], "cloudResourceMappings": []}, "isEditable": true}]}], "
actionResourceMappingAncillaries": [], "actionCustomParamSpecs": [{"paramName": "brokerHost", "displayName": "
BrokerHost", "helpText": "Ip Address or Hostname of Rabbit MQ cluster", "type": "string", "valueList": null, "
defaultValue": "", "confirmValue": "", "pathSuffixValue": "", "userVisible": true, "userEditable": true, "
systemParam": false, "exampleValue": null, "dataUnit": null, "optional": false, "deploymentParam": false, "
multiselectSupported": false, "useDefault": true, "valueConstraint": {"minValue": 0, "maxValue": 255, "maxLength":
255, "regex": null, "allowSpaces": true, "sizeValue": 0, "step": 0, "calloutWorkflowName": null}, "scope": null, "
webserviceListParams": {"url": "", "protocol": "", "username": "", "password": "", "requestType": null, "
contentType": null, "commandParams": null, "requestBody": null, "resultString": null}, "secret": null, "
tabularTypeData": null, "collectionList": [], "preference": "VISIBLE_UNLOCKED"}, {"paramName": "brokerPort", "
displayName": "BrokerPort", "helpText": "RabbitMQ Port number", "type": "string", "valueList": null, "
defaultValue": "", "confirmValue": "", "pathSuffixValue": "", "userVisible": true, "userEditable": true, "
systemParam": false, "exampleValue": null, "dataUnit": null, "optional": false, "deploymentParam": false, "
multiselectSupported": false, "useDefault": true, "valueConstraint": {"minValue": 0, "maxValue": 255, "maxLength":
255, "regex": null, "allowSpaces": true, "sizeValue": 0, "step": 0, "calloutWorkflowName": null}, "scope": null, "
webserviceListParams": {"url": "", "protocol": "", "username": "", "password": "", "requestType": null, "
contentType": null, "commandParams": null, "requestBody": null, "resultString": null}, "secret": null, "
tabularTypeData": null, "collectionList": [], "preference": "VISIBLE_UNLOCKED"}]}], "
repositoriesMappingRequired": false, "actionTypesCounts": [{"key": "EXECUTE_COMMAND", "value": "2"}]}
```

3. Access workload manager in your OLD cluster and navigate to the Actions Library page.
4. Import the actions.json file that you saved in Step 2 above. You should see two files (**AgentReconfig_Linux** and **AgentReconfig_Win**) as displayed in the following screenshot.



5. The files are disabled by default (OFF) enable both files by toggling each switch to **ON**.
6. Save the following script to a file in your local directory and name it **agentReconfig.sh**. This is the file to use for Linux environments.

The agentReconfig.sh file

```
#!/bin/bash

#Write to system log as well as to terminal
logWrite()
{
    msg=$1
    echo "$(date) ${msg}"
    logger -t "OSMOSIX" "${msg}"
    return 0
}
```



```

}

logWrite "Starting agent migrate..."

env_file="/usr/local/osmosix/etc/userenv"
if [ -f $env_file ];
then
    logWrite "Source the userenv file..."
    . $env_file
fi

if [ -z $brokerHost ];
then
    logWrite "Broker Host / Rabbit Server Ip not passed as action parameter"
    exit 3;
fi

if [ -z $brokerPort ];
then
    logWrite "Broker Port / Rabbit Server Port not passed as action parameter"
    exit 4
fi

replaceUserDataValue() {
    key=$1
    value=$2

    if [ -z $key ] || [ -z $value ];
    then
        logWrite "Command line arguments missing to update user-data file, key: $key, value:$value"
        return
    fi

    user_data_file="/usr/local/agentlite/etc/user-data"
    if [ -f $user_data_file ];
    then
        json_content=`cat $user_data_file`
        old_value=`echo $json_content | awk -F $key '{print $2}' | awk -F \" '{print $3}'`
        sed -i 's@"$old_value"@'$value"@g' $user_data_file
    fi
}

export AGENT_HOME="/usr/local/agentlite"

logWrite "Updating the user data file"
replaceUserDataValue "brokerClusterAddresses" "$brokerHost:$brokerPort"

logWrite "Updating config.json file"
sed -i '/AmqpAddress/c\    "AmqpAddress": "'${brokerHost}:${brokerPort}'"', ' "$AGENT_HOME/config/config.json"

cd $AGENT_HOME
echo "sleep 10" > execute.sh
echo "/usr/local/agentlite/bin/agent-stop.sh" >> execute.sh
echo "/usr/local/agentlite/bin/agent-start.sh" >> execute.sh
chmod a+x execute.sh
nohup bash execute.sh > /dev/null 2>&1 &

exit 0

```

7. Save the following script to a file in your local directory and name it **agentReconfig.ps1**. This is the file to use for Windows environments.

The agentReconfig.ps1 file

```
param (
    [string]$brokerHost = "$env:brokerHost",
    [string]$brokerPort = "$env:brokerPort"
)

$SERVICE_NAME = "AgentService"
$SYSTEM_DRIVE = (Get-WmiObject Win32_OperatingSystem).SystemDrive
. "$SYSTEM_DRIVE\temp\userenv.ps1"

if ($brokerHost -eq 0 -or $brokerHost -eq $null -or $brokerHost -eq "") {
    echo "Variable brokerHost not available in the env file"
    exit 1
}

if ($brokerPort -eq 0 -or $brokerPort -eq $null -or $brokerPort -eq "") {
    echo "Variable brokerPort not available in the env file"
    exit 2
}

$AGENTGO_PARENT_DIR = "$SYSTEM_DRIVE\opt"

echo "Check if AgentGo Parent directory exists. If not create it: '$AGENTGO_PARENT_DIR'"
if (-not (Test-Path $AGENTGO_PARENT_DIR)) {
    echo "Create $AGENTGO_PARENT_DIR..."
    mkdir $AGENTGO_PARENT_DIR
}
else {
    echo "$AGENTGO_PARENT_DIR already exists."
}

$AGENT_CONFIG="{0}\agentlite\config\config.json" -f $AGENTGO_PARENT_DIR
if (Test-Path $AGENT_CONFIG) {
    echo "Changing the config.json file with the new broker host $env:brokerHost and port $env:
brokerPort"
    $confJson = get-content $AGENT_CONFIG | out-string | convertfrom-json
    $confJson.AmqpAddress = "$($env:brokerHost): $($env:brokerPort)"
    $confJson | ConvertTo-Json | set-content $AGENT_CONFIG
}

$USER_DATA_FILE = "{0}\agentlite\etc\user-data" -f $AGENTGO_PARENT_DIR
if (Test-Path $USER_DATA_FILE) {
    echo "Changing user-data file with new broker host $env:brokerHost and port $env:brokerPort"
    $userDataJson = get-content $USER_DATA_FILE | out-string | convertfrom-json
    $userDataJson.brokerClusterAddresses = "$($env:brokerHost): $($env:brokerPort)"
    $userDataJson | ConvertTo-Json | set-content $USER_DATA_FILE
}

$AGENT_SERVICE_NAME = "AgentService"
echo "Stop-Service $AGENT_SERVICE_NAME" > $AGENTGO_PARENT_DIR\exec.ps1
echo "sleep 10" >> $AGENTGO_PARENT_DIR\exec.ps1
echo "Start-Service $AGENT_SERVICE_NAME" >> $AGENTGO_PARENT_DIR\exec.ps1

echo "Restarting agent"
Start-Process -filepath "powershell" -argumentlist "-executionpolicy bypass -noninteractive -file
`"$AGENTGO_PARENT_DIR\exec.ps1`""

echo "Agent set to restart after config changes"
```


8. Add these two files to a folder called **agent** (just an example) and compress the folder to create **agent.zip** with the same structure displayed here.

agent

agentReconfig.ps1

agentReconfig.sh

9. Move the **agent.zip** folder to an HTTP repository in your local environment that is accessible from the OLD and NEW clusters.

 This procedure uses the following URL as an example:
`http://10.0.0.3/repo/agent.zip`

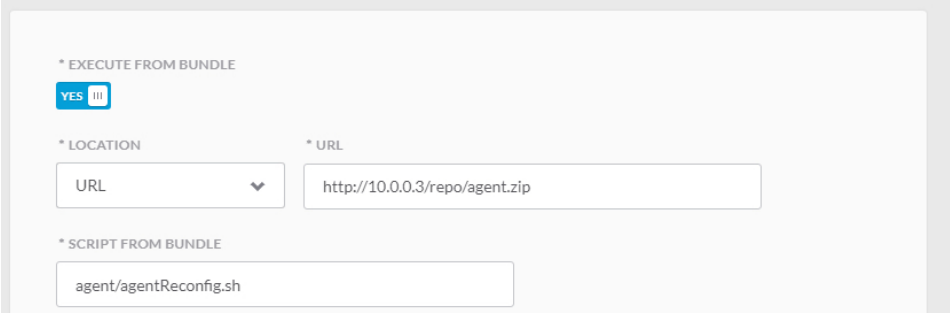
You have now ensured cluster connectivity and saved the required files for the migration procedure.

f. Migrate Deployments from the OLD Cluster to the NEW Cluster

To migrate the deployment details from the old cluster to the new cluster, follow this procedure.

1. Navigate to the workload manager **Actions Library** page and edit the **AgentReconfig_Linux** action. This procedure continues to use the Linux file
2. going forward. Scroll to the **Actions Definition** section and update the URL as displayed in the following screenshot.

Action Definition



* EXECUTE FROM BUNDLE
 YES

* LOCATION * URL

* SCRIPT FROM BUNDLE

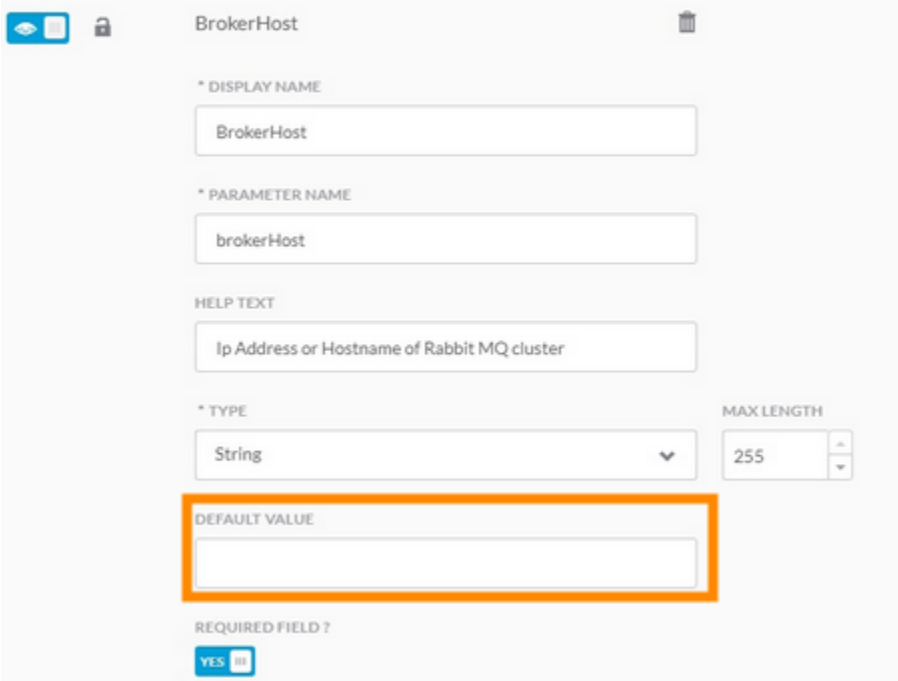
MAX LENGTH

 The URL and Script from Bundle fields in the above screenshot are in accordance with the steps above.

3. Scroll to the **Custom Fields** section and change the default value of the **Broker Host** to use the NEW cluster IP.

Custom Fields

If desired add custom fields to the action. They can be made to be user entered or defined here by you, locked and hidden



BrokerHost

* DISPLAY NAME

* PARAMETER NAME

HELP TEXT

* TYPE MAX LENGTH

DEFAULT VALUE

REQUIRED FIELD?

4. Scroll down to the **Broker Port** and change the default to use the NEW Worker AMQP IP port (for example, 26642 in Step 8 above).

BrokerPort

* DISPLAY NAME
BrokerPort

* PARAMETER NAME
brokerPort

HELP TEXT
RabbitMQ Port number

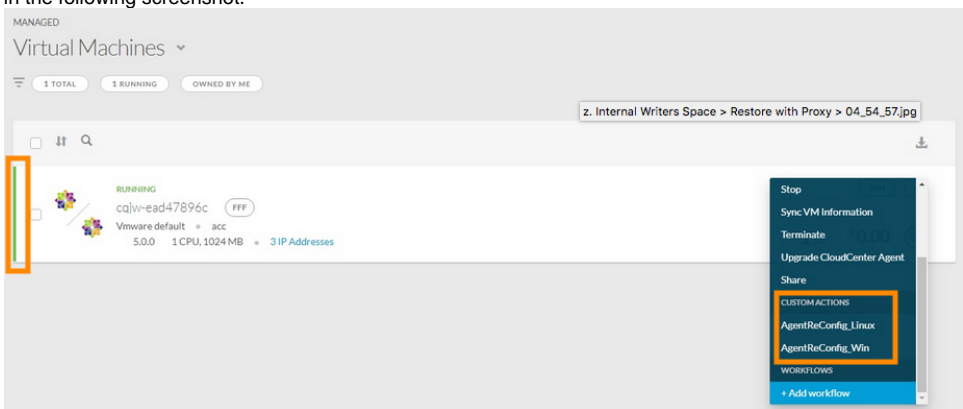
* TYPE
String

MAX LENGTH
255

DEFAULT VALUE

REQUIRED FIELD ?
YES

5. Click **Done** to save your default configuration changes in the OLD cluster.
6. Navigate to the **Virtual Machines** page and locate the VM to migrate to the new cluster.
7. Click the **Actions** dropdown and verify if your newly modified actions are visible under the Custom Actions section in the dropdown list as visible in the following screenshot.



8. Click one of the actions and verify that the configured defaults are displayed in the Broker host and Broker port fields as indicated earlier.
9. Click **Submit** to migrate this VM to the new cluster.
10. Verify that the migration is complete by going to the Deployment page in your NEW cluster and the VM is listed as RUNNING (green line).
11. Repeat Steps 6 through 10 for each VM that needs to be migrated to the NEW cluster.

You have now migrated the deployment details from the old cluster to the new cluster

You have now backed up and restored the CloudCenter Suite to an isolated environment using the

Minio server.

Troubleshooting

Troubleshooting

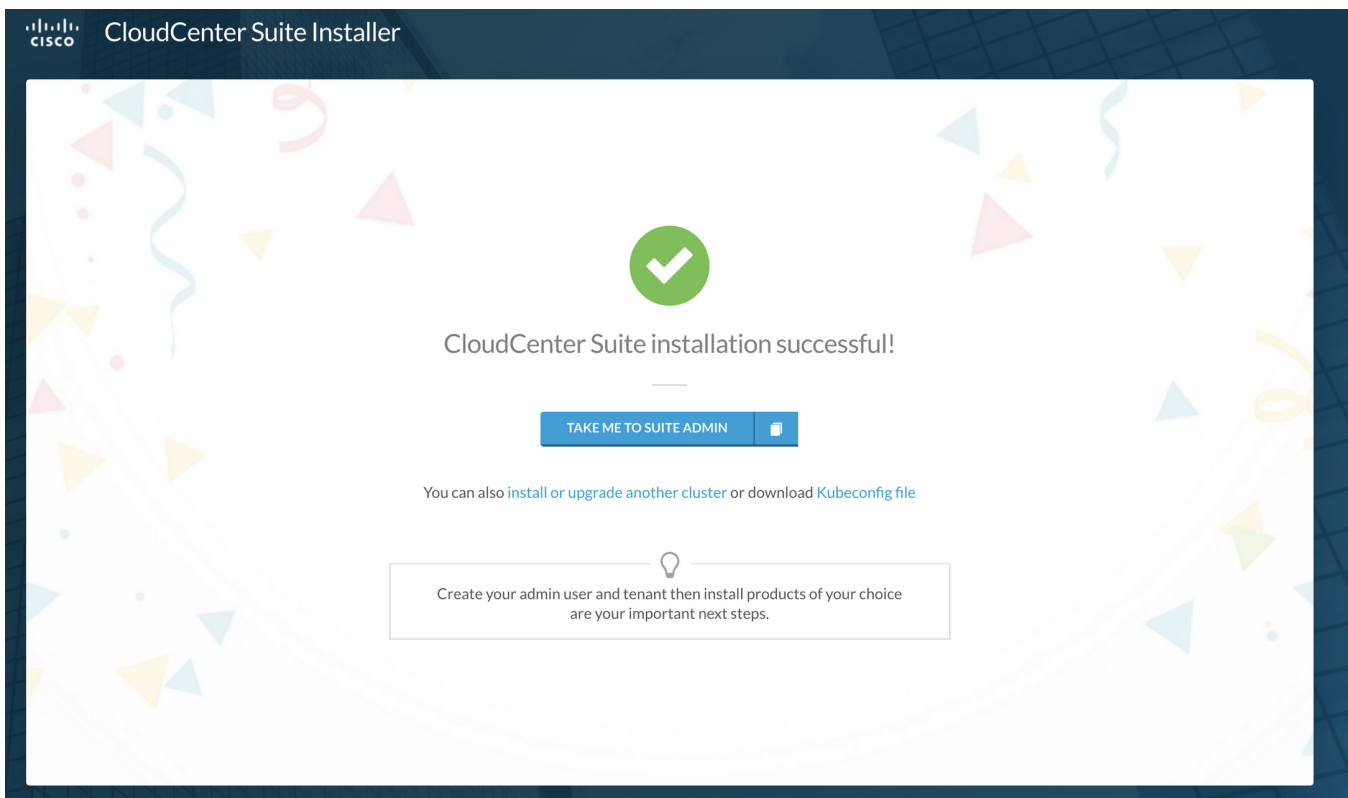
- [Overview](#)
- [Finding Kubernetes Resources](#)
- [Error during the Suite Installation Process](#)
- [The Kubernetes Cluster is installed successfully, but the progress bar for Suite Administration is stuck at Waiting for product to be ready](#)
- [After using Suite Admin for a while, users cannot login to Suite Admin if any of the cluster nodes are in a Not Ready state](#)
- [Download Logs](#)
- [Velero Issues](#)
- [vSphere Environments](#)
 - [A Pod has unbound PersistentVolumeClaims](#)
 - [The Progress bar for a Kubernetes Cluster is stuck at Launching cluster nodes on the cloud or Configuring the primary server cluster](#)
 - [Installation Failed: Failed to copy <script-name.sh> to remote host or any error related to SSH connection failure](#)
 - [When one of the workers is down a worker node scale up operation is stuck](#)

Overview

If you encounter issues during the installation process, be sure to review the tips provided in this page before calling the support team.

Finding Kubernetes Resources

For private clouds, the download link for the **Kubeconfig** file is available on the last page of the installer UI as displayed in the following screenshot.



While you may see this file for successful installations in the above screen, you will not be able to access this file if your installation was not successful. This file is required to issue any command listed in the <https://kubernetes.io/docs/reference/kubectl/cheatsheet/> section of the Kubernetes documentation.

By default, the **kubectl** command looks for the Kubeconfig file in the **\$HOME/.kube** folder.

- **Successful installation:** Copy the downloaded Kubeconfig file to your **\$HOME/.kube** folder and then issue any of the kubectl commands listed in the Kubernetes cheatsheet link above.
- **Stalled Installation:**
 - Private clouds and most public clouds: SSH into one of the primary server nodes and copy the Kubeconfig file from **/etc/kubernetes/admin.conf** to the **/root/.kube** folder.
 - GCP: Login to GCP, access the Kubernetes Engine, locate your cluster, click **Connect** to connect to the cluster, and click the **copy** icon as displayed in the following screenshot. You should have already installed gcloud in order to view this icon.

Connect to the cluster

You can connect to your cluster via command-line or using a dashboard.

Command-line access

Configure [kubectl](#) command line access by running the following command:

```
$ gcloud container clusters get-credentials pujanrc221-7220e62e-ca6f-4f08-963c-9e49b --zone us-east1-b --project
```

[Run in Cloud Shell](#)

Cloud Console dashboard

You can view the workloads running in your cluster in the Cloud Console [Workloads dashboard](#).

[Open Workloads dashboard](#)

OK

Error during the Suite Installation Process

At any time, if you your installation stalls due to a lack of resources, perform this procedure to analyze the error logs.

To fetch the logs for this pod run :

1. Locate the actual name of the container by running the following command:

```
kubectl get pods -all-namespaces | grep common-framework-suite-prod-mgmt-xxxx
```

2. Click the [Download Logs](#) link to download the installation logs for the failed service in case of an installation failure.
3. View the Logs for the container: common-framework-suite-prod-mgmt ...
4. Run the following command to view the error:

```
kubectl logs -f common-framework-suite-prod-mgmt-xxxx -n cisco
```

The Kubernetes Cluster is installed successfully, but the progress bar for Suite Administration is stuck at *Waiting for product to be ready*

This issue indicates that the CloudCenter Suite installation has some issue. SSH into one of the primary server nodes using the private key. To check the status of the pods, run `kubectl get pods --all-namespaces` for each pod. If the status does not display **Running**, run the following commands to debug further:

```
kubectl describe pod <pod-name> -n cisco
```

or

```
kubectl logs -f <pod-name> -n cisco
```

To SSH into each cluster node, SSH into the node using the private key and check if the system clock is synchronized on all nodes. Even if the NTP servers were initially synchronized verify if they are still active by using the following command.

```
ntpdate <ntp_server>
```

or

You may have provided the wrong proxy details at installation time test if the proxy is working on the installer VM and ensure that the repository is accessible.

or

Verify the offline CloudCenter Suite cluster to ensure that the installer is able to pull the image from the offline repository. Alternately, manually pull the images from the offline repo and verify if it works

After using Suite Adminfor a while, users cannot login to Suite Adminif any of the cluster nodes are in a *Not Ready* state

This issue may be the result of any of the following situations:

- Are all the cluster nodes up and running with a valid IP address?
- If the nodes are running, then SSH into one of the primary server nodes using the private key.
- Run the following command on the primary server to verify if all the nodes are in the **Ready** state.

```
kubect1 get nodes
```

Download Logs

Click the **Download Logs Download** link to download the installation logs for the failed service incase of an installation failure. See [Monitor Modules > Download Logs](#) for additional information.

Velero Issues


Refer to <https://heptio.github.io/velero/v0.11.0/> for Velero troubleshooting information.

vSphere Environments

The following issues are specific to vSphere environments.


A Pod has unbound PersistentVolumeClaims

The problem displayed in the following screenshot is usually caused when the cloud user does not have permissions to the configured storage. For example, a vSphere user may not have permissions to the selected datastore.



The Progress bar for a Kubernetes Cluster is stuck at *Launching cluster nodes on the cloud* or *Configuring the primary server cluster*

The issue displayed in the following screenshot could be an issue with the cloud environment. Refer to your cloud documentation for possible issues.



Other examples:

- If the target cloud is vSphere, check if the cloud account being used has permissions to launch a VM and if the VM is configured with a valid IPv4 address.
- If the cluster nodes are configured to use static IP, verify if the IP pool used is valid and if all the launched nodes have a unique IP from the pool.

Installation Failed: Failed to copy <script-name.sh> to remote host or any error related to SSH connection failure

If any of the nodes are **Not Ready** state, then run the following command on the node:

```
kubectl describe node <node-name>
```

This issue can occur when the installer node cannot SSH/SCP into launched cluster nodes. Verify if all the launched nodes have a valid IPv4 address and if the installer network can communicate with the Kubernetes cluster network (if they are on different networks). Also verify that the cluster nodes are able to connect to vSphere.

If none of the above methods work, retry the installation or contact your CloudCenter Suite admin.

When one of the workers is down a worker node scale up operation is stuck

When one of the workers is down, and you try to scale up the worker node, the node does not scaled up. The scale up operation remains stuck in scaling state.

Restart the operator POD of your environment by using the following command. The following example displays vSphere, and the corresponding operator will be the vSphere operator. Similarly, if you are working in an OpenStack environment, use the OpenStack operator as applicable.


```
kubectl delete pod kaas-ccp-vsphere-operator-<dynamic alphanumeric characters> -n ccp
```

```
#or
```

```
kubectl delete pod kaas-ccp-openstack-operator-<dynamic alphanumeric characters> -n ccp
```

By restarting this service on any worker node, you will start the shutdown VM and scale up the new node which was stuck during the scale operation.

Suite Admin Workflow

Suite Admin

The following table identifies the tasks to be performed on the Suite Admin once you install the CloudCenter Suite.

#	Required?	Goal	Task	Description
1	Yes	Onboarding	Create the suite administrator and root tenant.	See Initial Administrator Setup
			Navigate to the Suite Admin Dashboard.	See Suite Admin Dashboard
2	No	Language selection	Select your language choices to view the CloudCenter Suite UI.	See UI Language Availability
3	Yes	Module installation	Install module(s) of choice based on the list available in the Dashboard. This is optional, however, you cannot configure resources other than users/tenants/groups/roles /admin menu settings if you don't install modules!	See Install Module
4	Yes	User management	Create users	See Create and Manage Users
5	Yes	Group Management	Assign users to default groups. When the suite administrator installs any module, additional, default out-of-box groups become available. These groups vary based on the module.	See Create and Assign Groups
	Optional		Create a custom group If the out-of-box groups don't meet your requirements, you can create custom groups.	See Custom Groups by Admin
	Yes		Assign roles to a group For each custom group, you must assign at least one role.	See Understand Roles
6	Yes	Admin Management	Set up the base URL	See Base URL Configuration
	Yes		Set up email communication	See Email Settings
	Optional		Configure a dedicated alias hostname and use an external IdP to authenticate its users.	See SSO Setup
	Optional		Set up the proxy server	See Proxy Settings
7	Yes	Product Registration	Configure a license	See Configure Smart Licenses
8	Optional	Cluster Management	Modify the size of the cluster	See Manage Clusters
9	Optional	Troubleshooting	<ul style="list-style-type: none"> View log archives Download logs for troubleshooting purposes 	<ul style="list-style-type: none"> See Log Archive See Monitor Modules
10	Optional	Tenant/Sub-tenant Management	Manage your own tenant or create additional sub-tenants	See Manage Tenants
			Add users as additional tenant administrators to a group	See Create and Assign Groups
11	Optional	Admin Management	Backup CloudCenter Suite	See Backup
			Restore CloudCenter Suite	See Restore
			Setup Isolated (Air Gap) environment	See Without Internet Access

Initial Administrator Setup

InitialAdministratorSetup

- [Overview](#)
- [The Suite Administrator](#)
- [Configure an Admin User and Tenant](#)

Overview

Once the Suite Admin is installed you must perform the following tasks:

- Note or bookmark the IP address for the Suite Admin console.
- Set up the credentials for the Suite administrator.
- Configure a Root tenant.

The Suite Administrator

As the administrator for the Suite Admin, you can perform the following tasks from the Suite Admin dashboard:

- [Install Module\(s\)](#)
- [Create and Manage Users](#), including tenants and tenant administrators
- [Create and Assign Groups](#), including user-group(s) association
- [Configure Smart Licenses](#)
- [Manage Clusters](#), if the cluster was created by the suite administrator

Configure an Admin User and Tenant

To configure the admin user and tenant, follow this procedure:

1. Navigate to the Suite Admin console and complete the **Admin User and Tenant Credentials** form to enter details for the root user and tenant as displayed in the following screenshot.



The screenshot shows a login interface for Cisco CloudCenter Suite. On the left is a blue sidebar with the Cisco logo and the text 'WELCOME TO CloudCenter Suite'. The main content area is white and contains the following elements:

- Text: "Please log in to your account to proceed."
- Form field: "* EMAIL ADDRESS" with the value "joe.smith@cisco.com".
- Form field: "* PASSWORD" with the value "password" and a toggle icon.
- Form field: "* TENANT ID" with the value "tenant ID".
- Button: A blue "LOGIN" button.
- Footer: "Language: English" with a dropdown arrow, and a blue link "FORGOT PASSWORD?".


2. Besides the First and Last Name, Email Address, Password, Company Name, and Company Logo (defaults to the Cisco logo), you must enter a Tenant ID of your choice so you can log into the Suite Admin using this Tenant ID and password.
3. Click **Done** to save your settings and launch the Suite Admin Dashboard as displayed in the following screenshot.




Modules



Suite Admin
v5.1.0 • Installed: Aug 07, 2019



Workload Manager
v5.1.0 • Installed: Aug 07, 2019



Cost Optimizer
v5.1.0 • Installed: Aug 07, 2019



Action Orchestrator
v5.1.0 • Installed: Aug 07, 2019

Kubernetes Cluster Management

Kubernetes Cluster Management

- [Cluster Status](#)
- [Manage Clusters](#)

Cluster Status

Cluster Status

- [Overview](#)
- [Requirements](#)
- [The Cloud Icon Details](#)
- [Kubernetes Cluster Actions](#)
- [Modify Cluster Size](#)
- [Virtual Machines](#)

Overview


You can view the status of a Kubernetes cluster by clicking the *cloud* icon located in the header of the [Suite Admin Dashboard](#). The Cluster status popup displays. Click **View Details** to view detailed information about each node in the cluster.

Requirements

For private clouds, the HA cluster requires a minimum of 2 out of 3 master nodes to be running at any point, for the cluster to function as designed.

The Cloud Icon Details

Click the *cloud* icon to view and verify the number of nodes in the Kubernetes cluster. The **View Details** page displays detailed information about each node in the cluster. This information is retrieved from the Kubernetes cluster after you install the CloudCenter Suiteccs. The following screenshot displays details within this page.



Kubernetes Cluster Actions

The cluster-level actions allow you to download the following files.

- The SSH key file is used to connect to the cluster.
- The KubeConfig file is used to view cluster information.

Modify Cluster Size

Based on your environment requirements, you can modify the Kubernetes cluster size from the Suite Admin. See [Manage Clusters](#) for additional details.

Virtual Machines

This tab displays the VMs that make up the Kubernetes cluster accessed from this instance of CloudCenter Suite.

The colored status indicators identify the state of each VM in your Kubernetes cluster as described in the following table.

Cluster Status Color	Indication
Green	The node is functioning.
Red	The node is not functioning.

The color merely indicates the health of your Kubernetes cluster so you can make the required changes to your Kubernetes setup as required by your environment.

Manage Clusters


Manage Clusters

- [Overview](#)
- [Scale Up](#)
- [Scale Down](#)
- [Reconfigure Cloud Credentials](#)

Overview


If a cluster was created by the suite administrator as described in [Initial Administrator Setup](#), then this suite administrator can manage those clusters. Managing a cluster includes the following tasks.

- Scale this cluster.
- Monitor the cluster by viewing alerts.

 Suite administrators can only manage clusters that they installed.

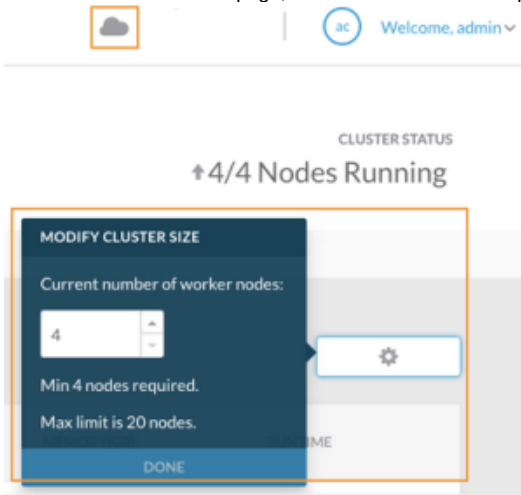
The suite administrator's ability to view a cluster is indicated by the green circle on the **cloud icon**. Clicking this icon provides additional information as displayed in the following screenshot.

Scale Up

 If you setup the CloudCenter Suite using static IPs, verify that the static IP range has free IPs available to support scale up operations. If IPs are not available in the static IP range (defined during installation) then the scale up process will not take place.

To increase the number of nodes in your cluster, perform this procedure.

1. Navigate to the [Suite Admin Dashboard](#) > [Tenants](#) page.
2. Click the [cloud icon](#) to access the [Cluster Status](#) > [View Details](#) page.
3. In the Kubernetes Cluster page, click the wheel icon to display the **Modify Cluster Size** popup as displayed in the following screenshot.



4. Increase the number as required in the **Current number of worker nodes**: field. You will see the status bar list a *Scaling operation successful* alert. It take a few minutes to increase the node count.
 - Initially, the node will be in the red state while it is still initializing. Once it has initialized, it will turn green.
 - The Runtime displays the length of time that this node has been running:
 - h = Upto 24 hours
 - d = Any number of days
 - The Status can only be up (red) or down (green).
 - The memory and CPU details are displayed as available in the Kubernetes cluster.
 - When complete, you see a subsequent alert notifying you of the Cluster node being added.

You have now increased the number of nodes in your cluster.

Scale Down

While you can scale up the number of nodes in the Kubernetes cluster from the Suite Admin, you *cannot* scale down using this process.

Reconfigure Cloud Credentials

OpenStack

If you installed CloudCenter Suite 5.1.1 as a fresh installation, this feature is not available in OpenStack environments.

If you upgraded CloudCenter Suite from 5.0.x to 5.1.0 or 5.1.1, the Cloud Account section is preserved and you can update the password.

vSphere

If you have updated your password in the vSphere console, be sure to update it in the Cloud Accounts tab (in the Kubernetes Cluster page), before the vSphere lockout period takes effect.

If you do not update the password, be aware that the vSphere policy will prevent you from proceeding with your CloudCenter Suite configuration and CloudCenter Suite will continue with its polling attempts with vSphere.

The Cloud Accounts tab, provides a way to change your cloud credentials for the [cloud where the CloudCenter Suite is installed](#).

You can change your cloud account password based on your cloud credentials for each supported cloud as listed in [New Cluster Installation](#).

Configure Smart Licenses

Configure Smart Licenses

- [Overview](#)
- [Cisco Smart Software Manager](#)
 - [Virtual Accounts](#)
 - [Smart Call Home](#)
- [Configuring Cisco Smart Software Licensing](#)
 - [Request a Smart Account](#)
 - [Adding Users to a Smart Account](#)
- [License Usage and Compliance](#)
- [Workflow of Cisco Smart Software Licensing](#)
 - [Generating a Registration Token](#)
 - [Configuring Transport Settings](#)
 - [Registering a CloudCenter Suite License](#)
 - [Renewing Authorization](#)
 - [Re-Registering a CloudCenter Suite License](#)
 - [De-Registering a CloudCenter Suite License](#)
- [Enable for Production](#)
- [Troubleshooting Licensing Issues](#)
 - [Invalid Token](#)
 - [Download Logs](#)

Overview

CloudCenter Suite integrates with the [Cisco Smart Software Licensing](#) solution. The CloudCenter Suite is available for a 90-day evaluation period after which, you must register with Cisco Smart Software Manager.

The number of licenses required depends on your deployment scenario. For example, the Workload Manager and Cost Optimizer define entitlements based on features used in those modules. These entitlements may apply to the use of a specific public/private cloud, the number of management units used when deploying applications (VMs and containers), the options purchased (essentials, advanced, premium), and so forth.



Smart licenses are already incorporated in the CloudCenter Suite SaaS offer, see [SaaS Access](#) for additional details.

Cisco Smart Software Manager

Cisco Smart Software Manager (Cisco SSM) enables the management of software licenses and Smart Account from a single portal. This interface allows you to activate your product, manage entitlements, renew and upgrade software. You must have a functioning Smart Account to complete the registration process and will need to exchange three key elements with the Cisco Smart Software Manager over HTTPS:

- **Trusted Unique Identifier** This is the Product ID (SUDI/SUVI/ID).
- **Organizational Identifier** In a numerical format to associate product with a Smart / Virtual Account.
- **Licenses consumed** Allows the Cisco Smart Software Manager to understand the license type and level of consumption.

Virtual Accounts

A Smart Account provides a single location for all Smart enabled products and entitlements. It assists to speed procurement, deployment and maintenance of Cisco Software. When creating a Smart Account the submitter must have the authority to represent the requesting organization. After submitting the request goes through a brief approval.

A Virtual Account exists as a sub-account within the Smart Account. Virtual Accounts are a customer defined structure based on organizational layout, business function, geography or any defined hierarchy. They are created and maintained by the Smart Account administrator(s).

Smart Call Home

Smart Call Home is feature to communicate with the Cisco Smart Software Manager. By default, Smart Call Home is enabled when you configure Smart Software Licensing. Smart Call Home creates a Cisco TAC-1 profile and sends associated Smart Call Home messages after the enablement. For platforms with Smart Software Licensing enabled by default, call-home is also enabled by default with associated messages.

Configuring Cisco Smart Software Licensing

You need to configure Cisco Smart Software Licensing to easily procure, deploy, and manage licenses for your CloudCenter Suite.

Smart Licensing is a cloud-based approach to licensing. The solution simplifies the purchase, deployment and management of Cisco software assets. Entitlements are purchased through your Cisco account via Cisco Commerce Workspace (CCW) and immediately deposited into a *Virtual Account* for usage. This process eliminates the need to install license files on every device using the product. Products that are smart enabled communicate directly to Cisco to report consumption. A single location is available to customers to manage Cisco software licenses the Cisco SSM. License ownership and consumption are readily available to help make better purchase decision based on consumption or business need.

Cisco SSM enables you to manage your Cisco Smart Software Licenses from one centralized website. With Cisco SSM, you can organize and view your licenses into *Virtual Account* groups. You can also use Cisco SSM to transfer licenses between virtual accounts as needed. You can access Cisco SSM from the Cisco Software Central homepage at software.cisco.com, under Smart Licensing.

If you do not want to manage licenses using Cisco SSM, either for policy reasons or network availability reasons, you can choose to install Cisco SSM Satellite at your premises. CloudCenter Suite registers and reports license consumption to the Cisco SSM Satellite as it does to Cisco SSM. Cisco SSM Satellite coordinates with the Cisco Smart Software Manager to manage software licenses on premises. Devices register locally to report license ownership and consumption.



Ensure that you use Cisco SSM Satellite version 5.0 or later. For more information on installing and configuring Cisco SSM Satellite, refer to <http://www.cisco.com/go/smartsatellite>.

Request a Smart Account

The creation of a new Smart Account is a one-time event and subsequent management of users is a capability provided through the tool. To request a Smart Account, visit software.cisco.com and follow this process.

1. After logging in, select **Request a Smart Account** in the Administration section as displayed in the following screenshot.



Administration

Request a Smart Account

Get a Smart Account for your organization.

Request a Partner Holding Account

Allows Cisco Partners to request a Holding Smart Account

Manage Smart Account

Modify the properties of your Smart Accounts and associate individual Cisco Accounts with Smart Accounts.

Learn about Smart Accounts

Access documentation and training.

2. Select the type of Smart Account to create using one of two options as displayed in the following screenshot.

Create Account

Would you like to create the Smart Account now?

- Yes, I have authority to represent my company and want to create the Smart Account.
- No, the person specified below will create the account:

* Email Address:

Message to Creator:

- Individual Smart Account requiring agreement to represent your company. By creating this Smart Account you agree to authorize, create, and manage product and service entitlements, users, and roles on behalf of your organization.
- Create the account on someone else's behalf

3. Provide the required domain identifier and the preferred account name as displayed in the following screenshot.

Account Information

The Account Domain Identifier will be used to **uniquely identify the account**. It is based on the email address of the person creating the account by default and must belong to the company that will own this account. [Learn More](#)

* Account Domain Identifier:	domainidentifier.com Edit
* Account Name:	<input type="text" value="Account Name"/>

4. The account request requires approval for the Account Domain Identifier as displayed in the following screenshot. An email will be sent to the requester to complete the setup process.

i Smart Account Request Pending
The account setup process is pending approval of an Account Domain Identifier. You will receive an email confirmation and a Cisco representative will contact you at the number provided below.

Adding Users to a Smart Account

Smart Account user management is available in the Administration section of software.cisco.com. To add a new user to a Smart Account, follow this process.

1. After logging in, select **Manage Smart Account** in the Administration section as displayed in the following screenshot.



Administration

[Request a Smart Account](#)

Get a Smart Account for your organization.

[Request a Partner Holding Account](#)

Allows Cisco Partners to request a Holding Smart Account

[Manage Smart Account](#)

Modify the properties of your Smart Accounts and associate individual Cisco Accounts with Smart Accounts.

[Learn about Smart Accounts](#)

Access documentation and training.

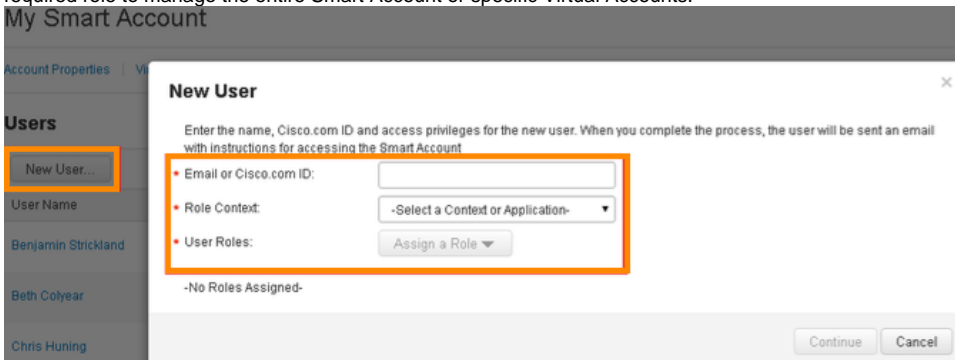
2. Select the **Users** tab as displayed in the following screenshot.

[Cisco Software Central](#) > **Manage Smart Account**

My Smart Account

[Account Properties](#) | [Virtual Accounts](#) | **[Users](#)** | [Account Agreements](#)

3. Select **New User** and provide the required email address, cisco.com ID, and role as displayed in the following screenshot. You can select the required role to manage the entire Smart Account or specific Virtual Accounts.



4. Click **Continue** to complete the process.

License Usage and Compliance

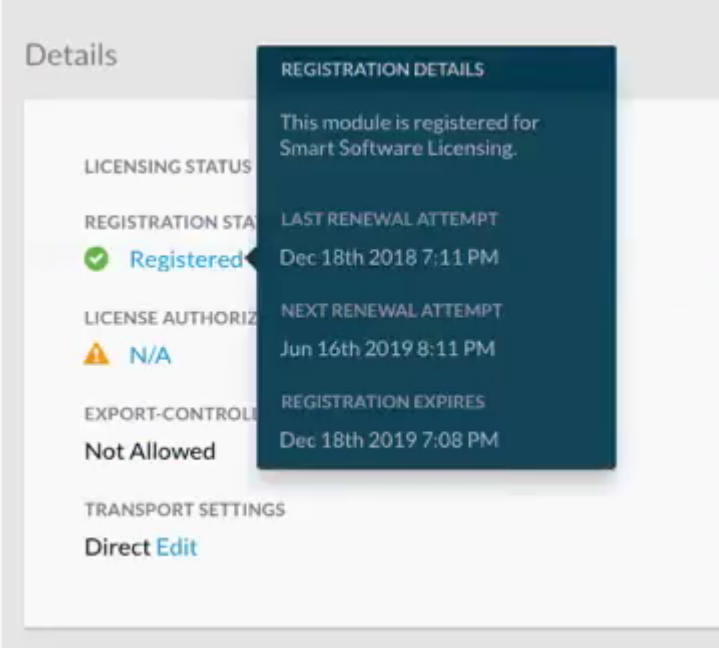
Once you register CloudCenter Suite with Cisco SSM, you will receive the CloudCenter Suite License.

If you use specific resources, the CloudCenter Suite reports each usage to the Cisco SSM to tally the number of times that this resource was used and report it in the **Count** column. By verifying this usage count, Cisco SSM calculates the license usage and compliance.

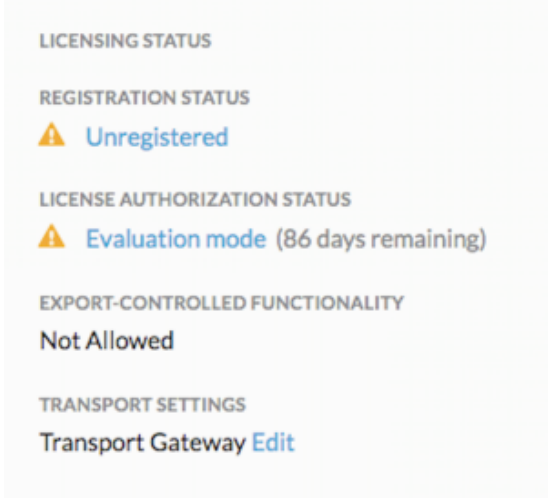
Cisco SSM or Cisco SSM Satellite totals the license requirements for all your CloudCenter Suite instances and compares the total license usage to the number of licenses purchased, on a daily basis.

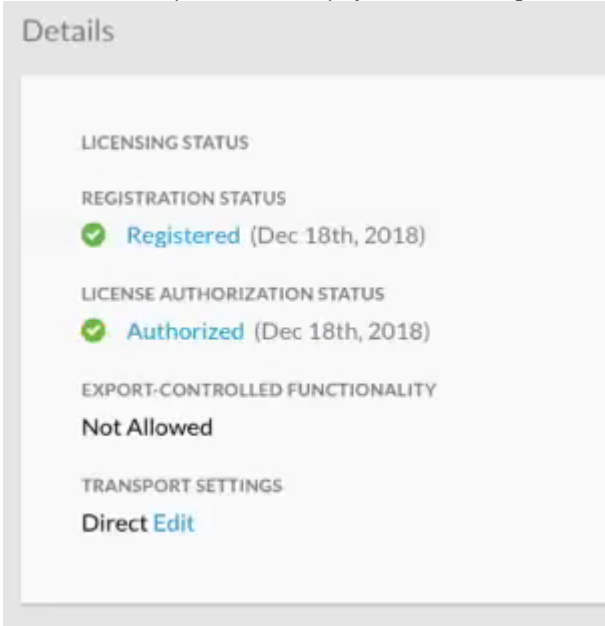
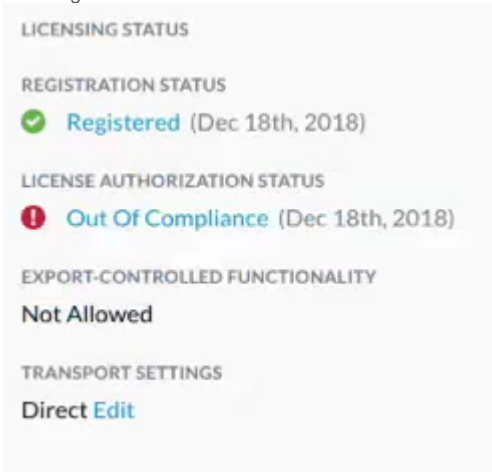
After the data synchronization, your CloudCenter Suite instance displays one of the **Registration Status** indicators listed in the following table.

Registration Status	Description
Unregistered	<p>The Smart Software Licensing is running in Evaluation mode and you have not yet registered the CloudCenter Suite. This status is identified in the following screenshot when you click on the Licensing icon in the orange, exclamation icon in the following screenshot.</p>

<p>Registered</p>	<p>The product registration was completed and an ID certificate was received and will be used for future communication with the Cisco licensing authority. This status is identified in the following screenshot.</p>  <p>The screenshot shows a 'Details' page for a CloudCenter Suite instance. A dark blue tooltip titled 'REGISTRATION DETAILS' is overlaid on the page. The tooltip contains the text: 'This module is registered for Smart Software Licensing.' The background page shows the following information:</p> <ul style="list-style-type: none"> LICENSING STATUS: (Not explicitly labeled in the screenshot) REGISTRATION STATUS: Registered (with a green checkmark icon) LICENSE AUTHORIZATION STATUS: N/A (with a yellow warning triangle icon) EXPORT-CONTROLLED FUNCTIONALITY: Not Allowed TRANSPORT SETTINGS: Direct Edit <p>The tooltip also displays the following registration details:</p> <ul style="list-style-type: none"> REGISTRATION DETAILS: This module is registered for Smart Software Licensing. LAST RENEWAL ATTEMPT: Dec 18th 2018 7:11 PM NEXT RENEWAL ATTEMPT: Jun 16th 2019 8:11 PM REGISTRATION EXPIRES: Dec 18th 2019 7:08 PM
--------------------------	--

After the data synchronization, your CloudCenter Suite instance displays one of the **Licensing Authorization Status** indicators as explained in the following table.

<p>License Authorization Status</p>	<p>Description</p>
<p>Evaluation Mode (countdown from 90 days)</p>	<p>You must register your CloudCenter Suite instance with Cisco SSM or Cisco SSM Satellite before the 90-day evaluation period expires. This state is displayed in the following screenshot.</p>  <p>The screenshot shows a 'Details' page for a CloudCenter Suite instance in Evaluation Mode. The page displays the following information:</p> <ul style="list-style-type: none"> LICENSING STATUS: (Not explicitly labeled) REGISTRATION STATUS: Unregistered (with a yellow warning triangle icon) LICENSE AUTHORIZATION STATUS: Evaluation mode (86 days remaining) (with a yellow warning triangle icon) EXPORT-CONTROLLED FUNCTIONALITY: Not Allowed TRANSPORT SETTINGS: Transport Gateway Edit

<p>Authorized</p>	<p>The number of licenses purchased is sufficient. Registration is complete and valid and the license consumption has started. This state indicates compliance and is displayed in the following screenshot.</p> 
<p>Authorization Expired</p>	<p>The product has not communicated with Cisco SSM or Cisco SSM Satellite for a period of 90 days.</p> <p>The product has been unable to communicate with the Cisco SSM for an extended period of time. This state is due to non-communication with Cisco SSM or Cisco SSM Satellite for more than 90 days. The product will attempt to contact the Cisco SSM every hour in order to renew the authorization until the registration period expires.</p>
<p>Out of Compliance</p>	<p>The number of licenses is insufficient. Consumption exceeds available licenses in the Virtual Account. This state is displayed in the following screenshot.</p> 

Workflow of Cisco Smart Software Licensing

The following table describes the workflow of Cisco Smart Software Licensing.

Task	See the Related Section
Generate a product instance registration token in your virtual account	Configure Smart Licenses#Generating a Registration Token
Configure the transport settings using which CloudCenter Suite connects to Cisco SSM or Cisco SSM Satellite	Configure Smart Licenses#Configuring Transport Settings
Register the CloudCenter Suite instance with Cisco SSM or Cisco SSM Satellite	Configure Smart Licenses#Registering a CloudCenter Suite License

Manage licenses

- [Configure Smart Licenses#Renewing Authorization](#)
- [Configure Smart Licenses#Re-Registering a CloudCenter Suite License](#)
- [Configure Smart Licenses#De-Registering a CloudCenter Suite License](#)

Generating a Registration Token

You need to generate a registration token from Cisco SSM or Cisco SSM Satellite to register the CloudCenter Suite instance.



Ensure that you have set up a Smart Account and a Virtual account on Cisco SSM or Cisco SSM Satellite.

To generate a registration token, follow this procedure.

1. Log in to your Smart Account using Cisco SSM or Cisco SSM Satellite.
2. Navigate to the Virtual account using which you want to register the CloudCenter Suite instance.
3. If you want to enable higher levels of encryption for the products registered using the registration token, check the **Allow export-controlled** functionality on the products registered with this token check box.



This option is available only if your smart account is enabled for Export Control.

4. Click **New Token** to generate a registration token.
5. Copy and save the token so you can use it when you register your CloudCenter Suite instance.
6. For more information on registering your CloudCenter Suite instance, see [Configure Smart Licenses#Registering a CloudCenter Suite License](#).

Configuring Transport Settings

By default, CloudCenter Suite directly communicates with the Cisco SSM. You can modify the mode of communication by configuring the transport settings.



Ensure that you have obtained the registration token for the CloudCenter Suite instance.

To configure the transport settings, follow this procedure.

1. Navigate to the [Suite Admin Dashboard](#).
2. Click **Licensing** in the left tree pane. If you are running CloudCenter Suite in the Evaluation mode, a license notification is displayed on the Smart Software Licensing pane.
3. If a license notification is displayed, click the **Edit Transport Settings** link that is highlighted in the following screenshot.

The screenshot shows the Cisco Suite Admin interface. The top navigation bar includes the Cisco logo, 'Suite Admin', and a user profile 'Welcome, first'. The main content area is titled 'Smart Software Licensing' and features an 'ACTIONS' dropdown menu. A prominent blue information box contains the following text: 'You are currently running in Evaluation Mode. To register your Cisco CloudCenter Suite with Cisco Smart Software Licensing:'. Below this, there are four bullet points: 'Ensure this product has access to the internet or a Smart Software Manager satellite installed on your network. This might require you to **EDIT TRANSPORT SETTINGS**', 'Log into your Smart Account in SMART SOFTWARE MANAGER or your Smart Software Manager satellite.', 'Navigate to the Virtual Account containing the licenses to be used by this Product Instance.', and 'Generate a Product Instance Registration Token (this identifies your Smart Account) and copy or save it.' At the bottom of the information box are two buttons: 'REGISTER' and 'LEARN MORE ABOUT SMART SOFTWARE'. Below the information box is a 'Details' section with two tabs: 'LICENSING STATUS' and 'ACCOUNT & PRODUCT DETAILS'.

Alternatively, click the **Licensing Status** tab, and then click the **View/Edit** link that appears under **Transport Settings**.

4. In the Transport Settings dialog box displayed in the following screenshot, perform one of these steps:

Transport Settings

Configure how the product will communicate with Cisco. Note that this setting is shared with Smart Call Home, so any changes made here will apply to other features using this service.

Direct (Communicate directly via Cisco URL)

https://tools.cisco.com/its/service/oddce/service/DDCEService

CONNECT USING HTTP/HTTPS PROXY?

III OFF

Transport Gateway (Communicate via proxy data or Smart Software Manager satellite)

URL

CANCEL SAVE

- To configure CloudCenter Suite to send the license usage information to Cisco SSM using the Internet (default):
 - a. Click the **Direct** switch to communicate directly using the Cisco URL.
 - b. Configure a DNS on CloudCenter Suite to resolve tools.cisco.com.
 - To configure CloudCenter Suite to send the license usage information to Cisco SSM using the Cisco SSM Satellite:
 - a. Click the **Transport Gateway** button.
 - b. Enter the URL of the Cisco SSM Satellite.
 - To configure CloudCenter Suite to send the license usage information to Cisco SSM using a proxy server. For example, an off-the-shelf proxy, such as Cisco Transport Gateway or Apache:
 - a. Toggle the **HTTP/HTTPS Proxy** switch.
 - b. Enter the IP address and port number of the proxy server.
5. Click **Save**.

Registering a CloudCenter Suite License

You need to register your CloudCenter Suite instance with Cisco SSM or Cisco SSM Satellite before the 90-day evaluation period expires.

- ✔ Ensure that you have configured the transport settings.

To register the CloudCenter Suite license, follow this procedure.

1. Navigate to the [Suite Admin Dashboard](#).
2. Click **Licensing** in the left tree pane.
3. In the license notification, click **Register**. The Smart Software Licensing Product Registration dialog box appears.
4. In the Product Instance Registration Token field, paste the registration token that you generated using the Cisco SSM or Cisco SSM Satellite. For more information on generating a registration token, see [Configure Smart Licenses#Generating a Registration Token](#).
5. Click **Register** to complete the registration process. The CloudCenter Suite sends a request to Cisco SSM or Cisco SSM Satellite to check the registration status and Cisco SSM or Cisco SSM Satellite reports back the status to CloudCenter Suite, on a daily basis. If registering the token fails, you can re-register the CloudCenter Suite instance using a new token. For more information on re-registering CloudCenter Suite, see [Configure Smart Licenses#Re-Registering a CloudCenter Suite License](#).


Renewing Authorization

By default, the authorization is automatically renewed every 30 days. However, CloudCenter Suite allows a user to manually initiate the authorization renew in case the automatic renewal process fails. The authorization expires if CloudCenter Suite is not connected to Cisco SSM or Cisco SSM Satellite for 90 days and the licenses consumed by CloudCenter Suite are reclaimed and put back to the license pool.

- ✔ Ensure that the CloudCenter Suite instance is registered with Cisco SSM or Cisco SSM Satellite.

To renew authorization, follow this procedure.


1. Navigate to the [Suite Admin Dashboard](#).
2. Click **Licensing** in the left tree pane.
3. From the **Actions** drop-down list, choose **Renew Authorization Now** as displayed in the Actions dropdown in the following screenshot.



4. Click **OK** in the **Renew Authorization** dialog box to confirm authorization renewal. The **CloudCenter Suite** synchronizes with **Cisco SSM** or **Cisco SSM Satellite** to check the license authorization status and **Cisco SSM** or **Cisco SSM Satellite** reports back the status to **CloudCenter Suite**, on a daily basis.

Re-Registering a CloudCenter Suite License

You can re-register **CloudCenter Suite** with **Cisco SSM** or **Cisco SSM Satellite** by de-registering it and registering it again, or by using a register force option.


 Ensure that you have obtained a new registration token from **Cisco SSM** or **Cisco SSM Satellite**.

To re-register **CloudCenter Suite** license, follow this procedure.

1. Navigate to the [Suite Admin Dashboard](#).
2. Click **Licensing** in the left tree pane.
3. From the **Actions** drop-down list, choose **Reregister**.
4. In the **Product Instance Registration Token** field of the **Smart Software Licensing Product Re-registration** dialog box, enter the registration token that you generated using **Cisco SSM** or **Cisco SSM Satellite**. For more information on generating a registration token, see [Generating a Registration Token](#).
5. Click **Register** to complete the registration process. The **CloudCenter Suite** sends a request to **Cisco SSM** or **Cisco SSM Satellite** to check the registration status and **Cisco SSM** or **Cisco SSM Satellite** reports back the status to **CloudCenter Suite**, on a daily basis.

De-Registering a CloudCenter Suite License

You can de-register the **CloudCenter Suite** instance from **Cisco SSM** or **Cisco SSM Satellite** to release all the licenses from the current **Virtual account** and the licenses are available for use by other products in the virtual account. De-registering disconnects **CloudCenter Suite** from **Cisco SSM** or **Cisco SSM Satellite**.

 Ensure that the **CloudCenter Suite** instance is registered with **Cisco SSM** or **Cisco SSM Satellite**.

To de-register **CloudCenter Suite** license, follow this procedure.

1. Navigate to the [Suite Admin Dashboard](#).
2. Click **Licensing** in the left tree pane.
3. From the **Actions** drop-down list, choose **Deregister**.
4. Click **Deregister** in the confirmation dialog box. The **CloudCenter Suite** sends a request to **Cisco SSM** or **Cisco SSM Satellite** to check the de-registration status and **Cisco SSM** or **Cisco SSM Satellite** reports back the status to **CloudCenter Suite**, on a daily basis.

Enable for Production

Toggle the **Enable for Production** switch to use the license in production mode displayed in the following screenshot. When you purchase one license for the **CloudCenter Suite**, you automatically receive a free non-production license as well. Both modes are independent of each other and you can switch from one mode to the other any number of times.

Details

LICENSING STATUS	ACCOUNT & PRODUCT DETAILS
REGISTRATION STATUS Registered (Jan 23rd, 2019)	PRODUCT INSTANCE NAME Cisco CloudCenter Suite
LICENSE AUTHORIZATION STATUS Out Of Compliance (Jan 23rd, 2019)	SMART ACCOUNT CloudCenterSuite
EXPORT-CONTROLLED FUNCTIONALITY Not Allowed	VIRTUAL ACCOUNT Default
TRANSPORT SETTINGS Direct Edit	ENABLE FOR PRODUCTION <input type="checkbox"/>

Usage

LICENSING	DESCRIPTION	STATUS	COUNT
CCS_BASE_PLATFORM	CloudCenter Suite Production Platform	Init	1

When the CloudCenter Suite is in non-production mode, the entitlement tags do not validate the license for usage, in which case, you can use it for development, testing, or staging purposes.

Troubleshooting Licensing Issues

This section identifies issues that you may encounter when dealing with licenses.

Invalid Token

When you see the message displayed in the following screenshot for your instance, verify if your token is still valid and if it needs to be renewed.

The screenshot shows the Cisco Suite Admin interface. At the top, there is a header with the Cisco logo and 'Suite Admin'. Below the header, there is a section titled 'Smart Software Licensing' with an 'ACTIONS' dropdown menu. A red error message is displayed: 'LAST ATTEMPT TO RENEW REGISTRATION FAILED.' Below this, there is a 'Details' section with a 'LICENSING STATUS' label.

Download Logs

If you have any issues with Smart Licenses, download the logs files by using the UI (see [Monitor Modules > Download Logs](#)) or the suite-logs/v2/api-docs (see [Logs Service API Calls](#)) and contact the [Smart License team](#).

Module Lifecycle Management

Module Lifecycle Management

- [Install Module](#)
- [Update Module](#)
- [Monitor Modules](#)

Install Module

Install Module

- [Overview](#)
- [Requirements](#)
- [Process](#)
- [Free License](#)
- [Module Actions](#)
- [Uninstall a Module](#)
- [Module States](#)

Overview

The [Suite Admin Dashboard](#) lists the available modules in the Display pane. If you are installing each module for the first time, you will see the **Install** button enabled. Once installed, each module may be in various lifecycle phases as described in this section.

Requirements

Be sure to adhere to the following requirements:

- If your current cluster does not have sufficient resources to meet the minimum requirements mentioned in the [Prepare Infrastructure](#) section, then the installation process will be blocked and you will need to resolve these issues by scaling up to these requirements (see [Manage Clusters > Scale Up](#) for details).
- Only a suite administrator can install a module. By installing the module, this suite administrator automatically inherits the module admin role as well.
- Be sure to synchronize the server time for all instances running the CloudCenter Suite as this can potentially cause module install or upgrade to fail.

Process


You can install multiple modules simultaneously.

To install a module, follow this procedure.

1. Navigate to the [Suite Admin Dashboard](#).
2. Click **Install** on the required module. This procedure uses the Cost Optimizer as an example. The following screenshot displays the available modules.



After installing Action Orchestrator, be aware that you must wait for 2-3 minutes before accessing the application.



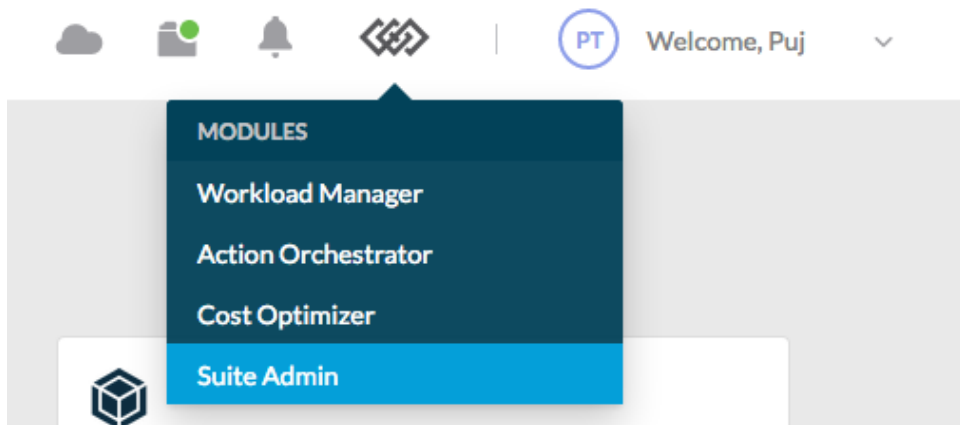
3. In the **You're updating *module name*** popup, select the required version from the dropdown list.

 Once installed, you cannot revert to a previous version.

4. The module starts its installation process and displays a progress bar indicator.

5. Once Installed, you can perform the following actions:

- Click a module to [Monitor Modules](#).
- Open the module or uninstall the module (see the section below).
- Navigate back and forth to other modules and the Suite Admin using the navigation icon in the header as displayed in the following screenshot.



You have now installed one of the modules in the CloudCenter Suite.

Free License

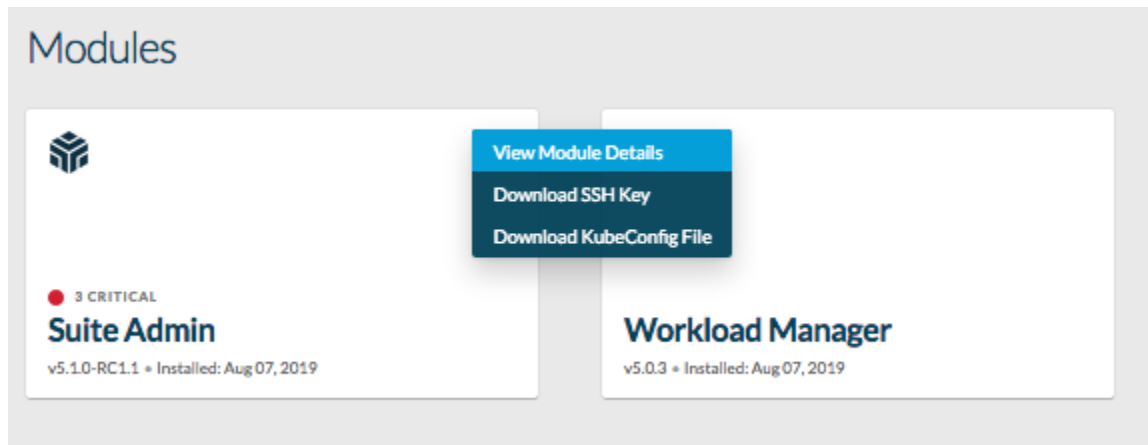
When you install any module, you see the countdown for the 90-day free license time remaining for the license in the top left portion of the module. See [Con figure Smart Licenses](#) for details.

Module Actions

Once installed, the suite administrator can perform the following actions on a module:

- [Update Module](#)
- [Monitor Modules](#)
- [Configure Smart Licenses](#)
- [Manage Module-Specific Content](#)

The Suite Admin module allows the additional actions displayed in the following screenshot:



- Download SSH Key (used to connect to the cluster).
- Download KubeConfig file (used to view cluster information).
- See [Cluster Status](#) for additional context.

Uninstall a Module




After you uninstall any module, verify that all dependent resources have been deleted.


Before re-installing a module that was previously installed, verify that the volumes, secrets, and other dependent details have been cleaned up.

To uninstall a module, follow this procedure.

1. Click the module's dropdown and select **Uninstall** as displayed in the following screenshot.



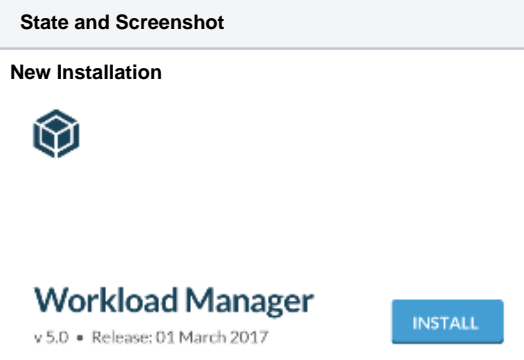
2. Confirm your intention to uninstall as all your content will be deleted as displayed in the following screenshot.





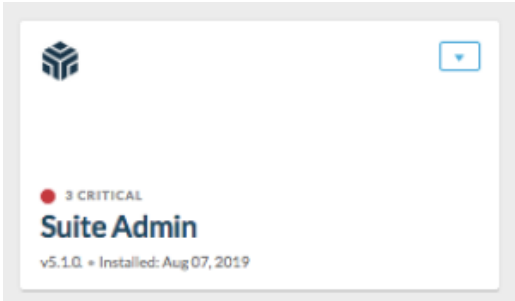



3. The module starts its uninstallation process. Uninstallation takes a few minutes as the CloudCenter Suite cleans up all aspects of the installation.

Module States

The following table provides details on the various module states.

State and Screenshot	Description
<p>New Installation</p> 	<p>A new module is available for installation in the Suite Admin Dashboard.</p>

<p>Installing (or updating)</p>  <p>Workload Manager Installing v 5.1 - 50%</p> 	<p>The module is being installed/updated and the installation process displays a progress bar indicator.</p>
<p>Licensed</p>  <p>Workload Manager v 5.0 • Installed: 05 July 2017</p>	<p>This screenshot identifies a module that is installed, registered, and licensed. See Configure Smart Licenses for details.</p>
<p>Update Available</p>  <p>Suite Admin v5.1.0-16256 • v5.1.0-16264 is available - Release Notes <input type="button" value="UPDATE"/></p>	<p>Once a new software version becomes available, the module displays the new version availability and provides a link to the documentation website. See Update Module for details.</p> <p>The release notes link for the available release is directly linked to the release notes for each module.</p> <p>The dropdown list also provides additional options for each module</p>
<p>Alerts</p> 	<p>When alerts are generated, they are displayed in the Suite AdminDashboard (dropdown list for this module) > View Module Details > Alerts tab.</p> <p>The number of alerts are also identified in the corresponding module tile that are displayed in the Suite AdminDashboard (the screenshot identifies that 3 Warning alerts are available for this module)</p> <p>See Monitor Modules for details.</p>
<p>Validation Error</p>  <p>Workload Manager Failed to install - Please Try again</p>	<p>The module installation resulted in an error. See Troubleshoot Suite Admin for additional details.</p>

Update Module

Update Module

- [Overview](#)
- [Considerations](#)
- [Limitations](#)
- [Process](#)
- [Module Actions](#)
- [Required Post-Kubernetes Upgrade Configuration Tasks](#)

Overview

The suite administrator can only upgrade the module to later versions of the software and will not be able to revert to an earlier version of the software.

Considerations

Before updating a module, see the following module considerations:

- [Workload Manager Installation Overview](#) > *Module Update Considerations*
- [Cost Optimizer Overview](#) > *Module Update Considerations*
- [Action Orchestrator](#) > *Migrating Database*

Limitations

Only a suite administrator can update a module.

Once a new software version becomes available, the module displays the new version availability and provides a link to the documentation website.

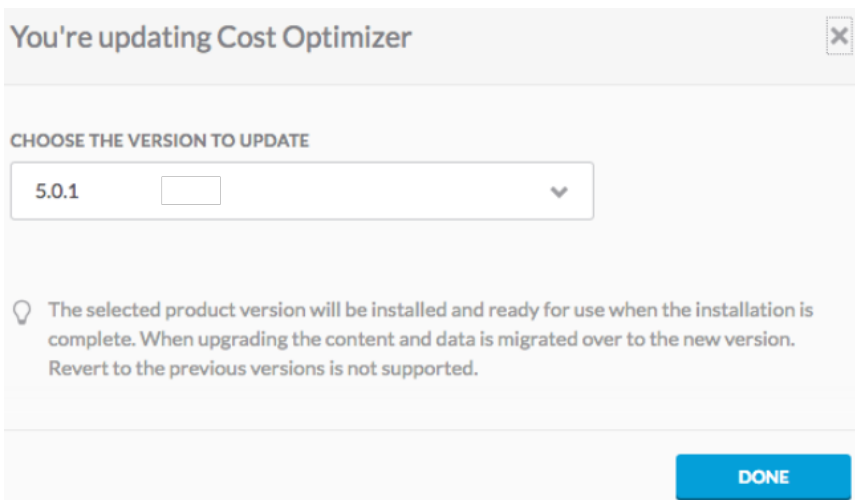
Process



- Before updating any module, verify that you have un-allocated CPU/Memory in your cluster to ensure that your environment has free CPU/Memory a module-update scenario requires additional resources for the old pod to continue running until the new pod initializes and takes over. This additional resource requirement is temporary and only required while a module update is in Progress. After the module is updated, the additional resources are no longer needed.
- You must update the Suite Admin module before you update any other CloudCenter Suite module.
- **Update only one module at a time. If you simultaneously update more than one module, your update process may fail due to limited resource availability. See [Prepare Infrastructure](#) for additional context.**
- You may see one or more error messages during the update process. Be aware that these messages will not affect the update itself.

To update a module, follow this process.

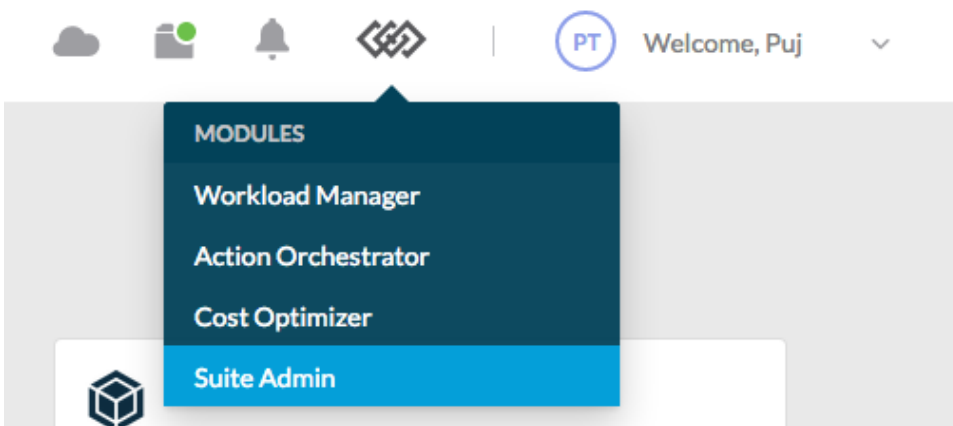
1. Navigate to the [Suite Admin Dashboard](#).
2. Select the required version and click **Done** to upgrade this module. The following screenshot displays Cost Optimizer as an example. All available releases are displayed in the dropdown list in descending order with the latest version at the start of the list.



3. The module starts its upgrade process and displays a progress bar indicator.
4. Once Installed, you can click the module to access the details of that module

or

Navigate to other modules using the module navigation icon in the header as displayed in the following screenshot.



You have now updated the modules in the CloudCenter Suite.

Module Actions

Once a module is upgraded, the suite administrator can perform the following actions on a module:

- [Monitor Modules](#)
- [Configure Smart Licenses](#)
- [Manage Module-Specific Content](#)

Required Post-Kubernetes Upgrade Configuration Tasks

The CloudCenter Suite connects to a Cisco hosted Helm repository and a Docker registry to check for available modules and updates. These repositories are fully compliant with export control and requires authentication for each user connecting to the repository. All CloudCenter Suite modules are packaged as Helm Charts and Docker images. The Helm Chart refers to Docker images via the image's SHA256 hash. The Helm Chart itself is signed and verified by the CloudCenter Suite upon installation or upgrade. This way the integrity of the Helm Chart and Docker images are guaranteed.

Due to changes in underlying Kubernetes versions, cert-manager & deprecated Kubernetes APIs, the Helm chart upgrades will not work directly. We need to migrate the references of deprecated APIs to new supported versions.

To ensure the chart works, perform the following steps:

1. Execute the below Helm commands with version helm-v2.16.3 from inside the common-framework-prod-mgmt pod which already has this version installed:

```
kubectl exec -it $(kubectl get pods -n cisco | grep prod-mgmt | cut -f1 -d' ') -n cisco /bin/sh
```

2. Run the following command:

```
mkdirp $(helm home)/plugins
```

3. Run the following command:

```
helm plugin install https://github.com/helm/helm-mapkubeapis --version=v0.1.0
```

4. Create a Map.yaml file with content shown below this procedure.

5. Run the following command to get all the helm releases

```
helm ls --tiller-namespace cisco
```

6. Repeat the above command for all helm releases.

```
helmmapkubeapis<replace with release name> namespace cisco v2 mapfileMap.yaml
```

7. Wait formapkubeapisto finish.

8. Update the CCS modules from SA UI.

For Airgap setups with no Internet, users can run the commands from a client with Internet access with helm-v2.16.3.

TheCloudCenterSuiteoffers granular control of access to eachCloudCenterSuite resource through role-based, module-level access control. Access toresources like services, clouds, application profiles, deployment environments, and otherCloudCenterSuiteresources can be managed based on rolesassociated with users or user groups. SeeUnderstand Rolesfor details.

Content required for step 4:

```
mappings:
- deprecatedAPI: "apiVersion: extensions/v1beta1\nkind: Deployment"
  newAPI: "apiVersion: apps/v1\nkind: Deployment"
  deprecatedInVersion: "v1.9"
  removedInVersion: "v1.16"
- deprecatedAPI: "apiVersion: apps/v1beta1\nkind: Deployment"
  newAPI: "apiVersion: apps/v1\nkind: Deployment"
  deprecatedInVersion: "v1.9"
  removedInVersion: "v1.16"
- deprecatedAPI: "apiVersion: apps/v1beta2\nkind: Deployment"
  newAPI: "apiVersion: apps/v1\nkind: Deployment"
  deprecatedInVersion: "v1.9"
  removedInVersion: "v1.16"
- deprecatedAPI: "apiVersion: apps/v1beta1\nkind: StatefulSet"
  newAPI: "apiVersion: apps/v1\nkind: StatefulSet"
  deprecatedInVersion: "v1.9"
  removedInVersion: "v1.16"
- deprecatedAPI: "apiVersion: apps/v1beta2\nkind: StatefulSet"
  newAPI: "apiVersion: apps/v1\nkind: StatefulSet"
  deprecatedInVersion: "v1.9"
  removedInVersion: "v1.16"
- deprecatedAPI: "apiVersion: extensions/v1beta1\nkind: DaemonSet"
  newAPI: "apiVersion: apps/v1\nkind: DaemonSet"
  deprecatedInVersion: "v1.9"
  removedInVersion: "v1.16"
- deprecatedAPI: "apiVersion: apps/v1beta2\nkind: DaemonSet"
  newAPI: "apiVersion: apps/v1\nkind: DaemonSet"
  deprecatedInVersion: "v1.9"
  removedInVersion: "v1.16"
- deprecatedAPI: "apiVersion: extensions/v1beta1\nkind: ReplicaSet"
```

```
newAPI: "apiVersion: apps/v1\nkind: ReplicaSet"  
deprecatedInVersion: "v1.9"  
removedInVersion: "v1.16"  
- deprecatedAPI: "apiVersion: apps/v1beta1\nkind: ReplicaSet"  
newAPI: "apiVersion: apps/v1\nkind: ReplicaSet"  
deprecatedInVersion: "v1.9"  
removedInVersion: "v1.16"  
- deprecatedAPI: "apiVersion: apps/v1beta2\nkind: ReplicaSet"  
newAPI: "apiVersion: apps/v1\nkind: ReplicaSet"  
deprecatedInVersion: "v1.9"  
removedInVersion: "v1.16"  
- deprecatedAPI: "apiVersion: extensions/v1beta1\nkind: NetworkPolicy"  
newAPI: "apiVersion: networking.k8s.io/v1\nkind: NetworkPolicy"  
deprecatedInVersion: "v1.8"  
removedInVersion: "v1.16"  
- deprecatedAPI: "apiVersion: extensions/v1beta1\nkind: PodSecurityPolicy"  
newAPI: "apiVersion: policy/v1beta1\nkind: PodSecurityPolicy"  
deprecatedInVersion: "v1.10"  
removedInVersion: "v1.16"  
- deprecatedAPI: "apiVersion: apiextensions.k8s.io/v1beta1\nkind: CustomResourceDefinition"  
newAPI: "apiVersion: apiextensions.k8s.io/v1\nkind: CustomResourceDefinition"  
deprecatedInVersion: "v1.16"  
removedInVersion: "v1.19"  
- deprecatedAPI: "apiVersion: extensions/v1beta1\nkind: Ingress"  
newAPI: "apiVersion: networking.k8s.io/v1beta1\nkind: Ingress"  
deprecatedInVersion: "v1.14"  
removedInVersion: "v1.22"  
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1alpha1\nkind: ClusterRole"  
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: ClusterRole"  
deprecatedInVersion: "v1.17"  
removedInVersion: "v1.22"  
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1alpha1\nkind: ClusterRoleList"  
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: ClusterRoleList"  
deprecatedInVersion: "v1.17"  
removedInVersion: "v1.22"  
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1alpha1\nkind: ClusterRoleBinding"  
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: ClusterRoleBinding"  
deprecatedInVersion: "v1.17"  
removedInVersion: "v1.22"  
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1alpha1\nkind: ClusterRoleBindingList"  
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: ClusterRoleBindingList"  
deprecatedInVersion: "v1.17"  
removedInVersion: "v1.22"  
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1alpha1\nkind: Role"  
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: Role"  
deprecatedInVersion: "v1.17"  
removedInVersion: "v1.22"  
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1alpha1\nkind: RoleList"  
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: RoleList"  
deprecatedInVersion: "v1.17"  
removedInVersion: "v1.22"  
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1alpha1\nkind: RoleBinding"  
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: RoleBinding"  
deprecatedInVersion: "v1.17"  
removedInVersion: "v1.22"  
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1alpha1\nkind: RoleBindingList"  
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: RoleBindingList"  
deprecatedInVersion: "v1.17"  
removedInVersion: "v1.22"  
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1beta1\nkind: ClusterRole"  
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: ClusterRole"  
deprecatedInVersion: "v1.17"  
removedInVersion: "v1.22"  
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1beta1\nkind: ClusterRoleList"  
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: ClusterRoleList"  
deprecatedInVersion: "v1.17"  
removedInVersion: "v1.22"  
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1beta1\nkind: ClusterRoleBinding"  
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: ClusterRoleBinding"  
deprecatedInVersion: "v1.17"  
removedInVersion: "v1.22"
```

- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1beta1\nkind: ClusterRoleBindingList"
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: ClusterRoleBindingList"
deprecatedInVersion: "v1.17"
removedInVersion: "v1.22"
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1beta1\nkind: Role"
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: Role"
deprecatedInVersion: "v1.17"
removedInVersion: "v1.22"
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1beta1\nkind: RoleList"
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: RoleList"
deprecatedInVersion: "v1.17"
removedInVersion: "v1.22"
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1beta1\nkind: RoleBinding"
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: RoleBinding"
deprecatedInVersion: "v1.17"
removedInVersion: "v1.22"
- deprecatedAPI: "apiVersion: rbac.authorization.k8s.io/v1beta1\nkind: RoleBindingList"
newAPI: "apiVersion: rbac.authorization.k8s.io/v1\nkind: RoleBindingList"
deprecatedInVersion: "v1.17"
removedInVersion: "v1.22"
- deprecatedAPI: "apiVersion: certmanager.k8s.io/v1alpha1\nkind: Certificate"
newAPI: "apiVersion: cert-manager.io/v1\nkind: Certificate"
deprecatedInVersion: "v1.15"
removedInVersion: "v1.16"
- deprecatedAPI: "apiVersion: certmanager.k8s.io/v1alpha1\nkind: Issuer"
newAPI: "apiVersion: cert-manager.io/v1\nkind: Issuer"
deprecatedInVersion: "v1.15"
removedInVersion: "v1.16"


Monitor Modules

Monitor Modules

- [Overview](#)
- [Accessing a Module](#)
- [View Logs in Kibana](#)
- [Download Logs](#)
- [The Grafana Dashboard Alert](#)
- [Default Alert Categories](#)
- [Type of Alerts](#)
- [Alert Types](#)
- [Viewing Alerts in Grafana](#)
- [Setup Grafana Email Alerts](#)


Overview

Once Installed, you can click a module to access the **Module Details** page displayed in the following screenshot. If you click the Workload Manager, the following screenshot displays the corresponding page to monitor this module.




The module name displays at the top of the page and you can perform the following actions on this page:

- Perform one of the actions listed in the **Dropdown** next to the *Module* name as displayed in the following screenshot:



Alternately, click the dropdown from the Module Alerts page as displayed in the following screenshot:



- View the **Alerts** Tab See the *Understand Dashboard Alerts* section below.
- Access the **License Usage** Tab

Accessing a *Module*

There are numerous ways for you to access a module in the CloudCenter Suite. However, your **User Levels** determine if you can access the module!

View Logs in Kibana

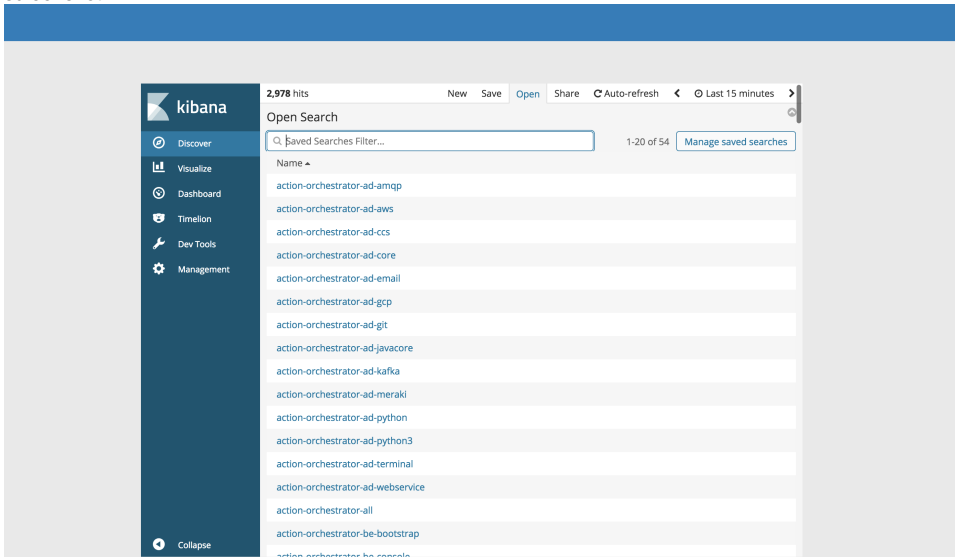
Kibana is a web interface that can be used to search and view the logs for **any of the CloudCenter Suite modules**.

CloudCenter Suite log file use the standard log format:

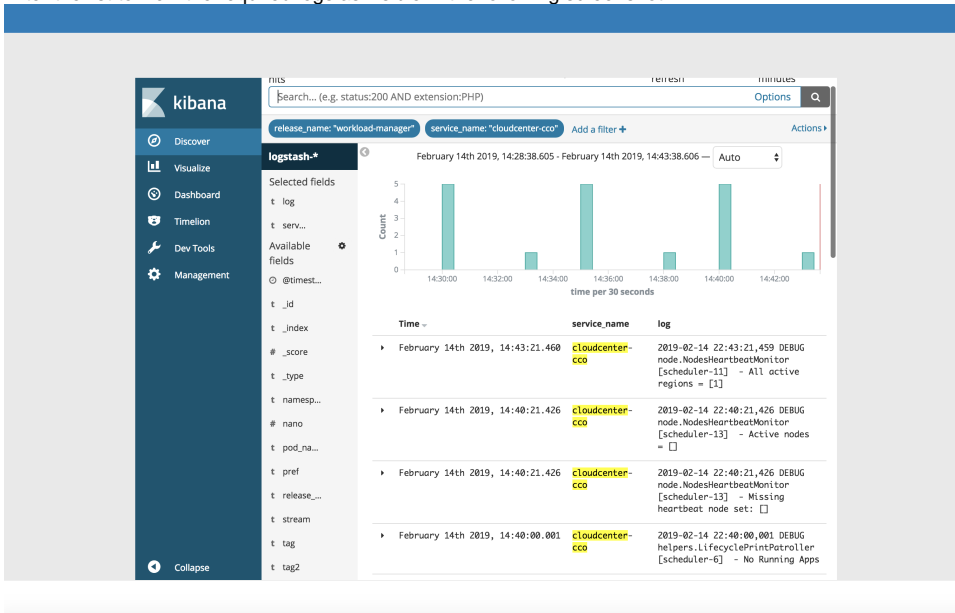
- Where relevant, modules display the user and tenant information.
- You can search by *userid* or *tenantId* when users view logs in Kibana.
- The log files support JSON format.

To view the Kibana logs, follow this procedure.

1. Click the module dropdown and select **View Logs in Kibana** from the dropdown to display the Kibana dashboard visible in the following screenshot.



2. Click **Discover > Open** to list and filter the available logs for this module.
3. Filter the list to view the required logs as visible in the following screenshot.



Download Logs

An alternative to viewing logs in Kibana is to download the log files by clicking a module and selecting **Download Logs** from the dropdown as displayed in the following screenshot.

Download Logs X

Choose the time period to download the logs

TIME PERIOD

1W 30D 60D 90D YTD CUSTOM

[VIEW LOGS IN KIBANA](#)
DOWNLOAD

The Grafana Dashboard Alert

Grafana is an open source visualization tool that allows you to create and edit dashboards.

Modules can create their own services to write custom alerts or create alerts in Grafana for services that they wish to monitor.

When alerts are generated, they are displayed in the Suite Admin's *module* details page > **Alerts** tab. When you acknowledge active alerts, they are moved to the Dismissed tab and stored there for 60 days before they are deleted.

Default Alert Categories

The **Alerts** tab lists two categories of alerts which are driven from Grafana.

- Active Alerts: Each active alert lists the following details:
 - A color-coded alert category
 - The alert title [click the alert link](#) to open the chart in Grafana using authorized credentials
 - An alert count only displayed when there is more than one alert
 - A brief description of the alert
 - The alert source
 - The impacted component
 - A snapshot of the chart in Grafana not available for application alerts
 - The timestamp when this alert was issued hovering over this timestamp displays the exact time
 - The option to multi-select multiple alerts the **Dismiss** button becomes visible when you multi-select alerts
- Dismissed Alerts

Type of Alerts

Alert types are described in the following table.

Alert Type	Description
Infrastructure	These alerts pertain to network, disk, CPU, and memory usage derived from module configured Grafana dashboards.
Application	These alerts are derived from application endpoints that provide the current health of the system.

Alert Types

You can filter alerts based on the type. Alert types are described in the following table.

Alert Type	Color	Description
Critical	Red	Red bar on the side. VM launch failure rate is increasing on the configured cloud.
Warning	Orange	The connection to the AMQP server is not stable and has been dropped <i>t</i> times in the last 45 minutes.

Info	Blue	Updates based on endpoint reports.
------	------	------------------------------------


Viewing Alerts in Grafana

When you access the Grafana dashboard, you will see the following sections:

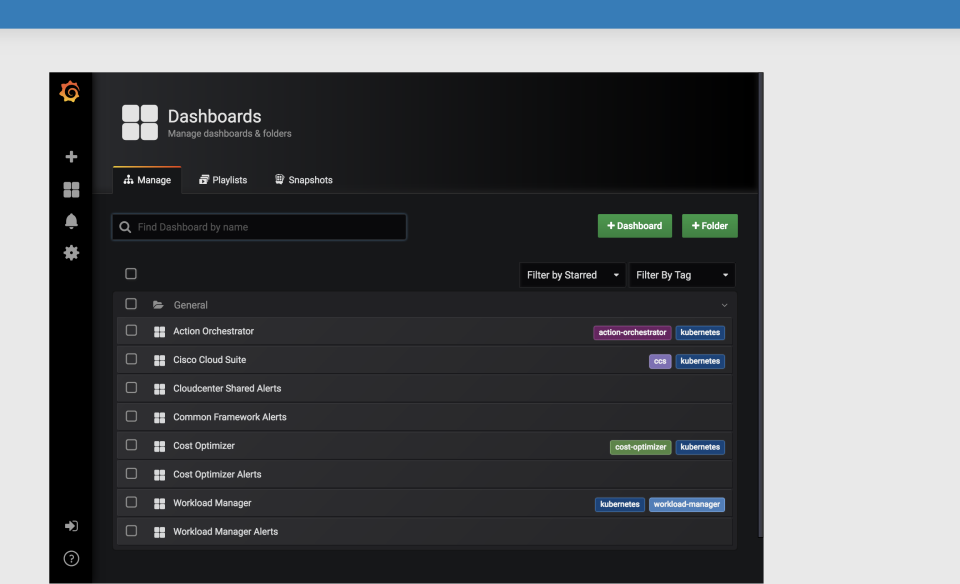
- **System metrics:** CPU usage, memory usage, and crash loops. You can also configure additional alerts in this section, refer to <http://docs.grafana.org/alerting/rules/>.
- **Visualization metrics:** Cluster health, deployments, nodes, pods (number of pods and pods status), containers, and jobs. You cannot configure additional alerts in this section.

To view the Grafana alerts, follow this procedure.

1. Click the module dropdown and select **View Alerts in Grafana** from the dropdown to display the Grafana dashboard visible in the following screenshot.



2. Click **Dashboard > Manage** to list and filter the available alerts for this module.
3. Filter the list to view the required alerts as visible in the following screenshot.



Setup Grafana Email Alerts

To setup email alerts in Grafana, follow this procedure.



Perform this procedure *each* time you upgrade the Suite Admin.

1. Use the following command to edit the configmap for Grafana:

```
kubectl edit configmap common-framework-grafana
```

2. Add the following block to the Grafana configmap:

```
grafana.ini: |
  [smtp]
  enabled = true
  host = smtp.gmail.com
  user = <your email address>@gmail.com
  password = <your password>
```

3. Use the following command to reload Grafana:

```
run kubectl delete po <grafana pod name> to reload grafana
```