



UCC SMI - Release Change Reference

- [New in Documentation](#), on page 2
- [Feature and Behavior Changes Quick Reference](#), on page 3
- [CPU Isolation](#), on page 5
- [Kubernetes 1.26 Upgrade](#), on page 6
- [UCS M7 Server Support](#), on page 7
- [XFS File System](#), on page 8
- [Cilium Addon Support](#), on page 8
- [CIMC Certificate Renewal](#), on page 9
- [Kubernetes 1.25 Upgrade](#), on page 10
- [Support for new Prometheus Parameters](#), on page 11
- [Updated Versions for Third Party Software](#), on page 13
- [Biased Terminologies Update](#), on page 14
- [Calico Version Upgrade](#), on page 15
- [CNDP Container Base Image Version Upgrade](#), on page 16
- [Log Forwarding to Grafana Cloud](#), on page 16
- [Pushing Prometheus Metrics to Grafana Cloud](#), on page 17
- [CNDP Support for cnUPF and cnMME](#), on page 18
- [Hardening System Reliability](#), on page 19
- [IPSec Monitoring](#), on page 20
- [Kubernetes 1.24 Upgrade](#), on page 20
- [Provisionable TLS Certificates](#), on page 21
- [Release Version Upgrade](#), on page 22
- [Ubuntu Version Upgrade](#), on page 22
- [Kubernetes 1.23 Upgrade](#), on page 23
- [UCS M6 Server Support](#), on page 24
- [UCS Server Status Alert](#), on page 24
- [vSphere Datacenter Folder Support](#), on page 25
- [Kubernetes 1.22 Upgrade](#), on page 26
- [Parallel Node Upgrade with Deployment Zone Strategy](#), on page 26
- [Path Based Routing for Inception Server](#), on page 31
- [CA Signed Certificate for Path-based Ingress](#), on page 33
- [OnDemand LDAP Connectivity Check](#), on page 36
- [Alerts for Node Disk Partition Usage](#), on page 38

- [Network Policy for K8s Pods](#), on page 39
- [Splitting Master and Additional Master VIPs into Separate VRRPs](#), on page 40
- [SSH Firewall Rules for Cluster Nodes](#), on page 41
- [Docker Subnet Override Support](#), on page 42
- [cluster connect Command Update](#), on page 43
- [IPSec Support for SMF N4 Interfaces](#), on page 44
- [Push KPIs to S3 Using Thanos](#), on page 48
- [virsh console Command Update](#), on page 50
- [Alert for Standby Cluster Manager Failure](#), on page 51
- [Ubuntu User Password Expiration Configuration Enhancement](#), on page 52
- [Unified RMA Procedure for Planned and Failure Events on Bare Metal](#), on page 53
- [Alert for KVM Node Unreachable](#), on page 53
- [CDL Data Slicing](#), on page 54
- [CDL Overload Protection Enhancement](#), on page 55
- [CLI Support for UPF IFTASK Forwarder Type](#), on page 58
- [Cluster Manager Now Uses Internal Network for HA Communications](#), on page 59
- [Configurable Option to Control Ping Properties](#), on page 61
- [Deleting Stale CDL Slot Data](#), on page 62
- [Dual Stack Support](#), on page 63
- [Emulator Pinning](#), on page 64
- [GR Failover Notifications](#), on page 65
- [Hostname and URL Path-Based Routing for Ingress](#), on page 65
- [Increased NotReady Detection Sensitivity for K8s Nodes](#), on page 66
- [Kubernetes 1.20.0 Upgrade](#), on page 67
- [Node Failure Notifications During RMA](#), on page 67
- [Silence 'Always On' vm-alive Alerts](#), on page 68
- [Smart Agent Upgrade](#), on page 69
- [Support VM Status Alerts on CNDP](#), on page 70
- [tac-debug-pkg CLI Enhancements](#), on page 74
- [User Role APIs](#), on page 75
- [VIP Config Enhancements](#), on page 76
- [VPP CPU Worker Count](#), on page 78

New in Documentation

Information on new features, enhancements, and behavior changes in the Release Change Reference (RCR) document will now be available under the **What's New in this Release** section in the 5G release notes.



Note This document will be deprecated in 2024.01 and later releases.

Feature and Behavior Changes Quick Reference

Features/Behavior Changes	Introduced/Modified
CPU Isolation, on page 5	2023.04.1
Kubernetes 1.26 Upgrade, on page 6	2023.04.1
UCS M7 Server Support, on page 7	2023.04.1
XFS File System, on page 8	2023.04.1
Cilium Addon Support, on page 8	2023.03.1
CIMC Certificate Renewal, on page 9	2023.03.1
Kubernetes 1.25 Upgrade, on page 10	2023.03.1
Support for new Prometheus Parameters, on page 11	2023.03.1
Updated Versions for Third Party Software, on page 13	2023.03.1
Biased Terminologies Update, on page 14	2023.02.1
Calico Version Upgrade, on page 15	2023.02.1
CNDP Container Base Image Version Upgrade, on page 16	2023.02.1
Log Forwarding to Grafana Cloud, on page 16	2023.02.1
Pushing Prometheus Metrics to Grafana Cloud, on page 17	2023.02.1
CNDP Support for cnUPF and cnMME, on page 18	2023.01.1
Hardening System Reliability, on page 19	2023.01.1
IPSec Monitoring, on page 20	2023.01.1
Kubernetes 1.24 Upgrade, on page 20	2023.01.1
Provisionable TLS Certificates, on page 21	2023.01.1
Release Version Upgrade, on page 22	2023.01.1
Ubuntu Version Upgrade, on page 22	2023.01.1
Kubernetes 1.23 Upgrade, on page 23	2022.03.1
UCS M6 Server Support, on page 24	2022.03.1
UCS Server Status Alert, on page 24	2022.03.1
vSphere Datacenter Folder Support, on page 25	2022.03.1
Kubernetes 1.22 Upgrade	2020.02.2.3.10

Features/Behavior Changes	Introduced/Modified
Parallel Node Upgrade with Deployment Zone Strategy	2020.02.2.3.10 2022.02.1
Path Based Routing for Inception Server	2022.02.1
CA Signed Certificate for Path-based Ingress	2022.02.1
OnDemand LDAP Connectivity Check	2022.02.1
Alerts for Node Disk Partition Usage	2022.01.1.02
Network Policy for K8s Pods	2022.01.1.02
Splitting Master and Additional Master VIPs into Separate VRRPs, on page 40	2022.01.1.02 2020.02.2.3.04
SSH Firewall Rules for Cluster Nodes	2022.01.1.02
Docker Subnet Override Support, on page 42	2020.02.7.07
cluster connect Command Update, on page 43	2020.02.2.47 2020.02.7.07
IPSec Support for SMF N4 Interfaces, on page 44	2020.02.2.47
Push KPIs to S3 Using Thanos	2020.02.2.47
virsh console Command Update	2020.02.2.47
Kubernetes 1.21 Upgrade	2020.02.2.41
Alert for Standby Cluster Manager Failure, on page 51	2020.02.2.41
Ubuntu User Password Expiration Configuration Enhancement, on page 52	2020.02.2.41
Unified RMA Procedure for Planned and Failure Events on Bare Metal, on page 53	2020.02.2.41
Alert for KVM Node Unreachable	2020.02.2.19
CDL Data Slicing	2020.02.2.19
CDL Overload Protection Enhancement	2020.02.2.19
CLI Support for UPF IFTASK Forwarder Type	2020.02.2.19
Cluster Manager Now Uses Internal Network for HA Communications	2020.02.2.19
Configurable Option to Control Ping Properties	2020.02.2.19
Deleting Stale CDL Slot Data	2020.02.2.19

Features/Behavior Changes	Introduced/Modified
Dual Stack Support	2020.02.2.19
Emulator Pinning	2020.02.2.19
GR Failover Notifications	2020.02.2.19
Hostname and URL Path-Based Routing for Ingress	2020.02.2.19
Increased NotReady Detection Sensitivity for K8s Nodes	2020.02.2.19
Kubernetes 1.20.0 Upgrade	2020.02.2.19
Node Failure Notifications During RMA	2020.02.2.19
Silence 'Always On' vm-alive Alerts	2020.02.2.19
Smart Agent Upgrade	2020.02.2.19
Support VM Status Alerts on CNDP	2020.02.2.19
tac-debug-pkg CLI Enhancements	2020.02.2.19
User Role APIs	2020.02.2.19
VIP Config Enhancements	2020.02.2.19
VPP CPU Worker Count	2020.02.2.19

CPU Isolation

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on Disabled – Configured Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2023.04.1

Feature Description

SMI provides a higher level of CPU isolation for VPP workers. With CPU isolation, no other processes can be scheduled on "isolcpu" CPUs where VPP workers are pinned.

SMI uses the host profile to define isolcpu that isolates CPUs from the kernel scheduler. It does not prevent K8s containers from changing their affinities to run on isolcpu. Depending on the deployment, SMI also provides the flexibility to use VPP workers and session managers for CPU isolation.

For more information, refer to the [UCC SMI Operations Guide > SMI Cluster Manager Operations > CPU Isolation](#) chapter/section.

Kubernetes 1.26 Upgrade

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Deployment Guide</i>

Revision History

Revision Details	Release
First introduced.	2023.04.1

Feature Description

With this release, the Kubernetes version is upgraded from 1.25 to 1.26.

UCS M7 Server Support

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	SMI
Applicable Platforms	Bare Metal
Feature Default Setting	Enabled – Always On
Related Changes in this Release	Not Applicable
Related Documentation	<ul style="list-style-type: none"> • <i>UCC CEE Configuration and Administration Guide</i> • <i>UCC SMI Deployment Guide</i> • <i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2023.04.1

Feature Description

SMI Bare Metal supports the UCS C220 M7 server with a single socket for Private 5G deployments in this release.



Note The M7 server is not supported for on-prem deployments.

The Cisco UCS C220 M7 Rack Server is a high-density, 1RU, general-purpose infrastructure and application server that provides industry-leading performance and efficiency.

For more information, see the *UCC SMI Deployment Guide*.

XFS File System

Feature Summary and Revision History

Revision History

Revision Details	Release
First introduced.	2023.04.1

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configured Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Feature Description

SMI utilizes the XFS filesystem to install the */data* partition using Mongo DB. XFS works only with new deployments. By default, all partitions are formatted using ext4.

For more information, refer to the [UCC SMI Operations Guide > SMI Cluster Manager Operations > XFS File System](#) chapter/section.

Cilium Addon Support

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on

Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Deployment Guide</i> <i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2023.03.1

Feature Description

Cilium is an open-source project that provides networking and security capabilities for Kubernetes clusters. It is used as a networking and security add-on for Kubernetes, replacing or augmenting the default Kubernetes networking components. Cilium leverages the extended Berkeley Packet Filter (eBPF) technology to provide high-performance networking and security features.

Cilium (version 1.13.2) must be installed as K8s add-on on top of Calico.



Note The Cilium add-on is not fully supported in this release.

To enable the cluster configuration, use the following command:

```
clusters cluster_name addons cilium { enabled | disabled }
```

CIMC Certificate Renewal

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2023.03.1

Feature Description

The Cisco® Integrated Management Controller (IMC) is a baseboard management controller that provides embedded server management for Cisco UCS® C-Series Rack Servers and Cisco UCS S-Series Storage Servers. The Cisco IMC enables system management in the data center and across distributed locations.

The CIMC certificates are valid only for 3 years. If the certificate expires in less than 90 days, it must be renewed.

To renew the CIMC certificate, use the following configuration:

```
config
  clusters cluster_name
    node-defaults ucs-server cimc certificate rehydrate { true | false}
  exit
```

NOTES:

- When the certificate is renewed, the CIMC drops connections for 15 to 60 seconds while the host key is updated.
- The default setting is **false**. When set to **true**, it renews the certificate that expires in less than 90 days.
- Every cluster synchronization log displays the expiry date of the certificate.

Kubernetes 1.25 Upgrade

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Deployment Guide</i>

Revision History

Revision Details	Release
First introduced.	2023.03.1

Feature Description

With this release, the Kubernetes version is upgraded from 1.24 to 1.25.

Support for new Prometheus Parameters

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configured Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC CEE Configuration and Administration Guide</i>

Revision History

Revision Details	Release
First introduced.	2023.03.1

Feature Description

The remote-write feature is enhanced to support new Prometheus parameters for CNDP Grafana Cloud integration. These parameters enable the user to fine tune their setup.

The enhanced CEE dashboard has new panels to display the remote write metrics.

This enhancement also supports new alerts `fluent-proxy-output-retries-failed` and `prometheus-remote-write-behind`. See the *UCC SMI Operations Guide > Alerts Reference* for more information on alerts.

- Remote Timeout:

The **remote-timeout-seconds** command sets the timeout for requests to the remote write endpoint, in seconds. Default: 30 seconds.

The following is a sample configuration:

```
prometheus remote-write target demo
  remote-timeout-seconds 60
  exit
```

- Queue Configuration:

The **queue-config** command configures the queue used to write to remote storage.

The following is a sample configuration:

```
prometheus remote-write target demo
  ...
  queue-config capacity 500
  queue-config max-shards 100
  queue-config min-shards 2
  queue-config max-samples-per-send 300
  queue-config batch-send-deadline-seconds 10
  exit
```

NOTES:

- **queue-config capacity:** Specify the number of samples to buffer per shard. Default: 2500.

It is recommended to have adequate capacity in each shard to buffer several requests. The adequate capacity can maintain the throughput while processing occasional slow remote requests.

- **queue-config max-shards:** Specify the maximum number of shards. Default: 200.
- **queue-config min-shards:** Specify the minimum number of shards. Default: 1.
- **queue-config max-samples-per-send:** Specify the maximum number of samples per send. Default: 500.
- **queue-config batch-send-deadline-seconds:** Specify the maximum time in seconds that a sample will wait in buffer. Default: 5 seconds.

- Relabel Configuration:

The **relabel-configs** command defines a list of relabel configurations before the metrics are written to remote storage. The relabeling feature in Prometheus rewrites the label set of a target dynamically.

The following is a sample configuration:

```
prometheus remote-write target demo
  ...
  relabel-configs test1
    target-label test1_label
    regex      (.+);(.+)
    replacement ${1}@${2}
    action      replace
    source-labels container
    source-labels pod
    exit
  exit
```

NOTES:

- **target-label:** Specify the label to which the resulting value is written in a replace action.
- **regex:** Specify the regular expression against which the extracted value is matched.
Default = (.*)

- **replacement:** Specify the replacement value against which a regex replace is performed if the regular expression matches.
Default = \$1
- **action:** Specify the replace, keep, or drop action to perform based on regex matching.
Default = replace
- **source-labels:** Specify the source label to select values from existing labels.
- Multiple relabeling steps can be configured per scrape configuration. The steps are applied to the label set of each target in order of appearance in the configuration file.
- Note that Prometheus will drop any label with empty value, hence use the labels with caution.

Updated Versions for Third Party Software

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
First introduced.	2023.03.1

Feature Description

The following software versions are upgraded in this release:

Software Package	Component(s)	Previous Version	Current Version
containerd	Inception server	1.6.4	1.7.2
	Base image		
	Cluster deployer		

Software Package	Component(s)	Previous Version	Current Version
Prometheus	Metrics	2.37.2	2.44.0
Docker	—	23.0.1	24.0.2

Biased Terminologies Update

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
First introduced.	2023.02.1

Feature Description

The biased terms in the SMI Kubernetes context only are updated in this release.

- The term **master** is replaced by **control plane**.

For release 2023.02.1, you can use both deprecated and replacement configurations but cannot configure both at the same time. When using deprecated biased configurations, a warning message will be displayed on cluster sync.

- The yang model configuration changes include:

Deprecated Configuration	New Configuration
k8s node-type master	k8s node-type control-plane
node-type-defaults master	node-type-defaults control-plane
kube-master-ip	kube-control-plane-ip

Ops Center Changes

The following is the current Ops Center configuration when node-type is set to **master**.

```
// show run in ops-center:

k8s nodes abc-master
  node-type  master
  worker-type master
exit
```

When the node-type is set to **control-plane**, the Ops Center also reflects the same configuration.

```
// show run in ops-center:

k8s nodes abc-test
  node-type  control-plane
  worker-type control-plane
exit
```

Calico Version Upgrade

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
First introduced.	2023.02.1

Feature Description

The Calico solution is upgraded to version 3.24 in this release.

CNDP Container Base Image Version Upgrade

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
First introduced.	2023.02.1

Feature Description

The container base image version for CNDP is upgraded from 18.04 LTS to 20.04 LTS.

The previous 18.04 base image version is not supported for 2023.02.1 and future releases.

Log Forwarding to Grafana Cloud

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configured Required
Related Changes in this Release	Not Applicable

Related Documentation	<i>UCC CEE Configuration and Administration Guide</i>
-----------------------	---

Revision History

Revision Details	Release
First introduced.	2023.02.1

Feature Description

Log Forwarding allows you to forward the log entries (including the host and container level log entries) stored in the JournalD to the external collectors. This release supports Fluent-Bit log forwarding to Grafana Cloud. Grafana Cloud will be the target host.

Fluent-Bit supports sending logs to Grafana Cloud by providing the appropriate URL and ensuring that TLS is enabled.

For more information, see the *Log Forwarding* section in the *UCC CEE Configuration and Administration Guide*.

Pushing Prometheus Metrics to Grafana Cloud

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configured Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC CEE Configuration and Administration Guide</i>

Revision History

Revision Details	Release
First introduced.	2023.02.1

Feature Description

The CEE leverages the existing remote-write feature to push the Prometheus server metrics to Grafana Cloud. This feature supports the **basic-auth** and **proxy-url** fields of the remote-write configuration to push the metrics to Grafana Cloud.

Configuring Remote Write to Push Prometheus Metrics

To push the Prometheus metrics to Grafana Cloud using remote-write, use the following sample configuration:

```
prometheus remote-write target demo
url https://prometheus-us-central1.grafana.net/api/prom/push
basic-auth username 725569
basic-auth password $8$ntCDRl2FkMD1m8mj9FohYwTuy/jo+7Cka0msfP2qW3Y=
proxy-url http://proxy-wsa.esl.cisco.com:80
exit
```

NOTES:

- **url**—Specify the target URL of Grafana Cloud.
- **basic-auth username**—Specify the username in Confd.
- **basic-auth password**—Specify the password in Confd. The password is encrypted in Confd and passed to the metrics helm chart.
- **proxy-url**—Specify the optional proxy URL to access Grafana Cloud in Confd.

CNDP Support for cnUPF and cnMME

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Deployment Guide</i> <i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2023.01.1

Feature Description

This release supports cloud native deployment of UPF and MME. CNDP supports all cloud-based network functions including cnUPF, cnMME, AMF, and SMF in the same cluster for P5G deployments. This functionality is supported on the UCS M5 and UCS M6 servers.

Hardening System Reliability

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2023.01.1

Feature Description

Currently, the grace period for nodes before it is marked as not ready is 20 seconds.

The grace period is changed from 20 seconds to 5 minutes. This grace time will result in more tolerance of network failures.



Note The pods will take six minutes to reschedule.

IPSec Monitoring

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2023.01.1

Feature Description

SMI allows monitoring of IPSec certificates—sends certificate expiry alerts and updates the certificate through strongSwan configuration. The strongswan configuration supports a new **server-secret** configuration field to pass an existing TLS secret.

For more information, see the *IPSec Support for SMF N4 Interfaces* section in *UCC SMI Operations Guide*.

Kubernetes 1.24 Upgrade

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable

Related Documentation	<i>UCC SMI Deployment Guide</i>
-----------------------	---------------------------------

Revision History

Revision Details	Release
First introduced.	2023.01.1

Feature Description

With this release, the Kubernetes version is upgraded from 1.22.7 or 1.23.8 to 1.24.6.

Provisionable TLS Certificates

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2023.01.1

Feature Description

This release supports provisioning of TLS certificates used for both REST APIs and K8s APIs. The certificates are configurable through cluster manager and Ops-center. You can configure a certificate and its corresponding private key to provision the certificate as a TLS secret using the existing yang container.

The provisioned certificates are also monitored for expiry by setting alerts.

For more information, see *SMI Cluster Manager Operations > CA Signed Certificate for Path-based Ingress* section in the *UCC SMI Operations Guide*.

Release Version Upgrade

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Deployment Guide</i> <i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2023.01.1

Feature Description

The 2023.01.1.x version is upgradable from 2020.02.2.3.12+ for CNDP deployments.

Contact your Cisco account representative for more information.

Ubuntu Version Upgrade

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on

Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Deployment Guide</i> <i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2023.01.1

Feature Description

SMI/CNDP uses a Linux kernel and a hardened Ubuntu base image on all nodes in a cluster.

The Linux source operating system is upgraded from Ubuntu 18.04 to Ubuntu 20.04. The Ubuntu 18.04 version is no longer supported (EoL).

Kubernetes 1.23 Upgrade

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Deployment Guide</i> <i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2022.03.1

Feature Description

With this release, the Kubernetes version is upgraded from 1.22 to 1.23.

UCS M6 Server Support

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	SMI
Applicable Platforms	Bare Metal
Feature Default Setting	Enabled – Always On
Related Changes in this Release	Not Applicable
Related Documentation	<ul style="list-style-type: none"> • <i>UCC CEE Configuration and Administration Guide</i> • <i>UCC SMI Deployment Guide</i> • <i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2022.03.1

Feature Description

SMI Bare Metal utilizes the UCS M6 server for SMI Cluster Manager deployments.

For more information, see the *UCC SMI Deployment Guide*.

UCS Server Status Alert

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
---	--

Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC CEE Configuration and Administration Guide</i>

Revision History

Revision Details	Release
First introduced.	2022.03.1

Feature Description

If the UCS server is powered down or non-accessible, an alert will be set up to report and notify the UCS server availability status.

The SMI metrics track and report faults on the UCS server. The **cimc_server_not_reachable_alert** metric tracks the availability status of the UCS server. To establish an HTTP connection during login, this metric is set to 1 or 0 based on success (response) or failure.

vSphere Datacenter Folder Support

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	SMI
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<ul style="list-style-type: none"> <i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2022.03.1

Feature Description

SMI supports datacenters at the root level and within folders. If the datacenter is within a folder, then the entire path from the root until the datacenter is mentioned in the *datacenter-path* field. If the vSphere cluster is organised within folders, SMI can auto-detect the cluster as long as the name is unique with the datacenter.

For more information, see the *UCC SMI Operations Guide*.

Kubernetes 1.22 Upgrade

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Deployment Guide</i> <i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2020.02.2.3.10

Feature Description

With this release, the Kubernetes is upgraded from 1.21 to 1.22.

Parallel Node Upgrade with Deployment Zone Strategy

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
---	--

	K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2020.02.3.10

Feature Description

The current in-service upgrade strategy only supports upgrading one node at a time. For bigger clusters, more than six nodes, this upgrade strategy leads to longer upgrade periods, which mostly exceed the maintenance window (MW) limits.

This feature enables you to perform parallel upgrades for multiple nodes concurrently for faster in-service upgrades without impacting the availability and replication for any NF.

Architecture

The following images show the high-level design of the group upgrade flow for K8s and KVMs.

Figure 1: Upgrade Flow for K8s Clusters

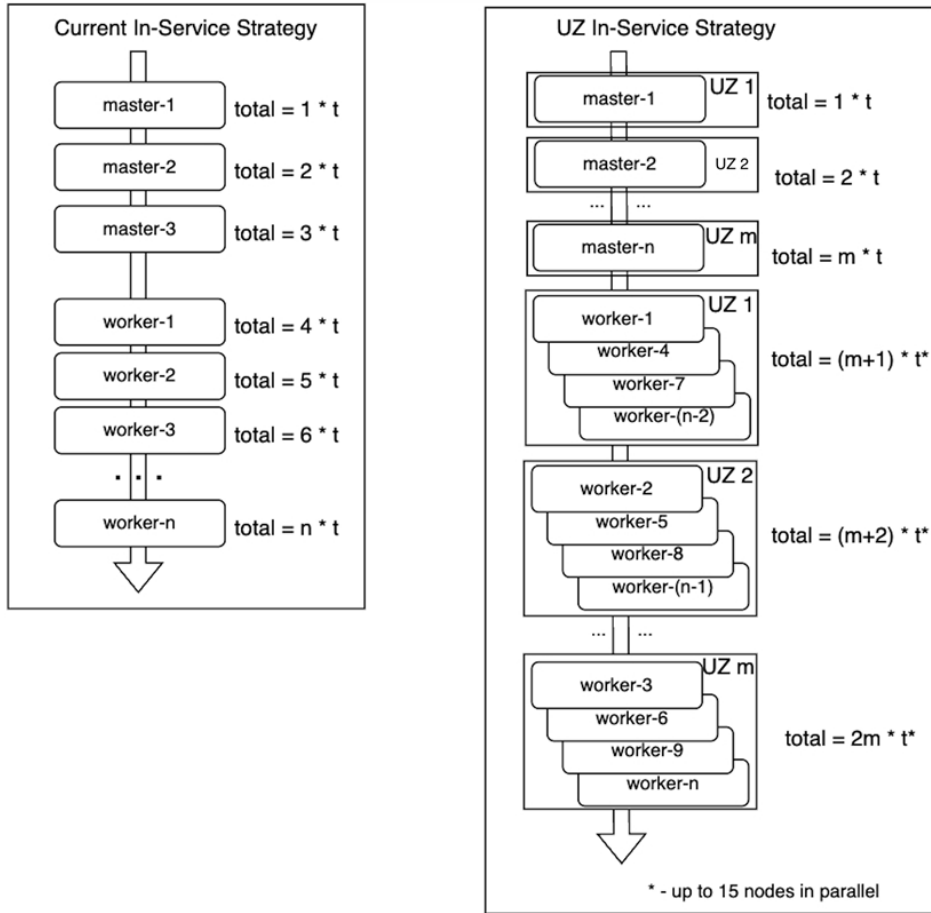
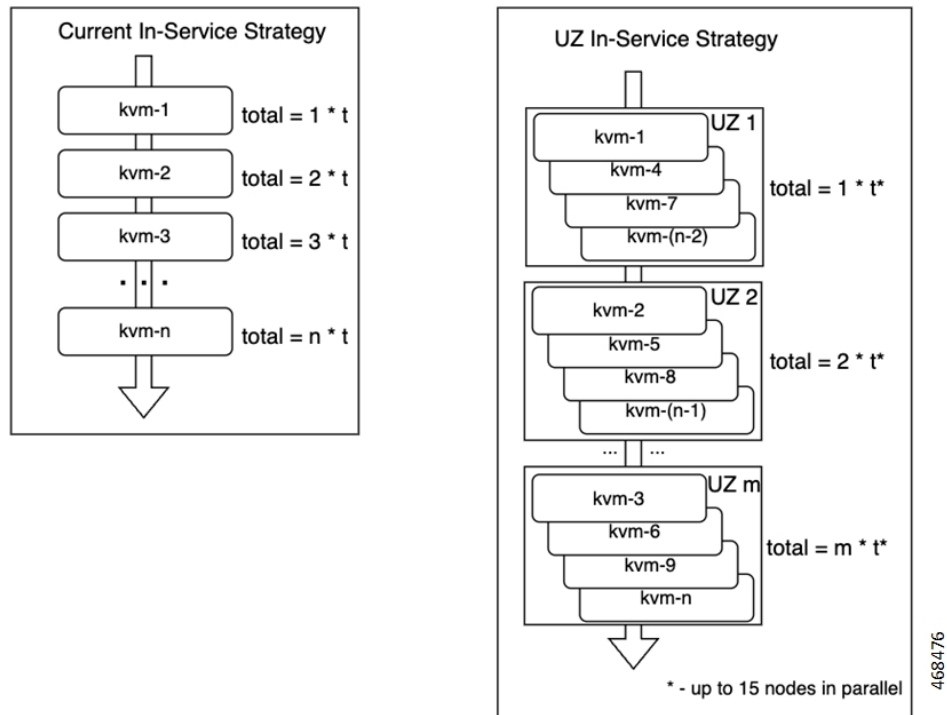


Figure 2: Upgrade Flow for KVM Clusters



How it Works

This section describes how the feature works.

This feature enables users to group servers into upgrade groups, which are similar to availability zones. The nodes in each upgrade group are upgraded in parallel (the maximum number of parallel nodes supported is 15).

The upgrade groups are upgraded in a sequential manner. For example, the control plane groups are not upgraded concurrently with the worker groups, but one at a time.

Requirements and Limitations

Some requirements and limitations associated with this feature are as follows.

- If the feature is disabled, the SMI reverts to the previous method of performing consecutive upgrade for the nodes.
- If the feature is enabled, the upgrade group configuration becomes mandatory for all nodes.

- There must be a majority of control plane or etcd nodes running at all time (for example, two out of three control planes should be always running).
- All the worker or KVM nodes should be distributed among the upgrade zone in a manner that ensures the majority of nodes never gets upgraded at the same time.
- The upgrade groups feature applies to the control plane, worker, KVM, and etcd node types, but doesn't apply to the CM-HA nodes.
- For the K8s clusters, the nodes include the upgrade group name as a new label. This label enables the NFs to use the affinity and anti-affinity rules to achieve proper HA and replication. The NFs can use the upgrade-zone provisioned node label or use custom defined ones to enable the application to align with the affinity rule.

Configuring the Deployment Zone Strategy

This section describes how to configure the upgrade groups for different nodes.

Use the following command to configure the upgrade groups for different nodes.

```
configuration enable-upgrade-zones true
  upgrade-zones zone_name
  exit
nodes node_name
  upgrade-zone zone_name
  exit
```



Note In this release, the zone upgrade strategy is applicable for only the **auto** option for cluster **upgrade-strategy**. See the following example configuration:

```
clusters foo actions sync run upgrade-strategy
Possible completions:
auto concurrent rolling
```

When **upgrade-strategy** is set to **auto** and calculated as **rolling**, Cluster Manager evaluates the upgrade zone configuration and performs a zone-based upgrade. If the **upgrade-strategy** is **auto** and calculated as **concurrent**, then it performs a concurrent upgrade regardless of the initial configuration.

Configuration Example:

```
clusters ott-bml-cl

configuration enable-upgrade-zones true

  upgrade-zones zone1
  exit
  upgrade-zones zone2
  exit
  upgrade-zones zone3
  exit

nodes mm1-controlplane1
  upgrade-zone zone1
  exit
nodes mm1-controlplane2
  upgrade-zone zone2
```

```

exit
nodes mm1-controlplane3
upgrade-zone zone3
exit
nodes mm1-etcd1
upgrade-zone zone1
exit
nodes mm1-etcd2
upgrade-zone zone2
exit
nodes mm1-etcd3
upgrade-zone zone3
exit
nodes mm1-worker1
upgrade-zone zone1
exit
nodes mm1-worker2
upgrade-zone zone2
exit
nodes mm1-worker3
upgrade-zone zone3
exit
nodes mm1-worker4
upgrade-zone zone1
exit
nodes mm1-worker5
upgrade-zone zone2
exit
nodes mm1-worker6
upgrade-zone zone3
exit
nodes mm1-worker7
upgrade-zone zone1
exit
commit
end

```

Path Based Routing for Inception Server

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2022.02.1

Feature Description

This feature enables the SMI to support path based URL routing for its nginx routing (external traffic) in Inception VM from the traditional host based approach for the following ingress.

- cli.smi-deployer.deployer.example.com
- restconf.smi-deployer.deployer.example.com

Configuring the Path Based Routing for Inception Server

This section describes how to enable the path based routing for Inception server.

Use the following argument in the deploy script to enable path based ingress for RESTCONF and CLI:

```
-i or --path-based-ingress
```

Configuration Example:

```
./deploy -p 209.165.200.224 -f Passwd@123 -i
```

After you enable the path based ingress, SSH and RESTCONF of inception server are accessible using the following URLs:

```
API: https://209.165.200.224/smi-deployer/restconf
```

If you provide the hostname in "--external-zone-name" along with the path based ingress argument, then the entire hostname is replaced with the provided host name.

Configuration Example:

```
./deploy -p 209.165.200.224 -f CSCO@123 -i --external-zone-name abc.com
```

After you enable the path based ingress, SSH and RESTCONF of inception server are accessible using the following URLs:

```
SSH (cli): ssh admin@127.0.0.1 -p 2022
```

```
API: https://abc.com/smi-deployer/restconf
```


CA Signed Certificate for Path-based Ingress

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
Added support for provisioning CA certificates.	2023.01.0
First introduced.	2022.02.1

Feature Description

This feature enables you to configure certificates signed by your own CA or external CA for path-based ingress URLs.

You can provision the certificates used for both REST APIs and K8s APIs through cluster manager and Ops-center. The recommended method to configure a certificate and its corresponding private key is to provision the certificate as a TLS secret using the existing yang container.

Certificate Expiry Check

The provisioned certificates must be monitored for expiry. The kube-certificate-expiring alert is automatically raised in advance to renew and update the certificate and key.

The alerts have the following severity levels:

- 30 days before expiry—Raise alert with Info severity
- 20 days before expiry—Raise alert with Major severity
- 15 days before expiry—Raise alert with Critical severity

Configuring Certificate for Path-based Ingress

This section describes how to configure TLS and CA certificates for path-based ingress.

Configuring TLS Certificate

Use the following procedure from cluster deployer to configure the certificates for path-based ingress.

1. Create a secret.

Use the following sample configuration to populate a certificate and its corresponding private key. The provided certificate and private key is stored as K8s TLS secret on the cluster under the mentioned namespace.

```
cluster cluster_name
  secrets tls namespace secret_name
    private-key private_key_content
    certificate certificate_content
  exit
exit
```

Example:

```
clusters sample-cluster
  secrets tls cee-global sample-secret
    private-key "$8$9n3U7OLEclVQoDpp/4VqkSLkeSmFbjx/
Mt6eEGN4EWOtPYlr9nqSWSZ40advmhDFsPFQZWfM\nhq/wpRzHXBZGp/
dNtNO+wpaQuxsT3CmkmRKFIHviUn4bEwBKfTCCsw7a5+66q3rm5vX4/nSw\
nNy4DrgTu4iFDzVYVKAYzoxWGzCqhKiaSqELjsW7gchEowC\n
  certificate "-----BEGIN CERTIFICATE-----\nMIID0zCCArugAw
IBAgIUPHTzpmTVUNVDQzJ/FM9tfCsAG2AwDQYJKoZIhvcNAQEL
\nBQAwaDELMAkGA1UEBhMCVVMxCzAJBgNVBAGMAkNBMQsw
\n-----END CERTIFICATE-----\n"
  exit
  exit
exit
```

2. Configure path-based ingress secret.

Use the following sample configuration to add the secret name for path-based ingresses.

```
clusters <cluster_name>
  ops-centers <opscenter_name> <instance_name>
    initial-boot-parameters path-based-ingress true
    initial-boot-parameters path-based-ingress-secret <secret_name>
  exit
  exit
exit
```



Note You must set **path-based-ingress** to **true** for getting the option to configure **path-based-ingress-secret**.

Example:

```
clusters sample-cluster
  ops-centers cee global
  initial-boot-parameters path-based-ingress true
  initial-boot-parameters path-based-ingress-secret sample-secret
```

```

    exit
  exit
exit

```

3. Run cluster sync to create and configure the secret as well as configure ingress to use the secret.

Verifying the Certificate for Path-based Ingress Configuration

This section describes how to verify the certificate for path-based ingress configuration.

Use the following CLI command to get the ingress in YAML and verify the configured secret name:

```
kubectl get ing -n <namespace> <ingress-name> -o yaml
```

Command Output Example:

```
cloud-user@sample-aio-controlplane:~$ kubectl get ing -n cee-global
cli-ingress-cee-global-ops-center -o yaml
```

```

apiVersion: networking.k8s.io/v1
kind: Ingress
...
spec:
  rules:
    - host: 10.x.x.x
      http:
        paths:
          - backend:
              service:
                name: ops-center-cee-global-ops-center
                port:
                  number: 7681
              path: /cee-global/cli
              pathType: ImplementationSpecific
      tls:
        - hosts:
            - 10.x.x.x.
          secretName: sample-secret

```

Run the **curl** command to verify the section "Server certificate:" to check whether the certificate is used properly.

```

cloud-user@satya-aio-controlplane:~$ curl -k -v https://10.x.x.x.nip.io/cee-global/cli
* Trying 10.x.x.x...
* TCP_NODELAY set
* Connected to 10.x.x.x.nip.io (10.x.x.x) port 443 (#0)
* ALPN, offering h2
* ALPN, offering http/1.1
* successfully set certificate verify locations:
*   CAfile: /etc/ssl/certs/ca-certificates.crt
*   CAPath: /etc/ssl/certs
* TLSv1.3 (OUT), TLS handshake, Client hello (1):
* TLSv1.3 (IN), TLS handshake, Server hello (2):
* TLSv1.3 (IN), TLS Unknown, Certificate Status (22):
* TLSv1.3 (IN), TLS handshake, Unknown (8):
* TLSv1.3 (IN), TLS Unknown, Certificate Status (22):
* TLSv1.3 (IN), TLS handshake, Certificate (11):
* TLSv1.3 (IN), TLS Unknown, Certificate Status (22):
* TLSv1.3 (IN), TLS handshake, CERT verify (15):
* TLSv1.3 (IN), TLS Unknown, Certificate Status (22):
* TLSv1.3 (IN), TLS handshake, Finished (20):
* TLSv1.3 (OUT), TLS change cipher, Client hello (1):
* TLSv1.3 (OUT), TLS Unknown, Certificate Status (22):
* TLSv1.3 (OUT), TLS handshake, Finished (20):

```

```

* SSL connection using TLSv1.3 / TLS_AES_256_GCM_SHA384
* ALPN, server accepted to use h2
* Server certificate:
* subject: C=US; ST=CA; L=SF; O=sample-signed.cisco.com; CN=10.x.x.x
* start date: Jul 12 04:19:56 2022 GMT
* expire date: Jul 11 04:19:56 2024 GMT
* issuer: C=US; ST=CA; L=SF; O=sample-signed.cisco.com; CN=10.x.x.x
* SSL certificate verify result: self signed certificate (18), continuing anyway.
* Using HTTP2, server supports multi-use
* Connection state changed (HTTP/2 confirmed)
* Copying HTTP/2 data in stream buffer to connection buffer after upgrade: len=0
* TLSv1.3 (OUT), TLS Unknown, Unknown (23):
* TLSv1.3 (OUT), TLS Unknown, Unknown (23):
* TLSv1.3 (OUT), TLS Unknown, Unknown (23):
* Using Stream ID: 1 (easy handle 0x56498909f550)
* TLSv1.3 (OUT), TLS Unknown, Unknown (23):
> GET /cee-global/cli HTTP/2
> Host: 10.x.x.x
> User-Agent: curl/7.58.0
> Accept: */*

```

Configuring CA Certificate

To configure the CA certificate, use the following configuration in Ops-center:

```

secrets ca-cert secret_name
  certificate certificate_content
exit

```

To configure the CA certificate, use the following configuration in cluster-manager:

```

cluster cluster_name
  secrets ca-cert namespace secret_name
    private-key private_key_content
    certificate certificate_content
  exit
exit

```

NOTES:

- If you add invalid certificate content and expired certificate, you will be prompted to correct the configuration.
- CA certificate is stored in generic (Opaque) secret type.
- The secrets are monitored and auto-healed if the user deletes the data by mistake.

OnDemand LDAP Connectivity Check

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
---	--

	K8s-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC CEE Configuration and Administration Guide</i> <i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2022.02.1

Feature Description

The SMI Ops Center provides an external authentication using LDAP support. The LDAP configuration can be configured in the SMI Ops Center using CLI or the RESTCONF APIs.

This feature enables you to validate a new LDAP configuration before adding it to the system or an existing LDAP configuration.

How it Works

This section describes how the feature works.

How to Validate a New Configuration

The steps to validate a new LDAP configuration are as follows.

1. Login to the SMI Ops Center.
2. Provide the LDAP new configuration inputs to validate (see the following example).

```
[pv/global] cee# smldap validate-security-config validate-new-security-config { ?
Possible completions:
base-dn          LDAP Base DN
bind-dn          LDAP Bind DN
group-attr       Group attribute
group-mapping    LDAP group to application security mapping
ldap-filter      LDAP Filter - use %s to sub username
ldap-server-url  LDAP Server URL (https://tools.ietf.org/html/rfc2255)
ldap-username-domain LDAP Username Domain
password         Password
username         Existing User name in LDAP server
```

3. Validate the LDAP new configuration (see the following example configuration).

```
cee(config)# smldap validate-security-config validate-new-security-config
{ base-dn dc=smi-lab,dc=com bind-dn cn=%s,ou=people,dc=smi-lab,dc=com group-attr
memberOf group-mapping { group admin ldap-group group1 } username user5 password
Passwd@123 ldap-filter cn=%s ldap-server-url ldap://209.165.200.224 }
```

```
Mon Jun 20 05:02:24.635 UTC+00:00
message accept "admin" external-user-group 1117 1117 /tmp
```

How to Validate an Existing LDAP Configuration

Use the following example configuration to validate an existing LDAP configuration.

```
cee# smildap validate-security-config validate-current-security-config
```

```
Mon Jun 20 05:07:41.765 UTC+00:00
```

```
Value for 'username' (<string>): user5
```

```
Value for 'password' (<string>): *****
```

```
message accept "admin" external-user-group 1117 1117 /tmp
```

Alerts for Node Disk Partition Usage

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	SMI Alerts
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always On
Related Changes in this Release	Not Applicable
Related Documentation	<ul style="list-style-type: none"> • <i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2022.01.1.02

Feature Description

In this release, two new alerts **node-disk-running-Low-24hours** and **node-disk-running-Low-2hours** are added to the **kube-prometheus-node-alerting.rules** rules group to notify the user when the node disk partition usage crosses the set thresholds.

Typically, the K8s performs a garbage collection when the disk usage is greater than 80%. In the new alerts, the node disk usage is set at a threshold of 75%. This way, the alerts are triggered before the K8s garbage collection.

This section describes the new alerts.

Alert	Description
node-disk-running-Low-24hours	The node disk partition is greater than 75%, and it will be greater than 80% in less than 24 hours.
node-disk-running-Low-2hours	The node disk partition is greater than 75%, and it will be greater than 80% in less than 2 hours.

Network Policy for K8s Pods

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	SMI
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always On
Related Changes in this Release	Not Applicable
Related Documentation	<ul style="list-style-type: none"> • <i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2022.01.1.02

Feature Description

In this release, the SMI provides a network policy for the K8s pods to secure communication between the pods.

Configuring the Network Policy for the K8s Pods

Use the following CLI configuration commands to enable or disable the network policy for the K8s pods. By default, the network policy is enabled.

```
clusterscluster_name
  configuration enable-network-policy { true | false }
```

Splitting Master and Additional Master VIPs into Separate VRRPs

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support K8s-based application deployment support
Applicable Platforms	Bare Metal
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Deployment Guide</i> <i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2022.01.1.02 2020.02.2.3.04

Feature Description

The CNDP allows you to configure an internal and external VIP as part of the cluster deployment. K8s uses the internal VIP while the ingress uses the external VIP to allow access to management interfaces such as Grafana, Ops-center, and Prometheus.

By default, both VIPs are part of one VRRP and failovers together. The VRRP verifies only the internal network for connectivity issues to prioritize the stability of the Kubernetes cluster. This can cause loss of connectivity to management interfaces if there's a network failure on the external network only.

This feature adds support for splitting the VIPs into different VRRP instances to allow them to failover independently.

To enable this feature, use the following configuration:

```
configuration separate-master-vip-vrrps true
```

You must configure **k8s additional-master-ip** for master nodes with the local external IP of the node that VRRP unicast uses.

For more information on deploying clusters, refer to the *UCC SMI Deployment Guide*.

SSH Firewall Rules for Cluster Nodes

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	SMI
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always On
Related Changes in this Release	Not Applicable
Related Documentation	<ul style="list-style-type: none"> • <i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2022.01.1.02

Feature Description

In this release, the SMI provides additional SSH firewall rules in the network policy for the cluster nodes for improved CM security.



Note It's recommended to set up an external firewall to secure access to the CM.

Configuring the SSH Firewall Rules in Network Policy

Use the following CLI configuration commands to enable or disable the SSH firewall rules. By default, the firewall rules are enabled.

```
clusterscluster_name
  configuration enable-ssh-firewall-rules { true | false }
```

Docker Subnet Override Support

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	cnBNG
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2020.02.7.07

Feature Description

By default, Docker uses the subnet range, 172.17.0.0/16 for container networking. If the same subnet range or an IP address from the range is already being used by some other resource in the same cluster environment, it might lead to a conflict.

This feature enables the user to configure and override the default value for the Docker subnet used by the SMI Cluster Manager (CM) or Inception VM. For the CM, this configuration is set by using the CM Ops-Center, whereas the Inception VM uses the `deploy.yaml` file to achieve the same configuration.

The `deploy.yaml` is enhanced to contain additional parameter, **configuration** with a sub-parameter, **docker-address-pools**. This YAML file contains a **base** for the CIDR range to use and a **size** for the size of the subnet to reserve for the new network.

Configuring the Docker Subnet Override

This section describes the configuration details for the Docker subnet override feature.

Use the following command to configure the Docker subnet override feature.

```
configuration docker-address-pools pool-name docker_bridge_address_pool_name [
base docker_bridge_subnet | size size ]
```

base *docker_bridge_subnet*

Specify the docker bridge subnet.

Must be a string in the ipv4-address-and-prefix-length pattern.

-Or-

Must be a string in the ipv6-address-and-prefix-length pattern.

Default Value: 172.17.0.0/16.

pool-name *docker_bridge_address_pool_name*

Specify the pool name of the docker bridge address pool.

Must be a string.

size *size*

Specify the size. For example, 16, 24, etc.

Must be an integer in the range of 8-24.

Default Value: 24.

cluster connect Command Update

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	SMI
Applicable Platforms	Bare Metal, OpenStack, VMware
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Revision Details	Release
First introduced.	2020.02.2.47 2020.02.7.07

Behavior Change Description

In this release, there is a change in the supported user-privilege level for using the **cluster connect** CLI command to access the containers in the cluster nodes.

Previous Behavior: The **cluster connect** CLI command enables any user to access all the pods and containers.

New Behavior: The **cluster connect** CLI command enables only users with administrator privileges to access only non-privileged containers.

IPSec Support for SMF N4 Interfaces

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
Added support for the following functionality: <ul style="list-style-type: none"> • IPSec Monitoring • Configuring IPSec certificates under strongSwan configuration 	2023.01.0
First introduced.	2020.02.2.47

Feature Description

This feature introduces strongSwan, a keying daemon, which uses the Internet Key Exchange (IKE) protocols, IKEv1 and IKEv2, to establish **security associations** (SA) between two peers in a network. Such an IKE session is denoted as **IKE_SA** in this chapter. The IKE provides strong authentication for both peers and derives unique cryptographic session keys. Besides authentication and key material, IKE also provides the means to exchange configuration information and to negotiate IPsec SAs, which are often called as **CHILD_SAs**. IPsec SAs define which network traffic is to be secured and how it has to be encrypted and authenticated.

The strongSwan feature is available as an add-on from the Cluster Manager (CM). Use the CM Ops-Center to configure this add-on. In the current release, the SMI uses strongSwan version 5.9.3.

SMI allows monitoring of IPSec certificates—sends certificate expiry alerts and updates certificate through strongSwan configuration.

Configuration Parameters

In this section, see the description for different configuration parameters available for the strongSwan add-on feature. Use the CM Ops-Center to configure these parameters.

- **name**: Specifies the name of the connection, which can be used for connection specific operations, for example, up or down.
- **auto { ignore | add | route | start }**: Specifies the operation, if any, that should be automatically performed at IPsec startup. The **add** option loads a connection without starting it, whereas **route** loads a connection and installs kernel traps. If traffic is detected between the leftsubnet and rightsubnet, a connection is established. The **start** option loads a connection and brings it up immediately. The **ignore** option ignores the connection and is the same as deleting a connection from the config file. The default value is **ignore**.
- **keyexchange { ikev1 | ikev2 }**: Specifies the method of key exchange and the protocol to use to initialize the connection.
- **type { tunnel | transport | transport_proxy | passthrough | drop }**: Specifies the type of the connection. Currently, the accepted values are **tunnel**, signifying a host-to-host, host-to-subnet, or subnet-to-subnet tunnel. The **transport** option signifies a host-to-host transport mode, whereas the **transport_proxy** option signifies the special Mobile IPv6 transport proxy mode. The **passthrough** option signifies that no IPsec processing should be done at all and **drop** signifies that packets must be discarded.
- **left** or **right { ip address ip_address | fqdn fqdn | %any | %any4 | %any6 | range | subnet }**: Specifies the IP address or FQDN of the participant public-network interface. The value **%any** for the local endpoint signifies an address to be filled in (by automatic keying) during negotiation. If the local peer initiates the connection setup, then the routing table is queried to determine the correct local IP address. If the local peer is responding to a connection setup, then any IP address that is assigned to a local interface is accepted. The value **%any4** restricts address selection to IPv4 addresses and **%any6** restricts address selection to IPv6 addresses.
- **leftsubnet** or **rightsubnet ip subnet**: Specifies the private subnet behind the left participant, expressed as either network or netmask.
- **leftid** or **rightid id**: Specifies how the left or right participant must be identified for authentication. The default values are left or right or the subject of the certificate configured. It must match the full subject DN or one of the subjectAltName extensions contained in the certificate.
- **leftsendcert { never | no | ifasked | always | yes }**: Defines whether a peer must send a certificate request (CR) payload in order to get a certificate in return.
- **leftauth** or **rightauth { pubkey | psk | eap | xauth }**: Specifies the authentication method to use locally (left) or require from the remote (right) side. The acceptable values are **pubkey** for public key encryption (RSA/ECDSA), **psk** for pre-shared key authentication, **eap** to use the Extensible Authentication Protocol, and **xauth** for IKEv1 eXtended Authentication.
Pubkey is the default option.
- **psk pre-shared key**: Specifies the required setting if leftauth or rightauth is configured as **psk**.
- **esp { cipher suites | aes128-sha256 }**: A comma-separated list of ESP encryption or authentication algorithms is used for the connection, for example, **aes128-sha256**. The notation is encryption-integrity[-dhgroup][-esnm]. For IKEv2, multiple algorithms (separated by -) of the same type can be included in a single proposal. IKEv1 only includes the first algorithm in a proposal.
aes128-sha256 is the default option.
- **ike { cipher suites | aes128-sha256-modp3072 }**: A comma-separated list of IKE/ISAKMP SA encryption or authentication algorithms is used, for example, **aes128-sha256-modp3072**.

The notation is encryption-integrity[-prf]-dhgroup. In IKEv2, multiple algorithms and proposals might be included, such as aes128-aes256-sha1-modp3072-modp2048 or 3des-sha1-md5-modp1024.

- **ikelifetime** { **time** *time* | **3h** }: Specifies how long the keying channel of a connection (ISAKMP or IKE SA) must last before being renegotiated.
- **lifetime** { **time** *time* | **1h** }: Specifies how long a particular instance of a connection should last, from successful negotiation to expiry.
- **dpdaction** { **none** | **clear** | **hold** | **restart** }: Specifies the action to be taken when dead peer is detected.
none is the default value.
- **dpddelay** { **time** *time* | **30s** }: Defines the period time interval with which INFORMATIONAL exchanges are sent to the peer. These are only sent if no other traffic is received.
- **dpdtimeout** { **time** *time* | **150s** }: Defines the timeout interval after which, all the connections to a peer are deleted in case of inactivity.
- **inactivity time** *time*: Defines the timeout interval after which, a CHILD_SA is closed if it did not send or receive any traffic.
- **closeaction** { **none** | **clear** | **hold** | **restart** }: Defines the action to take if the remote peer unexpectedly closes a CHILD_SA (see **dpdaction** for the description of different options). If the peer uses reauthentication or uniqueids checking, **closeaction** must not be used, these events might trigger the defined action when it's not desired.
- **nodes** *list_of_node_names*: Specifies the node names on which IPsec connection must be established.
- **serverCert** *server_certificate*: Specifies the content of Server certificate in the **pem** format to be used for this connection.



Note This keyword is not supported under strongSwan configuration.

- **serverPrivKey** *server_private_key*: Specifies the content of server private key in the **pem** format to be used for this connection.



Note This keyword is not supported under strongSwan configuration.

- **serverPrivKeyPassphrase** *passphrase*: Specifies the passphrase used to encrypt the **server-priv-key** value.
- **server-secret**: Pass an existing TLS secret for this connection.

Installing strongSwan

This section describes how to install the strongSwan feature.

Install strongSwan as an Add-on from the CM

Use the following steps to install strongSwan as an add-on from the CM Ops-Center:

1. Use the following CLI commands to enable the strongSwan add-on:

```
clusters cluster_name addons strongswan enabled
```

2. Set all the strongSwan parameters for **connection** (refer to the *Configuration Parameters* section for more details on available parameters).
3. Trigger the cluster sync operation.



Note The strongSwan pods run on all the nodes, however traffic is accepted only on those nodes, which are configured by using the "nodes" parameter in the CM Ops-Center. strongSwan does not accept or send any traffic on non-configured nodes.

Configuring IPsec Certificates

To configure IPsec certificates under strongSwan configuration, use the following procedure:

1. Create TLS associated secret for server and CA certificate.

Note: Create strongSwan-related secrets inside the smi-strongswan namespace.

Example:

```
[test-cm-controlplane] SMI Cluster Deployer# show running-config clusters secrets ca-cert
clusters test-aio
secrets ca-cert smi-strongswan 134-ca
certificate "-----BEGIN
CERTIFICATE-----\nMIIDqzCzQubm.....1Ac1L+s4M3ug==\n-----END
CERTIFICATE-----\n"
exit
secrets ca-cert smi-strongswan 135-ca
certificate "-----BEGIN
CERTIFICATE-----\nMIIFqzCCA5Og.....9XdMDiQANHg7w\n-----END
CERTIFICATE-----\n"
exit
secrets ca-cert smi-strongswan ca-1
certificate "-----BEGIN
CERTIFICATE-----\nMIID0TCCArmG.....UNvF0nAmIX0qxg4\n-----END
CERTIFICATE-----\n"
exit
secrets ca-cert smi-strongswan ca-2
certificate "-----BEGIN PRIVATE
KEY-----\nMIIEvQIBADAN.....tbNDzGANf29nus=\n-----END PRIVATE KEY-----\n"
exit
exit
```

2. Refer the secrets in strongSwan configuration. The strongSwan configuration shows the available TLS and certificates.

Example:

```
[test-cm-controlplane] SMI Cluster Deployer# show running-config clusters karan-aio
strongswan connections server-secret
clusters test-aio
strongswan connections a-to-b
server-secret a-to-b
exit
exit
```

```
[test-cm-controlplane] SMI Cluster Deployer# show running-config clusters karan-aio
strongswan ca-certs
clusters test-aio
strongswan ca-certs [ 134-ca 135-ca ]
exit
```

Push KPIs to S3 Using Thanos

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Common Execution Environment - Configuration and Administration Guide</i>

Revision History

Revision Details	Release
First introduced.	2020.02.2.47

Feature Description

In this feature, the CEE provides you the option to backup the local data stored in Prometheus to a remote storage object, for example, Amazon Web Services (AWS) S3, by using Thanos.

This feature provides the following two deployment models:

- Thanos Sidecar
- Thanos Receive

How it Works

Thanos with Sidecar

This section describes how to configure the Sidecar deployment with AWS S3.

Prerequisites

- S3 bucket in AWS



Note For more information about how to create an AWS S3 bucket, refer to the original product documentation.

Configuring the Sidecar

Use the following sample CLI commands in the CEE Ops-Center to set up the Sidecar deployment.

```
prometheus thanos-s3-object-store bucket zx-thanos-test
prometheus thanos-s3-object-store endpoint s3.us-east-1.amazonaws.com
prometheus thanos-s3-object-store access-key
prometheus thanos-s3-object-store secret-key
```

Thanos with Receive

This section describes how to configure the Remote-write target including the Receiver URL and enable TLS support for the same using the CEE Ops-Center for the Receive deployment with AWS S3.

Prerequisites

- S3 bucket in AWS
- Deploy Thanos Recieve



Note For more information about how to create an AWS S3 bucket, refer to the original product documentation.

Configuring the Remote Write Target with Receiver URL

Enter the URL of the Thanos Receiver in the CEE Ops-Center CLI.

A sample configuration for Prometheus to work with Thanos Receive with an HTTP endpoint is shown below.

```
[user/global] cee# config
Entering configuration mode terminal
[user/global] cee(config)# prometheus remote-write target demo
[user/global] cee(config-target-demo)# url http://thanos-receive-hi-res:10000/api/v1/receive
[user/global] cee(config-target-demo)# commit
Fri Dec 10 04:28:29.838 UTC+00:00
Commit complete.
[user/global] cee(config-target-demo)#
Message from confd-api-manager at 2021-12-10 04:28:31...
Helm update is STARTING. Trigger for update is CHANGE.
```

Configuring the Remote Write Target with TLS Enabled

Remote write to Thanos Receive or any other target with TLS enabled is also supported. You can input the necessary ca/cert/key file by using the CEE Ops-Center CLI.

A sample configuration about how to configure remote-write target with TLS enabled is shown below. This configuration enables you to configure Prometheus to work with Thanos Receive with an HTTPS endpoint.

Assume the target remote server has a self-signed server and user has the CA certificate for it.

```
[user/global] cee(config)# prometheus remote-write target demo
Fri Dec 3 20:58:39.735 UTC+00:00
[user/global] cee(config-target-demo)# url https://thanos-receive-hi-res:10908/api/v1/receive
Fri Dec 3 20:58:51.609 UTC+00:00
[user/global] cee(config-target-demo)# tls-config tls-
Possible completions:
  tls-ca      CA certificate to validate API server certificate with.
  tls-cert    Certificate file for client cert authentication to the server.
  tls-key     Key file for client cert authentication to the server.
[user/global] cee(config-target-demo)# tls-config tls-ca
Fri Dec 3 20:59:05.384 UTC+00:00
(<AES encrypted string>):
[Multiline mode, exit with ctrl-D.]
> *****
> *****
> *****
> *****
[user/global] cee(config-target-demo)# tls-config skip-verify
Possible completions:
  false true
[user/global] cee(config-target-demo)# tls-config skip-verify false
Fri Dec 3 20:59:40.188 UTC+00:00
[user/global] cee(config-target-demo)# commit
Fri Dec 3 20:59:42.797 UTC+00:00
Commit complete.
```

After the configuration, the Prometheus `remote_write` is configured as follows and the CA certificate from user input is created on the shown path in the Prometheus container.

```
remote_write:
- tls_config:
  ca_file: /etc/remote-write-certs-shared/demo-ca
  insecure_skip_verify: false
  url: https://thanos-receive-hi-res:10908/api/v1/receive
```

virsh console Command Update

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2020.02.2.47

Behavior Change Description

In this release, there is change in the CLI configuration command used for connecting to the UPF node through a KVM console post synchronization while installing a KVM and User Plane VM.

Previous Command:

```
virsh console <upf_name> serial1
```

New Command:

```
virsh console <upf_name> serial1 --force
```

Alert for Standby Cluster Manager Failure

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2020.02.2.41

Feature Description

In this release, new alert **backup-node-down** is added to the **cndp-ha** rules group to notify the user of any standby or secondary cluster manager failures. During a secondary cluster manager failure, this alert helps you to take necessary steps to create backup of the primary cluster manager to avoid a single point of failure. It constantly monitors the DRBD status of the HA cluster and alerts the user when there is a failure.

Configuring the Alert for Standby Cluster Manager Failure

Use the following CLI commands or configuration in the CEE Ops Center to enable or disable the DRBD monitoring feature:

```
config
  cm monitoring { true | false }
```

Ubuntu User Password Expiration Configuration Enhancement

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<ul style="list-style-type: none"> • <i>UCC SMI Operations Guide</i> • <i>UCC CEE Configuration and Administration Guide</i>

Revision History

Revision Details	Release
First introduced.	2020.02.2.41

Feature Description

In this release, the Cluster Sync operation updates the value of the password expiration days for the default user. The default user password expiration alert is enabled to monitor the password expiry. A critical alert is triggered when the password is about to expire in 30 days.

Unified RMA Procedure for Planned and Failure Events on Bare Metal

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal
Feature Default Setting	Enabled – Always On
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2020.02.2.41

Feature Description

In this release, the RMA (Return Merchandise Authorization) procedure for planned maintenance and any unplanned node failure events is unified for the SMI Bare Metal stacked cluster. The same procedure is applicable for both the master and worker nodes. This feature simplifies automation and MOP requirements for both the NSO and the user.

For more information on the unified RMA procedure on the SMI Bare Metal, refer to the SMI Cluster RMA section in the *UCC SMI Operations Guide*.

Alert for KVM Node Unreachable

Feature Summary and Revision History

Summary Data

Table 1: Summary Data

Applicable Product(s) or Functional Area	SMI Alerts
--	------------

Applicable Platforms	Bare Metal
Default Setting	Enabled—Always On
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Table 2: Revision History

Revision Details	Release
First introduced	2020.02.2.19
Note: This feature has not been fully qualified for this release.	

Feature Description

The following new alert has been added to the Node Network rules group to indicate when a KVM node is unreachable:

Alert: `kvm-node-not-ready`

Annotations:

Type: Communications Alarm

Summary: KVM node {{ \$labels.hostname }}({{ \$labels.ip}}) is not reachable.

Expression:

```
|
changes(kvm_metrics_tunnels_up[2m]) > 0
```

For: 0m

Labels:

Severity: major

For more information on the SMI Cluster RMA, refer to the [UCC SMI Operations Guide – Alerts Reference](#) chapter.

CDL Data Slicing

Feature Summary and Revision History

Summary Data

Table 3: Summary Data

Applicable Product(s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Default Setting	Disabled—Configuration Required

Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC CDL Configuration and Administration Guide</i>

Revision History

Table 4: Revision History

Revision Details	Release
First introduced	2020.02.2.19

Feature Description

Data slicing logically separates CDL as slices and stores the session data based on the slice name received from the Network Functions (NF).

With data slicing, one or more NFs can store different types of session data in dedicated slices of CDL. A default slice name called **session** is used if slice names are not configured.

```
cdl datastore <datastore name> slice-names [ <sliceName 1><sliceName 2> ... <sliceName n> ]
```

Sample configuration is as follows:

```
cdl datastore session slice-names [ session1 session2 ]
```



Note

- If the slice names are configured at the NF's ops-center or CDL's ops-center, every request from the NF must have a valid slice name. If the slice name is different from what is configured or empty, then the request is rejected with an error code.
- If the slice names are not configured, then the NF requests are routed to the default session
- The slice names cannot be updated in a running system post deployment.

For more information, refer to the [UCC CDL Configuration and Administration Guide – Common Data Layer](#) chapter.

CDL Overload Protection Enhancement

Feature Summary and Revision History

Summary Data

Table 5: Summary Data

Applicable Product(s) or Functional Area	KVM-based application deployment support
--	--

Applicable Platforms	Bare Metal, OpenStack, VMware
Default Setting	Disabled—Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC CDL Configuration and Administration Guide</i>

Revision History

Table 6: Revision History

Revision Details	Release
First introduced	2020.02.2.19

Feature Description

With this release, the existing CDL overload protection configuration is replaced with a new configuration.



Note In the previous release (Oct 2020/Jan 2021), the overload protection was enabled by default with hard coded limits of 1 million records per slot, 10 million records per index and 5 gb record size per slot. However, in this release, the feature is disabled by default and requires configuration to enable this feature and to configure the appropriate overload limits (optional) for each parameter.

Command	Changes
Old command: <code>cdl datastore session overload-protection disable true</code>	The old CDL overload protection command is deprecated. Note: The command is not functional but is available only for backward compatibility.
New command: <code>cdl datastore session features overload-protection enable <true/false></code>	The overload protection configuration is moved under <code>cdl datastore session features</code> configuration to configure overload-protection and alerts. CDL is now equipped to configure: <ul style="list-style-type: none"> • record-capacity per slot/index • record-capacity in bytes per slot • configure major and critical alarm %

If the overload protection is enabled then the alerts are also enabled. The `prometheus-rules-cdl` pod is spawned. If overload-protection is disabled then the alerts are disabled. The `prometheus-rules-cdl` pod is removed.

Configuring Overload Protection Parameter

The following parameters are configured to set limits for overload-protection:


```

cdl datastore session features overload-protection index-max-record-count <value>
cdl datastore session features overload-protection slot-max-record-count <value>
cdl datastore session features overload-protection slot-max-size <value>
cdl datastore session features overload-protection hard-limit-percentage <value>

```

The table below lists the configuration details:

CLI Command	Description
cdl datastore session features overload-protection enable <true/false>	(Optional) CDL overload-protection is disabled by default. The default value is false.
cdl datastore session features overload-protection index-max-record-count <value>	(Optional) Maximum number of records that can be stored in the Index shard. The default value is 60000000 (60M). The range is 100k to 100M. Note: The range from 100 to 1000 is applicable only for testing in the lab environment. It is not recommended for production environment.
cdl datastore session features overload-protection slot-max-record-count <value>	(Optional) Maximum number of records that can be stored in Slot shard. The default value is 2500000 (2.5M). The range is either 100 or 100k to 10M. Note: The value 100 is applicable only for testing in the lab environment. It is not recommended for production environment.
cdl datastore session features overload-protection slot-max-size <value>	(Optional) Maximum size of Slot shard in mega bytes. The default value is 16384 (16GB). The range is 1GB to 96GB.
cdl datastore session features overload-protection hard-limit-percentage <value>	(Optional) Additional capacity (percentage) in addition to the soft limit. This is used to determine when to reject the update requests at CDL endpoint. For eg: if index shard = 1, index-record-capacity = 100 and hard-limit-percentage = 5, then the create requests are rejected when number of index records = 100 and update requests are rejected only when it reaches 105. The default value is 5. The range is 0-10.
cdl datastore session features overload-protection major-alert-threshold <value>	(Optional) Threshold (percentage) at which CDL triggers an alert <code>cdlOverloaded-major</code> . The default value is 80. The range is 40-100.

CLI Command	Description
<code>cdl datastore session features overload-protection critical-alert-threshold <value></code>	(Optional) Threshold (percentage) at which CDL triggers an alert <code>cdlOverloaded-critical</code> . The default value is 90. The range is 40-100.

Configuring the alert percentage

Run the following to configure the threshold percentage:

```
cdl datastore session features overload-protection critical-alert-threshold <percentage>
cdl datastore session features overload-protection major-alert-threshold <percentage>
```

For more information, refer to the [UCC CDL Configuration and Administration Guide – Common Data Layer](#) chapter.

CLI Support for UPF IFTASK Forwarder Type

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	UCC SMI Operations Guide

Revision History

Revision Details	Release
First introduced.	2020.02.2.19

Feature Description

In past releases, SMI provided support for the Vector Packet Processing (VPP) data plane development kit (DPDK) forwarder for use with StarOS-based applications such as the 5G User Plane Function.

SMI now supports the use of the DPDK Internal Forwarder (IFTASK) for use with StarOS-based applications.

IFTASK support is enabled by the forwarder-type parameter as part of the day 0 configuration:

```
nodes master
```

```
[ no ] vm-defaults upf day0 forwarder-type { IFTASK | VPP }
```

If the forwarder type command is not issued, then VPP will be used as the forwarder type by default.

Once IFTASK has been set, the no variant of this command can be used to re-enable the VPP forwarder type.

When the forwarder type is set to IFTASK, the following additional parameters are used as part of the UPF day 0 configuration:

- IFTASK_SERVICE_TYPE=0
- IFTASK_CORES=44
- IFTASK_MCDMA_CORES=50

These parameters are hard-coded and set automatically by SMI during the deployment process.

For more information on deploying and monitoring KVM-based applications, refer to the [UCC SMI Operations Guide > SMI Cluster Manager Operations](#) chapter.

Cluster Manager Now Uses Internal Network for HA Communications

Feature Summary and Revision History

Summary Data

Table 7: Summary Data

Applicable Product(s) or Functional Area	Cluster Manager deployment
Applicable Platforms	Bare Metal, OpenStack, VMware
Default Setting	Disabled—Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i> <i>UCC SMI Deployment Guide</i>

Revision History

Table 8: Revision History

Revision Details	Release
First introduced	2020.02.2.19

Feature Description

Cluster Manager (CM) high-availability (HA) is provided through keepalived and Distributed Replicated Block Device (DRBD).

Prior SMI releases used the externally routable ssh-ip address to configure keepalived and DRBD communications between the active and standby CM HA nodes. This model left potential for a split-brain situation should the externally routable network become unstable or unavailable.

To reduce this potential, the CM HA nodes can now be configured to use the internal network for keepalived and DRBD communications.

This is done using the following command in the CM configuration file:

```
nodes <node_name>
    cm ha-ip <internal_address>
```

Below is an example of a configuration excerpt identifying the parameters for configuring internal and external addresses:

```
# The master-virtual-ip parameter contains the *internal* VIP address.
configuration master-virtual-ip 192.0.1.101
configuration master-virtual-ip-cidr 24
configuration master-virtual-ip-interface vlan1001
#
# The additional-master-virtual-ip parameter contains the details of the *externally*
available VIP address.
configuration additional-master-virtual-ip 203.0.113.214
configuration additional-master-virtual-ip-cidr 26
configuration additional-master-virtual-ip-interface vlan3540
#
#The additional cm ha-ip parameter needs to be added with the *internal* IP of the node.
#
```



Note node-ip in a CM HA config points to the internal master-virtual-ip

```
nodes cm1
    ssh-ip 203.0.113.212
    type k8s
    k8s node-type master
    k8s node-ip 192.0.1.101
    cm ha-ip 192.0.1.59
    ...
    initial-boot netplan vlans vlan3540
        addresses [ 203.0.113.212/26 ]
    exit
    os netplan-additions ethernet eno1
        addresses [ 192.200.0.29/8 ]
    exit
    os netplan-additions vlans vlan1001
        addresses [ 192.0.1.59/24 ]
    exit
    exit

nodes cm2
    ssh-ip 203.0.113.213
    type k8s
    k8s node-type backup
    k8s node-ip 192.0.1.101
    cm ha-ip 192.0.1.60
    ...
```

```

initial-boot netplan vlans vlan3540
addresses [ 203.0.113.213/26 ]
exit
os netplan-additions ethernet eno1
addresses [ 192.200.0.29/8 ]
exit
os netplan-additions vlans vlan1001
addresses [ 192.0.1.60/24 ]
exit
exit

```

For more information Cluster Manager HA deployments, refer to the [UCC SMI Deployment Guide > SMI Cluster Manager – Deployment](#) chapter.

Configurable Option to Control Ping Properties

Feature Summary and Revision History

Summary Data

Table 9: Summary Data

Applicable Product(s) or Functional Area	KVM-based application monitoring
Applicable Platforms	Bare Metal, OpenStack, VMware
Default Setting	Disabled—Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Table 10: Revision History

Revision Details	Release
First introduced	2020.02.2.19

Feature Description

SMI monitors applications such as the UPF that are deployed in KVM-based clusters. Monitoring, in part, is performed through pings sent at regular intervals to verify that the applications are alive. If an application is unresponsive for a certain number of times, then the monitoring service attempts to restart it.

The ping interval and the number of failure occurrences are now configurable using the following commands:

```

node-defaults kvm monitoring ping-interval <#_seconds>
node-defaults kvm monitoring failure-occurrence <#_instances>

```

<#_seconds> is any integer value. The minimum value is 3 and the default value is 10 seconds.

<#_instances> is any integer value. The minimum value is 3 and the default value is 10 times.

Upon changing this values and running a sync, the monitoring service is restarted and the new configuration is applied.

For more information on deploying and monitoring KVM-based applications, refer to the [UCC SMI Operations Guide > SMI Cluster Manager Operations](#) chapter.

Deleting Stale CDL Slot Data

Feature Summary and Revision History

Summary Data

Table 11: Summary Data

Applicable Product(s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Default Setting	Disabled—Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC CDL Configuration and Administration Guide</i>

Revision History

Table 12: Revision History

Revision Details	Release
First introduced	2020.02.2.19

Feature Description

In certain scenarios, the CDL records are found on Slot but not in the index pods. The notifications from Slot towards the application for such records do not receive the values correctly. The record in the slot is not deleted, if the index data is not deleted.

Ensure the following before deleting the CDL Slot Data:

- If the number of notifications to an application crosses a threshold value (default value of 3), a record is suspected to be stale.
- This triggers a validation check to find the corresponding record in any of the index pods (local or on any geo remote sites).
- If there is a mismatch in map ID from index, or if the map ID is not found in all index pods, then a clean-up is invoked to delete the record on local as well as remote sites.

The following parameters are introduced to delete stale records:

`disable-auto-deletion`: When set to true, the stale CDL records are not deleted. Auto deletion of stale records is enabled by default.

`notification-retry-count`: Specifies the minimum number of timer expiry notification retries sent to application without receiving an update from application. If there are no updates received even after `notification-retry-count` times, cdl proceeds to check if slot record is stale. The default number is 3.

The sample CDL configurations are as follows:

To disable the stale slot record auto deletion feature:

```
cdl datastore session
features slot-stale-session-detection disable-auto-deletion true
exit
```

You can change the `notification-retry-count` to a new value, for example 5. This indicates that the timer expiry notification tries 5 times, after which it proceeds for checking whether the data is stale.

```
cdl datastore session
features slot-stale-session-detection notification-retry-count 5
exit
```

Troubleshooting

To enable troubleshooting logs for this feature on endpoint and slot pods use the following configuration:

```
cdl logging logger ep.staleRecord.session
level info
exit
cdl logging logger slot.staleRecord.session
level info
exit
```

For more information, refer to the [UCC CDL Configuration and Administration Guide – Common Data Layer](#) chapter.

Dual Stack Support

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	UCC SMI Deployment Guide

Revision History

Revision Details	Release
First introduced.	2020.02.2.19

Feature Description

Dual stack enables networking devices to be configured with both IPv4 and IPv6 addresses. By setting the `ipv6-mode` to *dual-stack*, you can configure the Kubernetes and CM HA with IPv6 address.

For more information, refer to the [UCC SMI Deployment Guide – SMI Cluster Manager – Deployment](#) chapter.

Emulator Pinning

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	UCC SMI Deployment Guide

Revision History

Revision Details	Release
First introduced.	2020.02.2.19

Feature Description

SMI provides capability to set the emulator pinning during the VM CPU allocation.

For more information, refer to the [UCC SMI Operations Guide > SMI Cluster Manager Operations](#) chapter.

GR Failover Notifications

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Cisco Common Data Layer Configuration Guide

Revision History

Revision Details	Release
First introduced.	2020.02.2.19

Feature Description

The CDL is equipped with Geo Replication (GR) failover notifications, which can notify the timer expiry of session data and bulk notifications to the currently active site.

For information, refer to [UCC CDL Configuration and Administration Guide > Cisco Common Data Layer > CDL Geo Replication \(GR\) Deployment > Deploying CDL for Geo Replication \(GR\)](#).

Hostname and URL Path-Based Routing for Ingress

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	UCC SMI Deployment Guide

Revision History

Revision Details	Release
First introduced.	2020.02.2.19

Feature Description

SMI now supports configuring the ingress traffic with Fully Qualified Domain Names (FQDN) and Path-based URL routing to connect to an ops-center. The different ingresses are accessible through subdomain-based names and path-based names.

For more information, refer to the [UCC SMI Cluster Manager Deployment Guide > SMI Cluster Manager - Deployment](#) chapter.

Increased NotReady Detection Sensitivity for K8s Nodes

Feature Summary and Revision History

Summary Data

Table 13: Summary Data

Applicable Product(s) or Functional Area	SMI Alerts
Applicable Platforms	Bare Metal, OpenStack, VMware
Default Setting	Enabled—Always On
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Table 14: Revision History

Revision Details	Release
First introduced	2020.02.2.19
Note: This feature has not been fully qualified for this release.	

Feature Description

In previous SMI releases, the `k8s-node-not-ready` alert was triggered when a K8s node remains in the NotReady state for five minutes.

With this release, the `k8s-node-not-ready` alert triggers when a K8s node remains in the NotReady state for one minute.

For more information on the SMI Cluster RMA, refer to the [UCC SMI Operations Guide > Alerts Reference](#) chapter.

Kubernetes 1.20.0 Upgrade

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	UCC SMI Operations Guide

Revision History

Revision Details	Release
First introduced.	2020.02.2.19

Feature Description

With this release, the Kubernetes is upgraded from 1.19.0 to 1.20.0.

For more information, refer to the [UCC SMI Operations Guide](#).

Node Failure Notifications During RMA

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal
Feature Default Setting	Enabled – Always On

Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Revision Details	Release
First introduced.	2020.02.2.19

Feature Description

The SMI Cluster RMA procedure is enhanced to provide notifications to the NSO.

A notification is sent to the NSO, when a node becomes unresponsive due to unplanned node(s) outage and requires it to be removed from the cluster before putting it in maintenance mode. The NSO must subscribe to the notification stream **alert notification** to receive the *k8-node-not-ready* alert.

For more information on the SMI Cluster RMA, refer to [UCC SMI Operations Guide > Operating the SMI Cluster Manager on Bare Metal > SMI Cluster RMA](#).

Silence 'Always On' vm-alive Alerts

Feature Summary and Revision History

Summary Data

Table 15: Summary Data

Applicable Product(s) or Functional Area	SMI Alerts
Applicable Platforms	Bare Metal
Default Setting	Enabled—Always On
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Table 16: Revision History

Revision Details	Release
First introduced	2020.02.2.19
Note: This feature has not been fully qualified for this release.	

Feature Description

In previous SMI releases, the system would monitor for VM liveness every five seconds and trigger corresponding `vm-alive` alerts.

In this release, VM liveness is still monitored every five seconds, however, the alert goes away after VM liveness is stable for two minutes.

For more information, refer to the [UCC SMI Operations Guide > Alerts Reference](#) chapter.

Smart Agent Upgrade

Feature Summary and Revision History

Summary Data

Table 17: Summary Data

Applicable Product(s) or Functional Area	Smart Licensing
Applicable Platforms	Bare Metal, OpenStack, VMware
Default Setting	Enabled—Always On
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Common Execution Environment - Configuration and Administration Guide</i>

Revision History

Table 18: Revision History

Revision Details	Release
First introduced	2020.02.2.19

Feature Description

The Smart Agent is a software component within the Smart Licensing capability supported on SMI. It functions in the `smart-agent-cee-global-ops-center` pod.

With this release, the Smart Agent has been upgraded from 3.0.13 to 3.1.4.

For more information on configuring and using Smart Licensing in SMI, refer to the [UCC CEE Configuration and Administration Guide > Common Execution Environment](#) chapter.

Support VM Status Alerts on CNDP

Feature Summary and Revision History

Summary Data

Table 19: Summary Data

Applicable Product(s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal
Default Setting	Enabled—Always On
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Operations Guide</i>

Revision History

Table 20: Revision History

Revision Details	Release
First introduced	2020.02.2.19
Note: This feature has not been fully qualified for this release.	

Feature Description

The NSO needs to know the state of the NF (Network Function) that has been deployed and notifications for the following states:

- DEPLOYED
- ALIVE
- UNDEPLOYED
- ERROR
- RECOVERING
- RECOVERY_FAILED



Note The SMI supports VM status notification for UPF only.

The following parameters are introduced in the notification streams for the VM status notifications:

Parameter	Description
vm-state-notification stream	<p>The vm-state-notification stream has the following details:</p> <ul style="list-style-type: none"> • cluster_name: the cluster holds KVM nodes • node_name: KVM node name • vm_name: UPF name • state: VM state • message: Carries useful information such as error message or mgmt-ip when the state of the VM is ALIVE.
alert-notification stream	<p>The alert-notification stream has the following labels:</p> <ul style="list-style-type: none"> • node_name: the KVM node name • vm_name: UPF name • state: the VM state • message: Carries useful info such as error message or mgmt-ip when state is ALIVE.

Depending on the life cycle stage of the VM, notifications are generated from either the Cluster Manager or CEE Ops-Center and sent to the NSO. When VM gets deleted or redeployed, the UNDEPLOYED notification is sent to the Cluster Manager notification stream. All other notifications are generated by the Alert Manager and then sent to the CEE Ops-Center notification stream.

Cluster Manager Notification

The Cluster Manager sends vm-state-notification for the UNDEPLOYED VM state.

CEE Ops-Center Notification

The CEE Ops-Center alert-notification sends the following alerts for different VM states:

- vm-deployed: minor - DEPLOYED
- vm-alive: minor – ALIVE (alert lasts for a short time and disappears automatically)
- vm-error: major - ERROR
- vm-recovering: warning - RECOVERING
- vm-recovery_failed: critical - RECOVERY_FAILED

All required fields are included in alert labels for notification from alert-notification. All VM alerts are viewable on the Grafana dashboard.

VM Action Notifications

Delete Action: When delete VM action is triggered, CM sends notifications that the VM is deleted. The VM states are UNDEPLOYED and ERROR for vm delete action.

```
clusters abc-cluster-15 nodes kvm-1 vms upf1 actions delete
```

Redeploy Action: When VM is in RECOVERY_FAILED state, NSO sends a request to redeploy the VM. A redeploy action does both delete action and sync action.

```
clusters abc-cluster-15 nodes kvm-1 vms upf1 actions redeploy
```

Redeploy Action Notification: The redeploy action sends a notification to the CM. The redeploy vm action has the following 4 states: UNDEPLOYED, ERROR, REDEPLOYED, REDEPLOY_ERROR.

```
show notification stream vm-state

notification
eventTime 2021-02-23T21:27:28.692+00:00
vm-state-notification
  cluster_name cndp-testbed
  node_name kvm-1
  vm_name upf2
  state UNDEPLOYED
  message
!
!
notification
eventTime 2021-02-23T21:29:18.699+00:00
vm-state-notification
  cluster_name cndp-testbed
  node_name kvm-1
  vm_name upf2
  state REDEPLOYED
  message
!
!
```

Configuring the Alert Notification in CEE

The user must configure alert notifications when they deploy the UPF VMs. Log in to the CEE cli to add the following configuration:

```
config
bulk-stats prune-interval-days 3
prometheus kvm-metrics defaults private-key "-----BEGIN OPENSSSH PRIVATE
KEY-----LGXtil23N4YV=\n-----END OPENSSSH PRIVATE KEY-----\n"
prometheus kvm-metrics defaults user cloud-user
prometheus kvm-metrics monitor-server 10.194.62.41
  hostname abc-bm-15-master
exit
```



Note The user must replace the IP, hostname, private key and user details.

Sample Notification from the Alert Notification Stream

```
notification
eventTime 2021-01-08T03:28:54.501+00:00
smi-alert-notification
  starts-at 2021-01-08T03:28:24.493874101Z
```



```

ends-at 0001-01-01T00:00:00Z
alert-status firing
smi-alert-notification alert-label
  name alertname
  value vm-recovery-failed
!
smi-alert-notification alert-label
  name cluster
  value test-cee-kvm_cee-voice
!
smi-alert-notification alert-label
  name hostname
  value test-bm-15-master
!
smi-alert-notification alert-label
  name instance
  value metrics-proxy-test-bm-15-master:9100
!
smi-alert-notification alert-label
  name job
  value metrics-proxy
!
smi-alert-notification alert-label
  name message
  value 10.1.1.3
!
smi-alert-notification alert-label
  name monitor
  value prometheus
!
smi-alert-notification alert-label
  name node_name
  value master
!
smi-alert-notification alert-label
  name replica
  value test-cee-kvm_cee-voice
!
smi-alert-notification alert-label
  name severity
  value critical
!
smi-alert-notification alert-label
  name state
  value RECOVERY_FAILED
!
smi-alert-notification alert-label
  name vm_name
  value upf2
!
smi-alert-notification alert-annotation
  name summary
  value upf2 failed to recover.
!
smi-alert-notification alert-annotation
  name type
  value Equipment Alarm
!
!
!

```

For information, refer to the [UCC SMI Operations Guide > SMI Cluster Manager Operations](#) chapter.

tac-debug-pkg CLI Enhancements

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	CEE Debugging
Applicable Platforms	Bare Metal, OpenStack, VMware
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Common Execution Environment - Configuration and Administration Guide</i>

Revision History

Revision Details	Release
First introduced.	2020.02.2.19

Feature Description

The CLI syntax for the tac-debug-pkg has been enhanced to improve ease of use.

Previously, the command syntax required a user to specify a time period by entering from and to criteria. The **tac-debug-pkg create** and **tac-debug-pkg delete** commands now allows users to specify the duration relative to the current time using the last keyword:

```
tac-debug-pkg create last <time_to_now>
```

```
tac-debug-pkg delete last <time_to_now>
```

<time_to_now> specifies the time to now in terms of the number of:

- Days – Expressed as “D”, “d”, or “day”; for example “5D”
- Hours – Expressed as “H”, “h”, or “hour”; for example “3h”
- Minutes – Expressed as “M”, “m”, “min”, or “minute”; for example “18minute”
- Seconds - Expressed as “S”, “s”, “sec”, or “second”; for example “3600sec”

Additionally, omitting the to keyword from the from parameter instructs the system to collect the TAC package from the specified time until *now*:

```
tac-debug-pkg create from <time_to_now>
```

The FROM keyword no longer requires the use of the TO keyword if you are creating the TAC package from a specific time until now.

Usage examples:

User Intention	Command
collect tac-debug-package for last 50 seconds	<code>tac-debug-pkg create last 50s</code>
collect tac-debug-package for last 10 minutes	<code>tac-debug-pkg create last 10min</code>
collect tac-debug-package for last 3 hours	<code>tac-debug-pkg create last 3H</code>
collect tac-debug-package for last 7 days	<code>tac-debug-pkg create last 7day</code>
delete all collected tac-debug-package for the past 2 days	<code>tac-debug-pkg delete last 2</code>
collect tac-debug-package from 2019-08-09_01:00:00 to now	<code>tac-debug-pkg create from 2019-08-09_01:00:00</code>

For more information on CEE debugging, refer to the [UCC CEE Configuration and Administration Guide > Common Execution Environment](#) chapter.

User Role APIs

Feature Summary and Revision History

Summary Data

Table 21: Summary Data

Applicable Product(s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal, OpenStack, VMware
Default Setting	Disabled—Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 22: Revision History

Revision Details	Release
First introduced	2020.02.2.19

Feature Description

SMI is equipped to provide the following functionalities using the REST APIs:

- List all the SMI users
- List all the SMI groups
- Add a new User
- Delete a User
- Assign User to a Group
- Unassign User from a Group

VIP Config Enhancements

Feature Summary and Revision History

Summary Data

Table 23: Summary Data

Applicable Product(s) or Functional Area	Cluster Deployer configuration validation
Applicable Platforms	Bare Metal, OpenStack, VMware
Default Setting	Enabled—Always On
Related Changes in this Release	Not Applicable
Related Documentation	<i>UCC SMI Common Execution Environment - Configuration and Administration Guide</i>

Revision History

Table 24: Revision History

Revision Details	Release
First introduced	2020.02.2.19

Feature Description

Multiple virtual IP (VIP) groups can be configured for use by the applications being deployed in the K8s cluster.

SMI's cluster deployer logic has been enhanced to check if any IPv4 or IPv6 VIP address has been assigned to more than one VIP group.

If the same VIP address has been assigned to multiple VIP groups, the deployment configuration validation will fail.

The following is a sample erroneous VIP groups configuration and a sample of the resulting error message logged through the validation:

Example Erroneous keepalived Configuration	Example Error Message
<pre>show running-config clusters tb1-smi-blr-c3 virtual-ips clusters tb1-smi-blr-c3 virtual-ips rep2 vrrp-interface ens224 vrrp-router-id 188 ipv4-addresses 192.168.139.85 mask 24 broadcast 192.168.139.255 device ens224 exit ipv4-addresses 192.168.139.95 mask 24 broadcast 192.168.139.255 device ens256 exit hosts master2 priority 99 exit hosts master3 priority 100 exit exit virtual-ips rep3 vrrp-interface ens224 vrrp-router-id 189 ipv4-addresses 192.168.139.85 mask 24 broadcast 192.168.139.255 device ens224 exit</pre>	<pre>Manual validation: clusters tb1-smi-blr-c3 actions validate-config run 2021-04-27 15:21:45.967 ERROR __main__: Duplicate not allowed: ipv4-addresses 192.168.139.85 is assigned across multiple virtual-ips groups 2021-04-27 15:21:45.968 ERROR __main__: virtual-ips groups with same ip-addresses are rep3 and rep2 2021-04-27 15:21:45.968 ERROR __main__: Checks failed in the cluster tb1-smi-blr-c3 are: 2021-04-27 15:21:45.968 ERROR __main__: Check: ntp failed. 2021-04-27 15:21:45.968 ERROR __main__: Check: k8s-node-checks failed. 2021-04-27 15:21:45.968 ERROR __main__: Check: vip-checks failed. Auto-Validation actions sync run: clusters tb1-smi-blr-c3 actions sync run This will run sync. Are you sure? [no,yes] yes message Validation errors occurred: Error: An error occurred validating SSH private key for cluster: tb1-smi-blr-c3 Error: An error occurred validating node proxy for cluster: tb1-smi-blr-c3 Error: An error occurred validating node oam label config for cluster: tb1-smi-blr-c3</pre>

The `keepalived_config` container monitors the `configmap vip-config` for any changes at regular intervals and if a change is detected the `keepalived` configuration file is reloaded.

With this enhancement, either all or none of the VIP addresses configured in a VIP group must be present on a node. If only some of the addresses exist on the node, that `keepalived` process will be stopped and a new process is automatically started and apply the latest configuration. This ensures that the `keepalived` processes assign those IP addresses appropriately.

The following is an example of the resulting error message logged through the validation:

```
kubectll logs keepalived-zqlzp -n smi-vips -c keepalived-config --tail 50 --follow
container
INFO:root:group name :rep2
INFO:root:Ip address: 192.168.139.85 on interface ens224 found on this device: True
INFO:root:Ip address: 192.168.139.95 on interface ens256 found on this device: False
INFO:root:Error Occurred: All VIPs in /config/keepalived.yaml must be either present or
```

absent in this device

INFO:root:VIP Split brain Scenario: Restarting the keepalived process.

For more information on deploying clusters, refer to the [UCC SMI Deployment Guide > SMI Cluster Manager – Deployment](#) chapter.

VPP CPU Worker Count

Feature Summary and Revision History

Summary Data

Applicable Product (s) or Functional Area	KVM-based application deployment support
Applicable Platforms	Bare Metal
Feature Default Setting	Disabled – Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	UCC SMI Deployment Guide

Revision History

Revision Details	Release
First introduced.	2020.02.2.19

Feature Description

The SR-IOV or PCI_PT provides ability to configure number of NIC Rx & Tx queues per VF. A maximum of 16 Rx and Tx queues can be configured per VF in SR-IOV.

For more information, refer to the [UCC SMI Operations Guide > SMI Cluster Manager Operations > Operating the SMI Cluster on Bare Metal](#) section.