



Pods and Services Reference

- [Feature Summary and Revision History, on page 1](#)
- [Feature Description, on page 2](#)
- [Associating Pods to the Nodes, on page 9](#)
- [Viewing the Pod Details and Status, on page 9](#)

Feature Summary and Revision History

Summary Data

Table 1: Summary Data

Applicable Products or Functional Area	SMF
Applicable Platform(s)	SMI
Feature Default Setting	Enabled – Always-on
Related Changes in this Release	Not Applicable
Related Documentation	Not Applicable

Revision History

Table 2: Revision History

Revision Details	Release
First introduced.	Pre-2020.02.0

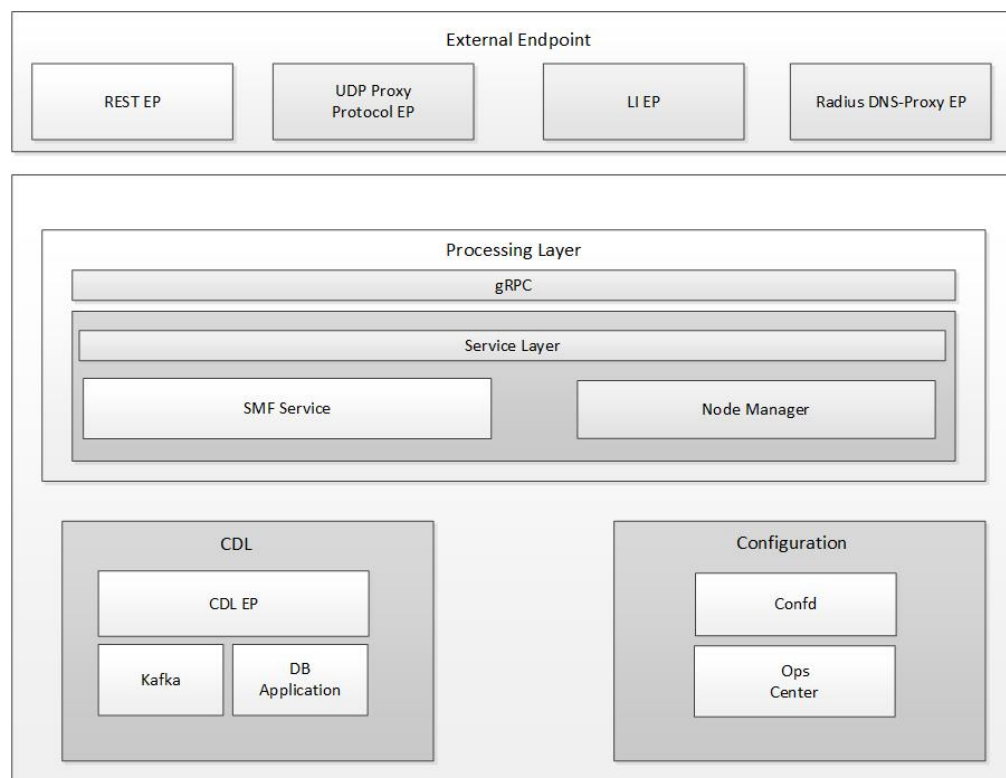
Feature Description

The SMF is built on the Kubernetes cluster strategy, which implies that it has adopted the native concepts of containerization, high availability, scalability, modularity, and ease of deployment. To achieve the benefits offered by Kubernetes, SMF uses the construct that includes the components, such as pods and services.

Depending on your deployment environment, the SMF deploys the pods on the virtual machines that you have configured. Pods operate through the services that are responsible for the intrapod communications. If the machine hosting the pods fails or experiences network disruption, the pods are terminated or deleted. However, this situation is transient and Kubernetes spins new pods to replace the invalid pods.

The following workflow provides a high-level visibility into the host machines, and the associated pods and services. It also represents how the pods communicate with each other. The representation may differ based on your deployment infrastructure.

Figure 1: Communication Workflow of Pods



Kubernetes deployment includes the `kubectl` command-line tool to manage the Kubernetes resources in the cluster. You can manage the pods, nodes, and services.

For information on the Kubernetes concepts, see the [Kubernetes documentation](#).

For more information on the Kubernetes components in SMF, see the following:

- Pods
- Services

Pods

A pod is a process that runs on your Kubernetes cluster. Pod encapsulates a granular unit known as a container. A pod contains one or multiple containers.

Kubernetes deploys one or multiple pods on a single node which can be a physical or virtual machine. Each pod has a discrete identity with an internal IP address and port space. However, the containers within a pod can share the storage and network resources.

The following table lists the SMF pod names and the hosts on which they are deployed depending on the labels that you assign. For information on how to assign the labels, see [Associating Pods to the Nodes](#).

Table 3: SMF Pods

Pod Name	Description	Virtual Machine Name
api-smf-ops-center	Functions as the <i>confD</i> API pod for the SMF Ops Center.	OAM
base-entitlement-smf	Supports Smart Licensing feature.	OAM
bgpspeaker	Dynamic routing for L3 route management and BFD monitoring	Protocol
cache-pod	Operates as the pod to cache any sort of system information that will be used by other pods as applicable.	Protocol
cdl-ep-session	Provides an interface to the CDL.	Session
cdl-index-session	Preserves the mapping of keys to the session pods.	Session
cdl-slot-session	Operates as the CDL session pod to store the session data.	Session
dns-proxy	Operates as DNS endpoint of SMF	Protocol
documentation	Contains the documentation.	OAM
etcd-smf-etcd-cluster	Hosts the etcd for the SMF application to store information, such as pod instances, leader information, NF-UUID, endpoints, and so on.	OAM
georeplication	Responsible for cache, etcd replication across sites, and site role management	Protocol
grafana-dashboard-cdl	Contains the default dashboard of CDL metrics in Grafana.	OAM
grafana-dashboard-smf	Contains the default dashboard of SMF service metrics in Grafana.	OAM
gtpc-ep	Operates as GTPC endpoint of SMF.	Protocol
kafka	Hosts the Kafka details for the CDL replication.	Protocol
li-ep	Operates as Lawful Intercept endpoint of SMF.	Protocol
oam-pod	Operates as the pod to facilitate Ops Center actions like show commands, configuration commands, monitor protocol monitor subscriber, and so on.	OAM
ops-center-smf-ops-center	Acts as the SMF Ops Center.	OAM

Pod Name	Description	Virtual Machine Name
smart-agent-smf-ops-center	Operates as the utility pod for the SMF Ops Center.	OAM
nodemgr	Performs node level interactions, such as N4 link establishment, management (heart-beat), and so on. Also, generates unique identifiers, such as UE IP address, SEID, CHF-ID, Resource URI, and so on.	Service
protocol	Operates as encoder and decoder of application protocols (PCFP, GTP, RADIUS, and so on) whose underlying transport protocol is UDP.	Protocol
radius-ep	Operates as RADIUS endpoint of SMF	Protocol
rest-ep	Operates as REST endpoint of SMF for HTTP2 communication.	Protocol
service	Contains main business logic of SMF.	Service
udp-proxy	Operates as proxy for all UDP messages. Owns UDP client and server functionalities.	Protocol
swift-smf-ops-center	Operates as the utility pod for the SMF Ops Center.	OAM
zookeeper	Assists Kafka for topology management.	OAM

For details on UDP proxy, see the [UDP Proxy Pod, on page 5](#) section.

These SMF pods communicate with the Common Execution Environment (CEE) pods. For the complete list of CEE pods, see the *UCC CEE Configuration and Administration Guide*.

Replicas

Each pod runs on a single instance of an application. To provide more resources by running more instances, you can use multiple Pods, one for each instance. This concept in Kubernetes is referred to as replication. Replicated Pods or replicas are usually created and managed as a group by a workload resource and its controller.

With multiple replicas, Kubernetes can distribute the load between them. During node failures, replicas can be used.



Note Replicas are based on the hardware and deployed call model.

UDP Proxy Pod

Feature Description

The SMF has UDP interfaces toward the UPF (N4) and SGW (s5 or s8 for EPS interworking). With the help of the protocol layer pods (smf-protocol and gtp-ep), the messages are encoded and decoded and exchanged on these UDP interfaces.

For achieving the functionalities mentioned on the 3GPP specifications:

- It is mandatory for the protocol layer pods to receive the original source and destination IP address and port number. But the original IP and UDP header is not preserved when the incoming packets arrive at the UDP service in the Kubernetes (K8s) cluster.
- Similarly, for the outgoing messages, the source IP set to the external IP address of the UDP service (published to the peer node) is mandatory. But the source IP is selected as per the egress interface when different instances of protocol layer pods send outgoing messages from different nodes of the K8s cluster.

The protocol layer POD spawns on the node, which has the physical interface configured with the external IP address to achieve the conditions mentioned earlier. However, spawning the protocol layer pods has the following consequences:

- It is not possible to achieve the node level HA (High Availability) because the protocol pods are spawned on the same node of the K8s cluster. Any failure to that node may result in loss of service.
- The protocol pods (smf-protocol, gtp-ep, and radius-ep) must include their own UDP client and server functionalities. In addition, each protocol layer pod may require labeling of the K8s nodes with the affinity rules. This restricts the scaling requirements of the protocol layer pods.

The SMF addresses these issues with the introduction of a new K8s POD called "udp-proxy." The primary objectives of this POD are:

- The "udp-proxy" POD acts as a proxy for all kinds of UDP messages. It also owns the UDP client and server functionalities.
- The protocol pods perform the individual protocol (PFCP, GTP, Radius) encoding and decoding and provide the UDP payload to the "udp-proxy" POD. The "udp-proxy" POD sends the UDP payload out after it receives the payload from the protocol pods.
- The "udp-proxy" POD opens the UDP sockets on a virtual IP (VIP) instead of a physical IP. This ensures that the "udp-proxy" POD does not have any strict affinity to a specific K8s node (VM). Thus, enabling node level HA for the UDP proxy.



Note One instance of the "udp-proxy" POD is spawned by default in all the worker nodes in the K8s cluster. The UDP proxy for SMF feature has functional relationship with the Virtual IP Address feature.

Architecture

The "udp-proxy" POD is placed in the worker nodes in the K8s cluster.

1. Each of the K8s worker node contains one instance of the "udp-proxy" POD. However, only one of the K8s worker node owns the virtual IP at any time. The worker node that owns the virtual IP remains in the active mode while all the other worker nodes remain in the standby mode.
2. The active "udp-proxy" POD binds to the virtual IP and the designated ports for listening to the UDP messages from the peer nodes (UPF and SGW).
3. The UDP payload received from the peer nodes are forwarded to one instance of the protocol, gtp-ep, or radius-ep pods. The payload is forwarded either on the same node or different node for further processing.
4. The response message from the protocol, gtp-ep, or radius-ep pods is forwarded back to the active instance of the "udp-proxy" POD. The "udp-proxy" POD sends the response message back to the corresponding peer nodes.
5. The SMF-initiated messages are encoded at the protocol, gtp-ep, or radius-ep pods. In addition, the UDP payload is sent to the "udp-proxy" POD. Eventually, the "udp-proxy" POD comprises of the complete IP payload and sends the message to the peer. When the response from the peer is received, the UDP payload is sent back to the same smf-protocol, gtp-ep, or radius-ep POD from which the message originated.

Protocol Pod Selection for Peer-Initiated Messages

When the "udp-proxy" pod receives the peer node (for instance UPF) initiated messages, it is load balanced across the protocol instances to select any instance of the protocol pod. An entry of this instance number is stored along with the source IP and source port number of the peer node. This ensures that the messages from the same source IP and source port are sent to the same instance that was selected earlier.

High Availability for the UDP Proxy

The UDP proxy's HA model is based on the keepalived virtual IP concepts. A VIP is designated to the N4 interface during deployment. Also, a keepalived instance manages the VIP and ensures that the IP address of the VIP is created as the secondary address of an interface in one of the worker nodes of the K8s cluster.

The "udp-proxy" instance on this worker node binds to the VIP and assumes the role of the active "udp-proxy" POD. All "udp-proxy" instances in other worker nodes remain in the standby mode.

Services

The SMF configuration consists of several microservices that run on a set of discrete pods. Microservices are deployed during the SMF deployment. SMF uses these services to enable communication between the pods. When interacting with another pod, the service identifies the pod's IP address to initiate the transaction and acts as an endpoint for the pod.

The following table describes the SMF services and the pod on which they run.

Table 4: SMF Services and Pods

Service Name	Pod Name	Description
base-entitlement-smf	base-entitlement-smf	Supports Smart Licensing feature.
bgpspeaker-pod	bgpspeaker	Dynamic routing for L3 route management and BFD monitoring
datastore-ep-session	cdl-ep-session	Responsible for the CDL session.

Service Name	Pod Name	Description
datastore-notification-ep	smf-rest-ep	Responsible for sending the notifications from the CDL to the <i>smf-service</i> through <i>smf-rest-ep</i> .
datastore-tls-ep-session	cdl-ep-session	Responsible for the secure CDL connection.
documentation	documentation	Responsible for the SMF documents.
etcd	etcd-smf-etcd-cluster-0, etcd-smf-etcd-cluster-1, etcd-smf-etcd-cluster-2	Responsible for pod discovery within the namespace.
etcd-smf-etcd-cluster-0	etcd-smf-etcd-cluster-0	Responsible for synchronization of data among the <i>etcd</i> cluster.
etcd-smf-etcd-cluster-1	etcd-smf-etcd-cluster-1	Responsible for synchronization of data among the <i>etcd</i> cluster.
etcd-smf-etcd-cluster-2	etcd-smf-etcd-cluster-2	Responsible for synchronization of data among the <i>etcd</i> cluster.
grafana-dashboard-app-infra	grafana-dashboard-app-infra	Responsible for the default dashboard of app-infra metrics in Grafana.
grafana-dashboard-cdl	grafana-dashboard-cdl	Responsible for the default dashboard of CDL metrics in Grafana.
grafana-dashboard-smf	grafana-dashboard-smf	Responsible for the default dashboard of SMF-service metrics in Grafana.
gtpc-ep	gtpc-ep	Responsible for inter-pod communication with GTP-C pod.
helm-api-smf-ops-center	api-smf-ops-center	Manages the Ops Center API.
kafka	kafka	Processes the Kafka messages.
li-ep	li-ep	Responsible for lawful-intercept interactions.
local-ldap-proxy-smf-ops-center	ops-center-smf-ops-center	Responsible for leveraging Ops Center credentials by other applications like Grafana.
oam-pod	oam-pod	Responsible to facilitate Exec commands on the Ops Center.
ops-center-smf-ops-center	ops-center-smf-ops-center	Manages the SMF Ops Center.
ops-center-smf-ops-center-expose-cli	ops-center-smf-ops-center	To access SMF Ops Center with external IP address.

Service Name	Pod Name	Description
smart-agent-smf-ops-center	smart-agent-smf-ops-center	Responsible for the SMF Ops Center API.
smf-sbi-service	smf-rest-ep	Responsible for routing incoming HTTP2 messages to REST-EP pods.
smf-n10-service	smf-rest-ep	Responsible for routing incoming N10 messages to REST-EP pods.
smf-n11-service	smf-rest-ep	Responsible for routing incoming N11 messages to REST-EP pods.
smf-n40-service	smf-rest-ep	Responsible for routing incoming N40 messages to REST-EP pods.
smf-n7-service	smf-rest-ep	Responsible for routing incoming N7 messages to REST-EP pods.
smf-nrf-service	smf-rest-ep	Responsible for routing incoming NRF messages to REST-EP pod.
smf-nodemgr	smf-nodemgr	Responsible for inter-pod communication with <i>smf-nodemgr</i> pod.
smf-protocol	smf-protocol	Responsible for inter-pod communication with smf-protocol pod
smf-radius-dns	smf-radius-dns	Responsible for inter-pod communication with smf-radius-dns pod
smf-rest-ep	smf-rest-ep	Responsible for inter-pod communication with smf-rest-ep pod
smf-service	smf-service	Responsible for inter-pod communication with smf-service pod
swift-smf-ops-center	swift	Operates as the utility pod for the SMF Ops Center
zookeeper	zookeeper	Assists Kafka for topology management
zookeeper-service	zookeeper	Assists Kafka for topology management

Open Ports and Services

The SMF uses different ports for communication purposes. The following table describes the default open ports and the associated services.

Table 5: Open Ports and Services

Port	Service	Usage
2024	SSH	SMF Ops Center uses this port to provide the ConfD CLI access.
8080	HTTP	SMF endpoint pods use this port for routing incoming messages on interfaces, such as N10, N11, N40, N7, and so on.

In addition to the preceding ports, SMF uses the ports that are destined for SMI for routing information between hosts. For information on SMI ports, see the *Ultra Cloud Core Subscriber Microservices Infrastructure Operations Guide*.

Associating Pods to the Nodes

This section describes how to associate a pod to the node based on their labels.

After you have configured a cluster, you can associate pods to the nodes through labels. This association enables the pods to get deployed on the appropriate node based on the key-value pair.

Labels are required for the pods to identify the nodes where they must get deployed and to run the services. For example, when you configure the protocol-layer label with the required key-value pair, the pods are deployed on the nodes that match the key-value pair.

To associate pods to the nodes through the labels, use the following sample configuration:

```
config
  k8 label vm_group key label_key value label_value
end
```

NOTES:

- **k8 label vm_group key label_key value label_value**: Configures the K8 node affinity label parameters.
 - **vm_group**: Specify the VM group. It must be one of the following:
 - cdl-layer
 - oam-layer
 - protocol-layer
 - service-layer
 - **key label_key**: Specify the label key. *label_key* must be a string.
 - **value label_value**: Specify the label value. *label_value* must be a string.
- If you choose not to configure the labels, then SMF assumes the labels with the default key-value pair.

Viewing the Pod Details and Status

If the service requires additional pods, SMF creates and deploys the pods. You can view the list of pods in your deployment through the SMF Ops Center.

You can run the **kubectl** command from the master node to manage the Kubernetes resources.

The pod details are available in YAML format.

Use the following sample configuration to view the comprehensive pod details:

```
kubectl get pods -n smf pod_name -o yaml
```

The output of this command displays the following information:

- The IP address of the host where the pod is deployed.
- The service and application that is running on the pod.
- The ID and name of the container within the pod.
- The IP address of the pod.
- The current state and phase in which the pod is.
- The start time from when the pod is in the current state.

To view all the pods in the SMF namespace, use the following sample configuration:

```
kp get pods -n smf_namespace -o wide
```

States

Understanding the pod's state lets you determine the current health and prevent the potential risks. The following table describes the pod's states.

Table 6: Pod States

State	Description
Running	The pod is healthy and deployed on a node. It contains one or more containers.
Pending	The application is in the process of creating the container images for the pod.
Succeeded	Indicates that all the containers in the pod are successfully terminated. These pods cannot be restarted.
Failed	One ore more containers in the pod have failed the termination process. The failure occurred as the container either exited with non zero status or the system terminated the container.
Unknown	The state of the pod could not be determined. Typically, this could be observed because the node where the pod resides was not reachable.