



# Cisco Common Data Layer

- [Feature Summary and Revision History, on page 1](#)
- [Feature Description, on page 2](#)
- [How it Works, on page 2](#)
- [Call Flows, on page 3](#)
- [Configuring the CDL Through SMF Ops Center, on page 4](#)

## Feature Summary and Revision History

### Summary Data

Table 1: Summary Data

Applicable Product(s) or Functional Area	SMF
Applicable Platform(s)	SMI
Feature Default Setting	Disabled – Configuration Required
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

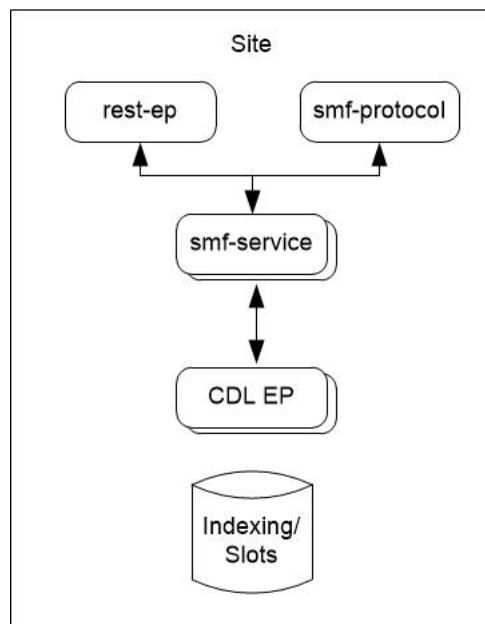
## Feature Description

The PCF extends support to the Geo Redundant (GR) version of the Cisco Common Data Layer (CDL). When the primary CDL endpoint fails, PCF attempts the same operation on the next highly rated secondary endpoint thus providing a non-disrupted N7 or Diameter message handling. If the next rated endpoint is unavailable, then PCF reattempts the operation on the subsequent endpoint that has the highest rating and so on.

## Architecture

The following figure depicts the failover that happens when the SMF service is unable to access the CDL datastore endpoint.

**Figure 1: CDL Datastore Architecture**



With relevance to this architecture, you can configure CDL through SMF Ops Center. When the SMF connects to the CDL, it uses the local endpoints.

## How it Works

When you configure the CDL in SMF through the SMF Ops Center, SMF gets enabled to support multiple CDL datastore endpoints. You can configure the endpoints by specifying the IP addresses, ports, and assigning ratings to each endpoint. By default, SMF considers the local endpoint as the primary endpoint, which has the highest rating. SMF performs CDL API operations on the primary endpoint. If this endpoint is unavailable, then SMF routes the operations to the next highest rated endpoint. SMF keeps failing over to the accessible secondary endpoint or until all the configured secondaries are exhausted. It does not reattempt a query on the next rated endpoint if the endpoint is reachable but responds with error or timeout.

If SMF is unable to access any of the endpoints in the cluster, then CDL operation fails with the "Datastore Unavailable" error.

# Call Flows

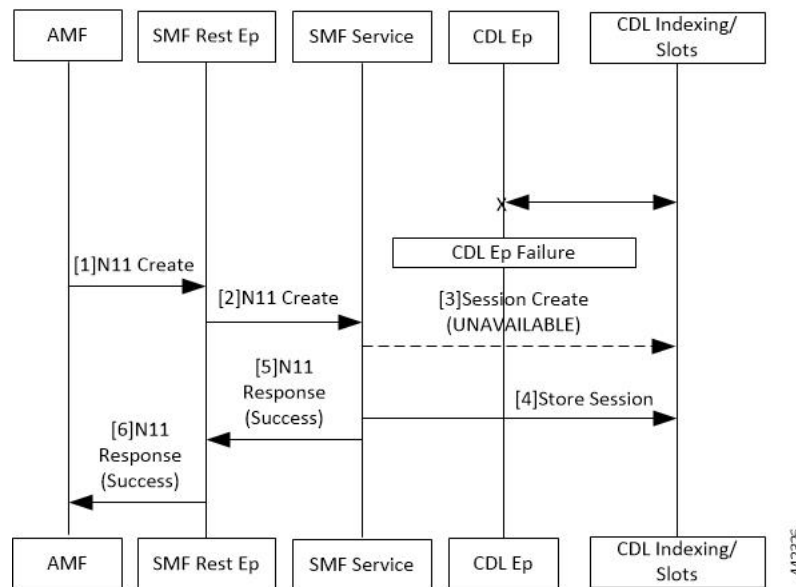
This section describes the call flow that is associated with this feature.

- [CDL Endpoint Failure Call Flow, on page 3](#)

## CDL Endpoint Failure Call Flow

This section describes the SMF local data store endpoint failure call flow.

**Figure 2: CDL Endpoint Failure Call Flow**



**Table 3: CDL Endpoint Failure Call Flow Description**

Step	Description
1	The AMF sends a Create Request to SMF REST endpoint over the N11 interface.
2	After receiving the request, the SMF REST endpoint forwards the Create Request to the SMF service.
3	The SMF service attempts to reach the CDL endpoint to send the session creation request. However, the CDL endpoint is unreachable.
4	The Create Request is evaluated in the stored session and the SMF service forwards the request to the CDL endpoint.
5	After the call request is successful, the SMF service notifies the Success Message to the SMF REST endpoint.
6	The SMF REST endpoint forwards the Success Message to the AMF.

## Limitations

The CDL configuration in SMF has the following limitations:

- The SMF service attempts to reroute the calls only when it encounters gRPC errors such as UNAVAILABLE. It does not acknowledge errors that the datastore endpoint returns and actual gRPC timeouts such as DEADLINE\_EXCEEDED gRPC status code.
- The SMF service does not resolve failures occurring with the datastore such as indexing and slot failures. The CDL layer must resolve these failures and if necessary, send an API call on the remote.

## Configuring the CDL Through SMF Ops Center

The configuration of the CDL using SMF Ops Center involves the following steps:

1. [Configuring the CDL Session Database and Defining the Base Configuration, on page 4](#)
2. [Configuring the Zookeeper in CDL, on page 5](#)

## Configuring the CDL Session Database and Defining the Base Configuration

This section describes how to configure the CDL session database and define the base configuration in SMF.

1. From the SMF Ops Center, run the following command to configure the CDL session database and base configuration.

```
configure
cdl system-id system_id
cdl node-type node_type
cdl zookeeper replica zookeeper_replica_id
exit
cdl logging default-log-level debug_level
cdl datastore session
cluster-id cluster_id
endpoint replica 1
endpoint replica num_replica
index map map_value
slot replica num_replica
slot map num_map/shards
slot write-factor write_factor
slot notification host host
slot notification port port
slot notification limit tps
index replica num_replica
index map num_map/shards
index write-factor write_factor
end
```

NOTES:

- **cdl system-id** *system\_id*: This is an optional command. Specifies the system or Kubernetes cluster identity. The default value is 1.
- **cdl node-type** *node\_type*: This is an optional command. Specifies the Kubernetes node label to configure the node affinity. The default value is “session.” Accepted length of the value is 0–64 alphabets.
- **cdl zookeeper replica** *zookeeper\_replica\_id*: Specifies the zookeeper replica server's ID.
- **endpoint replica** *num\_replica*: This is an optional command. Specifies the number of replicas to be created. The default value is 1. Must be an integer in the range of 1–16.
- **slot replica** *num\_replica*: This is an optional command. Specifies the number of replicas to be created. The default value is 1. *num\_replica* must be an integer in the range of 1–16.
- **slot map** *num\_map/shards*: This is an optional command. Specifies the number of partitions in a slot. The default value is 1. *num\_map/shards* must be an integer in the range of 1–1024.
- **slot write-factor** *write\_factor*: This is an optional command. Specifies the number of copies to be written before successful response. The default value is 1. *write\_factor* must be an integer in the range of 0–16. Make sure that the value is lower than or equal to the number of replicas.
- **slot notification host** *host*: This is an optional command. Specifies the notification server hostname or IP address. The default value is `datastore-notification-ep`.
- **slot notification port** *port*: This is an optional command. Specifies the notification server Port number. The default value is 8890.
- **slot notification limit** *tps*: This is an optional command. Specifies the notification limit per second. The default value is 2000.
- **index replica** *num\_replica*: This is an optional command. Specifies the number of replicas to be created. The default value is 2. *num\_replica* must be an integer in the range of 1–16.
- **index map** *num\_map/shards*: This is an optional command. Specifies the number of partitions in a slot. The default value is 1. *num\_map/shards* must be an integer in the range of 1–1024. Avoid modifying this value after deploying the CDL.
- **index write-factor** *write\_factor*: This is an optional command. Specifies the number of copies to be written before successful response. The default value is 1. *write\_factor* must be an integer in the range of 0–16.

## Configuring the Zookeeper in CDL

This section describes how to configure the Zookeeper in CDL.

1. Open the Policy Ops Center console and navigate to the datastore CLI.
2. Run the following command to define the parameters.

```
configure
cdl zookeeper data-storage-size data_storage_size_in_gb
log-storage-size log_storage_size_in_gb
replica number_of_replicas
enable-JMX-metrics boolean_value
```

```

    enable-persistence boolean_value
end

```

#### NOTES:

All the following parameters are optional.

- **cdl zookeeper data-storage-size** *data\_storage\_size\_in\_gb*: Specifies the size of the Zookeeper data storage in gigabyte. The default value is 20 GB. Accepted value is an integer in the range of 1-64.
- **log-storage-size** *log\_storage\_size\_in\_gb*: Specifies the size of the Zookeeper data log's storage in gigabyte. The default value is 20 GB. Accepted value is an integer in the range of 1-64.
- **replica num\_replicas**: Specifies the number of replicas that must be created. The default value is 3. Accepted value is an integer in the range of one to 16.
- **enable-JMX-metrics** *boolean\_value*: Specifies the status of the JMX metrics. The default value is true.
- **enable-persistence** *boolean\_value*: Specifies the status of the persistent storage for Zookeeper data. The default value is *false*.

## Sample Configuration

This section shows a sample configuration of CDL in a HA environment.

```

cdl system-id system_id
cdl zookeeper replica num_zk_replica
cdl datastore session
  endpoint replica ep_replica
index map index_shard_count
  slot replica slot_replica
  slot map slot_shard_count
exit

```