



Cisco Common Data Layer

- [Feature Summary and Revision History, on page 1](#)
- [Feature Description, on page 2](#)
- [How it Works, on page 2](#)
- [Configuring CDL through PCF Ops Center, on page 6](#)
- [Configuring the CDL Endpoints, on page 10](#)

Feature Summary and Revision History

Summary Data

Table 1: Summary Data

Applicable Product(s) or Functional Area	5G-PCF
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.	2020.01.0

Feature Description

The PCF extends support to the Geographic Redundancy (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 nondisrupted 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.

For more information on the CDL concepts, see the *Ultra Cloud Core Common Data Layer Configuration Guide*.

Limitations

In the current release, this feature has the following limitations:

- The PCF attempts to reroute the calls only when it encounters gRPC errors such as UNAVAILABLE. It does not acknowledge errors that the datastore returns and actual gRPC timeouts such as DEADLINE_EXCEEDED gRPC status code.
- The PCF Engine 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.

How it Works

When you configure the CDL in PCF through the PCF Ops Center, PCF gets enabled to support multiple CDL datastore endpoints. You can configure the endpoints by specifying the IP addresses, port numbers, and assigning ratings to each endpoint. By default, PCF considers the local endpoint as the primary endpoint, which has the highest rating. PCF performs CDL API operations on the primary endpoint. If this endpoint is unavailable, then PCF routes the operations to the next highest rated endpoint. PCF 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.

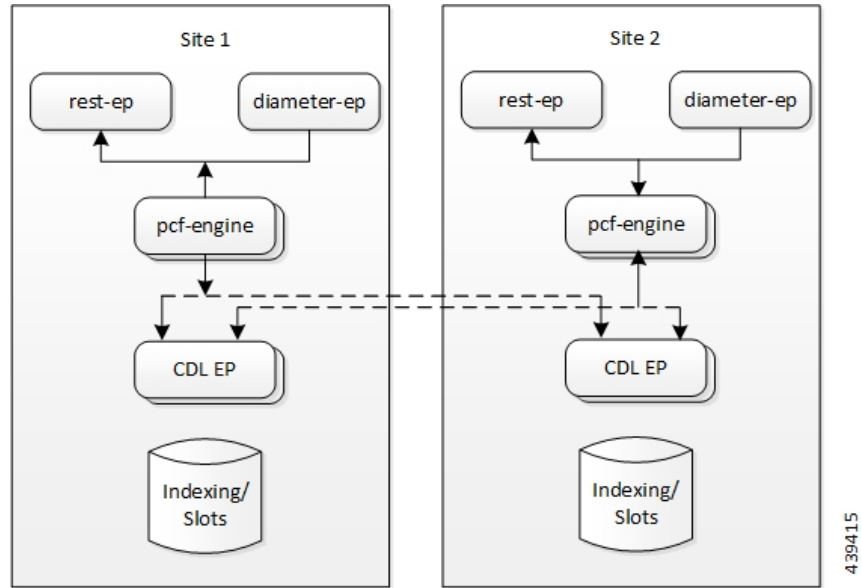
If PCF is unable to access any of the endpoints in the cluster, then CDL operation fails with the "Datastore Unavailable" error.

Architecture

You can configure CDL through PCF Ops Center. CDL in the GR mode replicates the session data across the configured sites. When PCF connects to the CDL, it always treats the local CDL endpoints as the primary endpoint and the remote endpoints as secondaries (with the appropriate rating). PCF uses the secondary endpoints when the connection to the primary endpoint fails.

The following illustration depicts the failover that happens when the PCF Engine is unable to access the primary CDL datastore endpoint.

Figure 1: CDL Datastore Architecture



Call Flows

This section describes the following call flow.

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

CDL Endpoint Failure Call Flow

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

Figure 2: CDL Endpoint Failure Call Flow

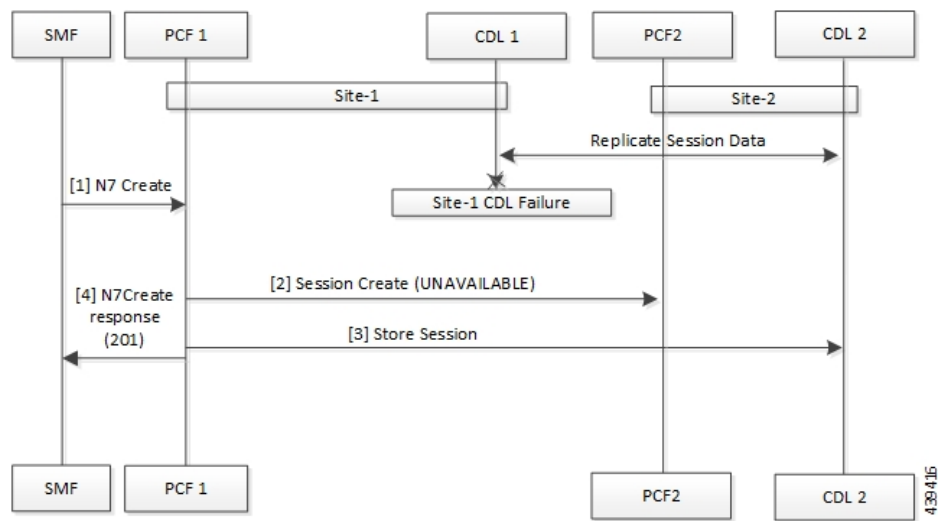


Table 3: CDL Endpoint Failure Call Flow Description

Step	Description
1	In the Site 1 environment, the SMF sends a N7 Create Request to the PCF 1 over the N7 interface.
2	The PCF 1 sends Session Create Request to the PCF 2.
3	The PCF 1 sends a Session Store Request to the CDL2.
4	The PCF 1 sends N7 Create Response to the SMF.

GR Call Flows

This section describes the possible CDL GR mode scenarios that could initiate a failover to another site.

This section describes the following call flows.

- [Indexing Shard Failure](#), on page 4
- [Slot Replica Set Failure](#), on page 5

Indexing Shard Failure

This section describes how the failover happens when two index replicas that belong to the same shard are down or unavailable.

The indexing shard failure is an example of two points-of-failure scenario where the two replicas reside on different virtual machines or hosts.

The PCF REST endpoint and PCF Engine redirect the traffic to the secondary CDL endpoint site (Site 2) based on the highest rating when the primary CDL site (Site 1) is unavailable.


Figure 3: Indexing Shard Failure Call Flow Description

Table 4: Indexing Shard Failure

Step	Description
1	In the Site 1 environment, index replica 1 and replica 2 for a configured shard has failed or unavailable. Since both the replicas for the shard are unavailable, the CDL endpoint in Site 1 is shut down and all the subsequent requests are directed to the CDL endpoint on Site 2. In the Site 1 environment, the SMF sends a Create Request to PCF REST endpoint over the N7 interface.
2	After receiving the request, the PCF REST endpoint forwards the Create Request to the PCF Engine.
3	The PCF Engine attempts to reach the CDL endpoint to send the Session Create Request. However, the CDL endpoint is unreachable. The PCF Engine sorts the CDL points across Site 1 and Site 2 to recognize the endpoint with the highest rating or priority.
4	The Create Request is evaluated in the stored session and the PCF Engine forwards the request to the CDL endpoint residing in Site 2.
5	After the call request is successful, the PCF Engine notifies the Success Message to the PCF REST endpoint.
6	The PCF REST endpoint forwards the Success Message to the SMF.

Slot Replica Set Failure

This section describes how the failover happens when two slot replicas that belong to the same replica set are down or unavailable.

The slot failure is an example of two points-of-failure scenario where the two slot replicas reside on different virtual machines or hosts.

Figure 4: Slot Failure Call Flow Description




Table 5: Slot Failure

Step	Description
1	In the Site 1 environment, slot replica 1 and replica 2 for a configured shard is down or unavailable. Since both the replicas for the shard are unavailable, the CDL endpoint in Site 1 is shut down and all the subsequent requests are directed to the CDL endpoint on Site 2. In the Site 1 environment, the SMF sends a N7 Create request to PCF REST endpoint over the N7 interface.
2	The PCF REST endpoint receives the request and forwards it to the PCF Engine.
3	The PCF Engine attempts to connect the CDL endpoint to send the Session Create request. If the CDL endpoint is unreachable, the PCF Engine sorts the CDL points across Site 1 and Site 2 to recognize the endpoint with the highest rating or priority.
4	The Create Request is evaluated in the stored session and the PCF Engine forwards the request to the CDL endpoint residing in Site 2.
5	After the call request is successful, the PCF Engine notifies the Success message to the PCF REST endpoint.
6	The PCF REST endpoint forwards the Success message to the SMF.

Configuring CDL through PCF Ops Center

This section describes how to configure the CDL endpoints.

Configuring the CDL using PCF Ops Center involves the following steps:

1. [Configuring the CDL Session Database and Defining the Base Configuration, on page 6](#)
2. [Configuring Kafka in CDL, on page 8](#)
3. [Configuring Zookeeper in CDL, on page 9](#)

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 PCF.

To configure the CDL session database and define the base configuration in CDL, use the following configuration in the Policy Ops Center console:

```

config
  cdl system-id system_id
  cdl node-type node_type
  cdl enable-geo-replication [ true | false ]
  cdl zookeeper replica zookeeper_replica_id
  cdl remote-site remote_system_id
    db-endpoint host host_name
    db-endpoint port port_number
    kafka-server remote_kafka_host1 remote_port1

```

```

kafka-server remote_kafka_host2 remote_port2
kafka-server remote_kafka_host3 remote_port3
exit
cdl logging default-log-level debug_level
cdl datastore session
cluster-id cluster_id
geo-remote-site remote_site_value
endpoint replica replica_number
endpoint external-ip ip_address
endpoint external-port port_number
index map map_value
slot replica replica_slot
slot map map/shards
slot write-factor write_factor
slot notification host host_name
slot notification port port_number
slot notification limit tps
index replica index_replica
index map 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.” *node_type* must be an alphabetic string of 0-64 characters.
- **cdl enable-geo-replication** [**true** | **false**] – This is an optional command. Specifies the geo replication status as enable or disable. The default value is false.
- **cdl zookeeper replica** *zookeeper_replica_id* – Specifies the Zookeeper replica server ID.
- **cdl remote-site** *remote_system_id* – Specifies the endpoint IP address for the remote site endpoint. Configure this command only when you have set the **cdl enable-geo-replication** to true.
- **db-endpoint host** *host_name* – Specifies the endpoint IP address for the remote site. Configure this command only when you have set the **cdl enable-geo-replication** to true.
- **db-endpoint port** *port_number* – Specifies the endpoint port number for the remote site endpoint. The default port number is 8882. Configure this command only when you have set the **cdl enable-geo-replication** to true.
- **kafka-server** *remote_kafka_host1 remote_port1* – Specifies the Kafka server’s external IP address and port number of the remote site that the **remote-system-id** identifies. You can configure multiple host address and port numbers per Kafka instance at the remote site. Configure this command only when you have set the **cdl enable-geo-replication** to true.
- **endpoint replica** *replica_number* – This is an optional command. Specifies the number of replicas to be created. The default value is 1. *replica_number* must be an integer in the range of 1 – 16.
- **endpoint external-ip** *ip_address* – This is an optional command. Lists the external IP address to expose the database endpoint. Configure this command only when you have set the **cdl enable-geo-replication** to true.

- **endpoint external-port** *port_number* – This is an optional command. Specifies the external port number to expose the database endpoint. Configure this command only when you have set the `cdl enable-geo-replication` to true. The default value is 8882.
- **slot replica** *replica_slot* – This is an optional command. Specifies the number of replicas to be created. The default value is 1. *replica_slot* must be an integer in the range of 1 – 16.
- **slot map** *map/shards* – This is an optional command. Specifies the number of partitions in a slot. The default value is 1. *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_name* – This is an optional command. Specifies the notification server hostname or IP address. The default value is `datastore-notification-ep`.
- **slot notification port** *port_number* – 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** *index_replica* – This is an optional command. Specifies the number of replicas to be created. The default value is 2. *index_replica* must be an integer in the range of 1 – 16.
- **index map** *map/shards* – This is an optional command. Specifies the number of partitions in a slot. The default value is 1. *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 Kafka in CDL

This section describes how to configure Kafka in CDL.

To configure the Kafka in CDL, use the following configuration:

1. Open the Policy Ops Center console and navigate to the datastore CLI.
2. To configure Kafka, use the following configuration:

```

config
  cdl kafka replica number_of_replicas
  enable-JMX-metrics [ true | false ]
  external-ip ip_address port_number
  enable-persistence [ true | false ]
  storage storage_size
  retention-time retention_period
  retention-size retention_size
end

```

NOTES:

All the following parameters are optional.

- **cdl kafka replica** *number_of_replicas* – Specifies the number of replicas to be created. The default value is 3. *number_of_replicas* must be an integer in the range of 1 – 16.
- **enable-JMX-metrics** [**true** | **false**] – Specifies the status of the JMX metrics. The default value is true.
- **external-ip** *ip_address port_number* – Lists the external IPs to expose to the Kafka service. Configure this command when you have set the **enable-geo-replication** parameter to true. You are required to define an external IP address and port number for each instance of the Kafka replica. For example, if the **cdl kafka replica** parameter is set to 3, then specify three external IP addresses and port numbers.
- **enable-persistence** [**true** | **false**] – Indicates whether to enable or disable persistent storage for Kafka data. The default value is false.
- **storage** *storage_size* – Specifies the Kafka data storage size in gigabyte. The default value is 20 GB. *storage_size* must be an integer in the range of 1-64.
- **retention-time** *retention_period* – Specifies the duration (in hours) for which the data must be retained. The default value is 3. *retention_period* must be an integer in the range of 1 – 168.
- **retention-size** *retention_size* – Specifies the data retention size in megabyte. The default value is 5120 MB.

Configuring Zookeeper in CDL

This section describes how to configure Zookeeper in CDL.

To configure Zookeeper in CDL, use the following configuration:

1. Open the Policy Ops Center console and navigate to the datastore CLI.
2. To configure the parameters, use the following configuration:

```

config
  cdl zookeeper data-storage-size data_storage
  log-storage-size log_storage
  replica number_of_replicas
  enable-JMX-metrics [ true | false ]
  enable-persistence [ true | false ]
  end

```

NOTES:

All the following parameters are optional.

- **cdl zookeeper data-storage-size** *data_storage* – Specifies the size of the Zookeeper data storage in gigabyte. The default value is 20 GB. *data_storage* must be an integer in the range of 1-64.
- **log-storage-size** *log_storage* – Specifies the size of the Zookeeper data log's storage in gigabyte. The default value is 20 GB. *log_storage* must be an integer in the range of 1-64.
- **replica** *number_replicas* – Specifies the number of replicas that must be created. The default value is 3. *number_replicas* must be an integer in the range of 1-16.

- **enable-JMX-metrics** [**true** | **false**] – Specifies the status of the JMX metrics. The default value is true.
- **enable-persistence** [**true** | **false**] – Specifies the status of the persistent storage for Zookeeper data. The default value is false.

Sample Configuration

The following is a sample configuration of CDL in the HA environment.

```
cdl system-id system_i
cdl enable-geo-replication true
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
cdl kafka replica kafka_replica
```

Configuring the CDL Endpoints

This section describes how to configure the CDL endpoints.

Configuring the CDL endpoints involves the following steps:

1. [Configuring the External Services, on page 10](#)
2. [Associating the Datastore with the CDL Endpoint Service, on page 11](#)

Configuring the External Services

This section describes how to configure the external services in PCF.

CDL gets deployed in the GR environment as part of the SMI deployment procedure. By default, the CDL endpoints are available in the Datastore CLI node of the PCF Ops Center. However, you are required to configure these endpoints.

For each CDL site and instance, configure external service with the IP address and port number that corresponds to the site and instance.

1. Open the Policy Ops Center console and navigate to the datastore CLI.
2. To configure the parameters, use the following configuration:

```
config
  external-services site_name
  ips ip_address
  ports port_number
end
```

NOTES:

- **external-services** *site_name* – Specifies the CDL site or instance name.

- **ips** *ip_address* – Specifies the IP address on which the CDL endpoint is exposed.
- **ports** *port_number* – Specifies the port number on which the CDL endpoint is exposed.

Associating the Datastore with the CDL Endpoint Service

This section describes how to configure the external service for each CDL endpoint service that you plan to use.

To configure the external service for each CDL endpoint service, use the following configuration:

1. Open the Policy Ops Center console and navigate to the datastore CLI.
2. To associate the datastore with CDL endpoint service, use the following configuration:

```
config
  datastore external-endpoints service_name
  port port_number
  rating rating_priority
end
```

NOTES:

- **datastore external-endpoints** *service_name* – Specifies the service name that belongs to the external services.
- **port** *port_number* – Specifies the port number where the external service resides.
- **rating** *rating_priority* – Specifies the rating or priority of the external service. PCF gives preference to the endpoints with the higher ratings.

