



Cisco Common Data Layer

- [Feature Summary and Revision History, on page 1](#)
- [Feature Description, on page 2](#)
- [How it Works, on page 4](#)
- [Configuring Cisco Common Data Layer, on page 11](#)
- [Configuring the CDL Engine, on page 15](#)
- [Configuring the CDL Endpoints, on page 15](#)
- [Starting the Remote Index Synchronization, on page 17](#)
- [Configuring the Stale Session Cleanup Using the Unique Key, on page 18](#)
- [Stale Sessions Cleanup Troubleshooting Information, on page 19](#)
- [OAM Support, on page 19](#)

Feature Summary and Revision History

Summary Data

Table 1: Summary Data

Applicable Product(s) or Functional Area	PCF
Applicable Platform(s)	SMI
Feature Default Setting	Disabled – Configuration required to enable
Related Documentation	Not Applicable

Revision History

Table 2: Revision History

Revision Details	Release
Enhancement introduced. PCF supports N5 Interface.	2022.02.0

Revision Details	Release
Enhancement introduced. PCF can handle issues of multiple CDL entry updates when multiple RxSTR received on PCF within a short gap. Added configuration support for: <ul style="list-style-type: none"> • Stale sessions cleanup • Remote index synchronization 	2021.04.0
First introduced.	2020.01.0

Feature Description

Geographic Redundancy

The PCF extends support to the Geographic Redundancy (GR) version of the Cisco Common Data Layer (CDL). When the highest rated CDL endpoint fails, PCF attempts the same operation on the next highly rated CDL endpoint thus providing a nondisrupted 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.

PCF can handle issues of multiple CDL entry updates when multiple RxSTR received on PCF within a short gap.



Note It is recommended to enable this feature after upgrading both local and remote sites to the latest PCF version.

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

Limitations

This GR support 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.

Stale Sessions Cleanup

In the CDL sessions, PCF adds the unique session key SupiDnnKey and the pre-existing unique keys that include FramedIpv6PrefixKey. With the CDL's index overwrite detection command in the PCF Ops Center,

the administrators can configure the ability to delete the old session using the same unique key while the new session is created.

The unique keys that should be used in the overwrite detection configuration are SupiDnnKey and FramedIpv6PrefixKey with the action as delete_record.



Note If two unique keys (one key mapped to the notify action and the other to the delete action) point to the same primary key, then only the notify action is considered for the primary key.

For more information on CDL components, see *Cisco Common Data Layer* documentation.

Limitations

This Stale Sessions Cleanup feature has the following limitations:

- Operations that depend on indexes for the stale sessions require either of the following:
 - The sessions must be present at the same subscriber that is reconnecting with the same DNN.
 - The associated framed IPv6 prefix is assigned to the same or the different subscriber session.

If the subscriber has reconnected with a different DNN or framed IPv6 prefix is not reassigned to a different session, the sessions are not identified as stale.

- The stale detection and cleanup procedures use the SupiDnnKey values. Indexes of the older session are not created based on the SupiDnnKey values.

If the stale session is created before an upgrade, and a new session is created for the same SUPI and DNN combination postupgrade, then the older session is not identified as stale.

- If the system has multiple stale sessions with the framed IPv6 prefix key, the corresponding index is associated only with the latest session.

When a new session is created with then same key then only one session gets associated.

Synchronizing the Index Records

Sometimes after the local site is reinstated, the index data on both the sites may not be consistent. To reconcile the records and eliminate the discrepancy in the sites, configure the sync operation that initiates index data synchronization on the site with its remote peers.

For information on how site isolation works in PCF, see [Site Isolation](#).



Note Configuring the sync operation may cause a negative performance impact. It is recommended to perform this operation in a production environment that experiences a high number of inconsistent index records.


A sync operation cannot be initiated for an index instance where the remote sync is in progress.

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



How it Works

This section describes how this feature works.

Geographic Redundancy

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 endpoints 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.

When Rx STR or N5 Delete messages are received on two different sites (site A and site B) for the same subscriber session, a conflict occurs while each PCF site tries to update and replicate the session data. In this situation:

- PCF receives notification from CDL with session record from both the sites.
- After receiving the notification from CDL based on the session creation state only one site must process the notification to resolve the conflict and save the session.

Processing of CDL Conflict Notification

The local and remote sites receive the same CDL conflict notification. The site where session is created will process the notification and the other site ignores the notification.



Note Based on GeoSiteName, PCF identifies whether the session is created at current site or not.

PCF decodes the records (local and remote) into session objects available in the CDL notification.

PCF considers the decoded local session object as a base and checks whether the Rx or N5 SessionIds are available in LastActionList of remote session object. If Rx or N5 SessionIds are available, PCF removes the following from base session object.

- Rx or N5 device session.
- Rx or N5 session tags (secondary keys).
- Rx or N5 session rules.

PCF then saves the modified local session.

Call Flows

This section describes the key call flows for this feature.

CDL Endpoint Failure Call Flow

This section describes the CDL 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 call flows scenarios that could start a failover to another site.

Indexing Shard Failure Call Flow

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

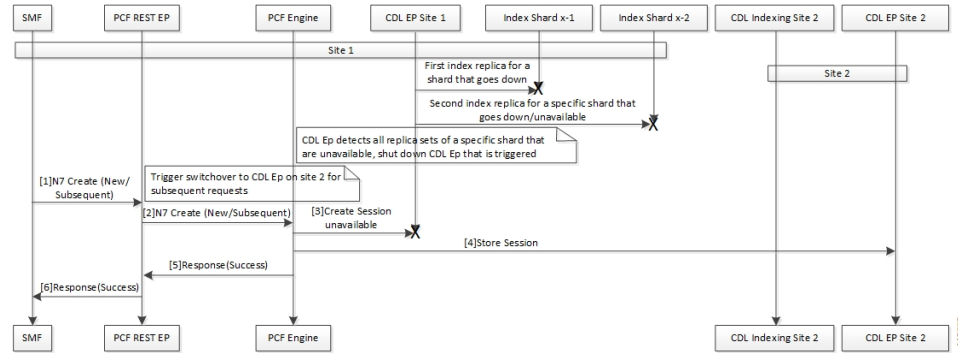


Table 4: Indexing Shard Failure Call Flow Description

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 Call Flow

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 Replica Set Failure Call Flow

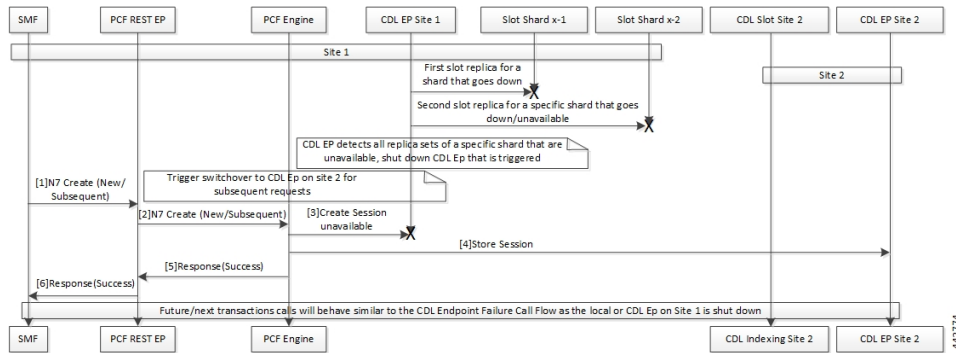


Table 5: Slot Replica Set Failure Call Flow Description

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.

Local and Remote Sites Receive Rx_STR Without Any Time Gap Call Flow

This section describes the local and remote sites receive Rx_STR without any time gap call flow.

Figure 5: Local and Remote Sites Receive Rx_STR Without Any Time Gap Call Flow




Table 6: Local and Remote Sites Receive Rx_STR Without Any Time Gap Call Flow Description

Step	Description
1	The SMF sends a N7 Create Request to the PCF 1 over the N7 interface.
2	The AF(n) sends a request Rx-AAR (R1) to the PCF 1.
3	The AF(n) sends a request Rx-AAR (R2) to the PCF 1.
4	The AF(n) sends the Rx Session-Termination-Request R1 to the PCF 1.
5	The AF(n) sends the Rx Session-Termination-Request R2 to the PCF 2.

Local and Remote Sites Receive N5 Delete Request Without Any Time Gap Call Flow

This section describes the local and remote sites receive N5 Delete Request without any time gap call flow.

Figure 6: Local and Remote Sites Receive N5 Delete Request Without Any Time Gap Call Flow




Table 7: Local and Remote Sites Receive Rx_STR Without Any Time Gap Call Flow Description

Step	Description
1	The SMF sends a N7 Create Request to the PCF 1 over the N7 interface.
2	The AF(n) sends a request N5 Create Request (R1) to the PCF 1.

Step	Description
3	The AF(n) sends a request N5 Create Request (R2) to the PCF 1.
4	The AF(n) sends the N5 Delete Request R1 to the PCF 1.
5	The AF(n) sends the N5 Delete Request R2 to the PCF 2.

Configuring Cisco Common Data Layer

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
2. Configuring Kafka in CDL
3. Configuring Zookeeper in CDL

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
    node-type node_type
    enable-geo-replication [ true | false ]
    zookeeper replica zookeeper_replica_id
    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
slot notification include-conflict-data [ true | false ]
index replica index_replica
index map map/shards
index write-factor write_factor
end

```

NOTES:

- **system-id** *system_id*—(Optional) Specify the system or Kubernetes cluster identity. The default value is 1.
- **node-type** *node_type*—(Optional) Specify 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.
- **enable-geo-replication** [true | false] —(Optional) Specify the geo replication status as enable or disable. The default value is false.
- **zookeeper replica** *zookeeper_replica_id*—Specify the Zookeeper replica server ID.
- **remote-site** *remote_system_id*—Specify 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*—Specify 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*—Specify 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*—Specify 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*—(Optional) Specify 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*—(Optional) Specify 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*—(Optional) Specify 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*—(Optional) Specify 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*—(Optional) Specify 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*—(Optional) Specify 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*—(Optional) Specify the notification server hostname or IP address. The default value is `datastore-notification-ep`.
- **slot notification port** *port_number*—(Optional) Specify the notification server port number. The default value is 8890.
- **slot notification limit** *tps*—(Optional) Specify the notification limit per second. The default value is 2000.
- **slot notification include-conflict-data** [**true** | **false**]—(Optional) Specify whether to receive the original data and the data from the request along with the DB conflict notification. This command is used to send conflict record data from CDL.
- **index replica** *index_replica*—(Optional) Specify 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*—(Optional) Specify 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*—(Optional) Specify 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*—Specify 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**]—Specify the status of the JMX metrics. The default value is true.
- **external-ip** *ip_address* *port_number*—Specify 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]—Specify whether to enable or disable persistent storage for Kafka data. The default value is false.
- **storage storage_size**—Specify 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**—Specify 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**—Specify 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*—Specify 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*—Specify 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*—Specify 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]—Specify the status of the JMX metrics. The default value is true.
- **enable-persistence** [true | false]—Specify 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 Engine

To configure this feature use the following configuration:

```
config
  cdl
    engine
      properties
        enable.conflict.merge [ true | false ]
        GeoSiteName geosite_name
        conflict.tps conflict_number
        conflict.resolve.attempts
      end
```

NOTES:

- **properties**—Specify the system properties.
- **enable.conflict.merge [true | false]**—Specify to enable the feature at application end.
- **GeoSiteName geosite_name**—Specify which site notification to be processed.
- **conflict.tps conflict_number**—Specify the rate limit of the conflict notification. The default value is considered as '5 tps'.
- **conflict.resolve.attempts**—Specify the number of attempts that application can try to merge the record. The default value is considered as '2 attempts'.



Note The **enable.conflict.merge [true | false]**, **conflict.tps conflict_number**, and **conflict.resolve.attempts** are to be configured manually.

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
2. Associating the Datastore with the CDL Endpoint Service

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*—Specify the CDL site or instance name.
- **ips** *ip_address*—Specify the IP address on which the CDL endpoint is exposed.
- **ports** *port_number*—Specify 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*—Specify the service name that belongs to the external services.
- **port** *port_number*—Specify the port number where the external service resides.

- **rating** *rating_priority*—Specify the rating or priority of the external service. PCF gives preference to the endpoints with the higher ratings.

Starting the Remote Index Synchronization

This section describes how to start the remote index synchronization.

Before configuring the remote index sync, ensure that the geo-remote-site parameter for CDL is configured.

To start the remote index synchronization, use the following configuration:

```
cdl
  actions
    remote-index-sync start [ map-id map_id | slice-name slice_name ]
  end
```

NOTES:

- **cdl**—Enters the CDL configuration mode.
- **remote-index-sync start**—Specify the remote index sync feature.
- **map-id** *map_id*—Specify the index mapID for which the remote index sync procedure should start. By default, remote index sync is initiated for all the index instances.
Using this parameter you can specify a maximum of 5 mapIDs.
- **slice-name** *slice_name*—Specify the slice name for which the remote index sync procedure should start. By default, remote index sync is initiated for all the sliceNames. There is no limit to the number of sliceNames that you can specify.

Sample Configuration

```
cdl actions remote-index-sync start map-id { 1 } map-id { 2
} slice-name { session-1 } slice-name { session-2 }
```

Viewing the Remote Index Synchronization Status

This section describes how to view the status of the index synchronization procedure that you have executed.

To view the status of the index sync procedure, use the following configuration:

```
cdl
  actions
    remote-index-sync status
  end
```

NOTES:

- **remote-index-sync status**—Displays the status of the index instances for which the syncing with the remote peers is in progress.

Sample Output

```
syncing-instances 'index-mapID-1-instanceID-1, index-mapID-1-
instanceID-2, index-mapID-2-instanceID-1, index-mapID-2-
instanceID-2'
```

Configuring the Stale Session Cleanup Using the Unique Key

The section describes how to configure stale session cleanup using the unique key.

To configure the stale session cleanup, use the following configuration:

```
config
  cdl
    datastore session datastore_name
      features
        index-overwrite-detection [ max-tps | queue-size |
unique-keys-prefix [ SupiDnnKey | FramedIpv6PrefixKey ]
          action [ notify-record | log-record | delete-record ]
        exit
      end
```

NOTES:

- **cdl**—Enter the CDL configuration mode.
- **datastore session *datastore_name***—Specify the CDL datastore session.
- **index-overwrite-detection [*max-tps* | *queue-size* | *unique-keys-prefix*]**—Configures the index keys overwrite detection capability. The parameter has the following subparameters:
 - **max-tps**—Specify the TPS per cdl-endpoint at which the stale record notification is sent. This parameter is applicable only when the action is "notify-record". The accepted value range is 1..2000. The default value is 200.
 - **queue-size**—Specify the queue size for each cdl-endpoint. The default value is 1000.
 - **unique-keys-prefix [*SupiDnnKey* | *FramedIpv6PrefixKey*]**—Specify the list of uniqueKey prefixes for the index overwrite detection and the action that must be performed on successful detection.
- **action [*log-record* | *delete-record*]**—Specify the action that must be taken on detecting a stale record.



Note If configuring the stale session cleanup feature for the first time on your system, Cisco recommends performing the configuration after both the GR sites are upgraded to the same software version.



Important The delete-record action on key SupiDnnKey command takes effect only when the required key SupiDnnKey is added in the CDL sessions.

Sample Configuration

The following is a sample configuration:

```
cdl datastore session
features index-overwrite-detection unique-keys-prefix SupiDnnKey
action delete-record
exit
features index-overwrite-detection unique-keys-prefix FramedIpv6PrefixKey
action delete-record
exit
end
```

Stale Sessions Cleanup Troubleshooting Information

To view the status of the clean up status of the stale sessions, review the warning logs in index pods.

You can review the logs to debug the stale sessions issues by setting the `index.overwrite.session` log to INFO level.

Example:

```
cdl logging logger index.overwrite.session
level info
exit
```

OAM Support

This section describes operations, administration, and maintenance support for this feature.

Statistics

Following are the list of counters that are generated for scenarios where the stale session cleanup process is initiated.

The following metrics track the counter information:

- `overwritten_index_records_deleted` - Captures the total number of records deleted due to overwritten or duplicate unique keys at index

Sample query: `overwritten_index_records_deleted`

The following labels are defined for this metric:

- `errorCode` - The error code in the DB response. Example: 0, 502.
- `sliceName` - The name of the logical sliceName. Example: session

- `overwritten_index_records_skipped` - Captures the total number of records detected as stale, but dropped when the queue is full while processing the records for notify or delete.

Sample query: `overwritten_index_records_skipped`

The following labels are defined for this metric:

- `action` - The action that was supposed to be performed for the stale record. Example: delete, notify.

- sliceName - The name of the logical sliceName. Example: session

The following statistics are supported to handle issues of multiple CDL entries updates when multiple Rx STR or N5 Delete received on PCF within a short gap feature:



Note The following values apply to all the statistics:

- Unit - Int64
- Type - Counter
- Nodes - Service

-
- record_conflict_merge_total - Captures the total count of conflict merge actions.

The following label is defined for this metric:

- action

The "action" label supports the following values:

- ok: Captures success processing of conflict notification.
- submit: Captures the number of messages submitted to the engine when conflict notification is received from CDL.
- retry: Captures the number of retry operations occurred during conflict merge.
- skip_<reason for skip>: Indicates that the PCF is expecting some data validation before the records are merged when CDL notification is received. If that data is missing, PCF logs these skip counters with reason to skip the data.

Reasons to skip the data are:

- throttle: Due to throttle check.
- feature_disabled: Feature is disabled.
- no_geositename: GeoSiteName is not configured.
- flag1_mismatch: Mismatch of CDL flag 1 at both sites.
- unsupported_record: Invalid data records of CDL notification.
- retry_unsupported_record: Invalid data records available during retry.
- unsupported_lastaction: Invalid lastAction objects are available in notification records.
- retry_unsupported_lastaction: During retry lastAction objects are invalid.
- no_sessionid - Session id is not available in remoteLastAction object.
- retry_no_sessionid: Session id is not available in remoteLastActoin object.
- no_remotesessionid: Remote session id is not available to remove in local record object.

- `retry_no_remotesessionid`: Remote session id not available to remove in local record object during retry.
- `attemptsdone`: Total number of attempts completed.
- `error_<cause of error>`: Indicates while merging the records some error/exception occurred in that case pcf logs this error related counter.

Following are the types of cause of errors:

- `deletesession`: Remote rx session delete operation from the local record failed.
- `retry_deletesession`: During retry remote rx session delete operation from the local record failed.
- `removeflags`: after deleting remote rx session from local record while removing corresponding flags from local record failed.
- `addactionlist`: Failed to consolidated all action list objects from local and remote records.
- `nopk`: Primary key is not available in the record.
- `retry_nopk`: Primary key is not available in the record during retry.

For information on statistics, see *Ultra Cloud Core Common Data Layer Configuration and Administration Guide*.

