



CDL Config Mode Command Reference

- [cdl](#), on page 2
- [cdl datastore](#), on page 3
- [cdl datastore endpoint](#), on page 4
- [cdl datastore endpoint settings](#), on page 6
- [cdl datastore features index-overwrite-detection unique-keys-prefix](#), on page 7
- [cdl datastore features instance-aware-notification](#), on page 7
- [cdl datastore features instance-aware-notification system-id](#), on page 8
- [cdl datastore features overload-protection](#), on page 8
- [cdl datastore features slot-stale-session-detection](#), on page 10
- [cdl datastore index](#), on page 11
- [cdl datastore metrics](#), on page 12
- [cdl datastore overload-protection](#), on page 13
- [cdl datastore slot](#), on page 13
- [cdl datastore slot metrics](#), on page 15
- [cdl datastore slot notification](#), on page 15
- [cdl kafka](#), on page 17
- [cdl kafka actions remote-index-sync start](#), on page 19
- [cdl kafka actions remote-index-sync start map-id](#), on page 19
- [cdl kafka actions remote-index-sync start slice-name](#), on page 19
- [cdl kafka actions remote-index-sync status](#), on page 20
- [cdl kafka clear sessions](#), on page 20
- [cdl kafka clear sessions filter](#), on page 21
- [cdl kafka describe](#), on page 21
- [cdl kafka describe consumer-groups](#), on page 21
- [cdl kafka describe topics](#), on page 22
- [cdl kafka external-ip](#), on page 22
- [cdl kafka label-config](#), on page 23
- [cdl kafka list](#), on page 23
- [cdl kafka list consumer-groups](#), on page 23
- [cdl kafka list topics](#), on page 24
- [cdl kafka log-compaction](#), on page 24
- [cdl kafka mirror-maker](#), on page 24
- [cdl kafka reset-offsets](#), on page 25

- [cdl kafka reset-offsets consumer-group](#), on page 25
- [cdl kafka reset-offsets-dry-run](#), on page 26
- [cdl kafka reset-offsets-dry-run consumer-group](#), on page 26
- [cdl kafka show indexes](#), on page 27
- [cdl kafka show sessions count detailed](#), on page 28
- [cdl kafka show sessions count detailed filter](#), on page 28
- [cdl kafka show sessions count summary](#), on page 29
- [cdl kafka show sessions count summary filter](#), on page 29
- [cdl kafka show sessions detailed](#), on page 30
- [cdl kafka show sessions detailed filter](#), on page 31
- [cdl kafka show sessions summary](#), on page 32
- [cdl kafka show sessions summary filter](#), on page 33
- [cdl kafka show status](#), on page 34
- [cdl kafka ssl-settings](#), on page 34
- [cdl label-config](#), on page 35
- [cdl label-config endpoint](#), on page 35
- [cdl label-config index map](#), on page 35
- [cdl label-config slot map](#), on page 36
- [cdl logging](#), on page 36
- [cdl logging logger](#), on page 38
- [cdl remote-site](#), on page 38
- [cdl remote-site db-endpoint](#), on page 39
- [cdl remote-site kafka-server](#), on page 39
- [cdl ssl-config](#), on page 40
- [cdl ssl-config certs](#), on page 40
- [cdl zookeeper](#), on page 41

cdl

Configures the next generation Cisco Data Layer (CDL).

Command Modes

Exec > Global Configuration

Syntax Description

```
cdl { datastore datastore_name { cluster-id cluster_id | { endpoint {
external-ip ipv4_address | external-port port_number } | index { map
index_partitions | write-factor write_factor } | slot { map slot_partitions |
notification { host host_name | limit notification_tps | port notification_server_port
} | write-factor write_factor } } { go-debug go_debug | patch-url patch_url |
replica ha_instances } | geo-remote-site list } } | deployment-model { large
| small } | enable-geo-replication { false | true } | kafka {
enable-JMX-metrics { false | true } | enable-persistence { false | true
} | external-ip { ipv4_address port_number } | replica ha_instances | retention-size
data_retention_size | retention-time data_retention_time | storage data_storage_size
} | logging { { default-log-level | logger cdl_module_name level
default_log_level } | node-type k8s_node_label_value | remote-site system_id {
db-endpoint { host db_endpoint_hostname | port db_endpoint_port_number } |
kafka-server { host port } } | system-id system_id | zookeeper {
```

```
data-storage-size data_storage_size | { enable-JMX-metrics | enable-persistence
} { false | true } | log-storage-size log_storage_size | replica ha_instances
} } }
```

deployment-model *deployment_model*

Specify the k8s deployment resource requirement. The default value is large. NOTE: In a production environment, you must always use the large deployment model.

Must be one of the following:

- large
- small

enable-geo-replication { false | true }

Specify to enable or disable geo-replication.

Must be one of the following:

- false
- true

Default Value: false.

node-type *k8s_node_label_value*

Specify the K8s node label value for affinity.

Must be a string of 0-64 characters.

Default Value: .

system-id *system_id*

Specify the deployment system ID.

Must be an integer.

Default Value: 1.

Usage Guidelines

Use this command to configure CDL parameters for deployment. If CDL container is not configured, the CDL chart will not be deployed.

cdl datastore

Configures the list of datastores to be deployed.

Command Modes

Exec > Global Configuration

Syntax Description

```
cdl datastore datastore_name { advanced-config-profile profile_name | cluster-id
cluster_id | copies-per-node copies_per_node | label-config label_config | { {
endpoint { external-ip ipv4_address | external-port port_number } | index {
```

```
map index_partitions | write-factor write_factor } | slot { map slot_partitions |
notification { host host_name | limit notification_tps | port notification_server_port
} | write-factor write_factor } } { go-debug go_debug | patch-url patch_url |
replica ha_instances } } | geo-remote-site list }
```

advanced-config-profile *profile_name*

Specify the advanced configuration parameters profile name.

cluster-id *cluster_id*

Specify the K8s cluster ID.

Must be an integer in the range of 1-10.

Default Value: 1.

geo-remote-site *list*

Specify the list of geo replication sites.

label-config *label_config*

Specify the label configuration name.

slice-names *slice_names_list*

Specify the list of slice names.

Must be a string of 1-16 characters.

Default Value: session.

datastore_name

Specify the name of the datastore.

Must be a string of 1-16 characters.

Usage Guidelines

Use this command to configure the list of datastores to be deployed.

You can configure a maximum of 10 elements with this command.

cdl datastore endpoint

Configures endpoint pod parameters.

Command Modes

Exec > Global Configuration

Syntax Description

```
cdl datastore datastore_name endpoint { go-debug go_debug_variables | patch-url
patch_url | replica ha_instances }
```

copies-per-node *copies_per_node*

Specify copies of cdl-endpoint pods per node.

Must be an integer in the range of 1-4.

Default Value: 1.

cpu-limit *max_cpu*

Specify the maximum CPU allowed in milli CPU.

Must be an integer in the range of 200-32000.

Default Value: 4000.

cpu-request *min_cpu*

Specify the minimum CPU required in milli CPU.

Must be an integer in the range of 100-16000.

Default Value: 2000.

external-ip *ipv4_address*

Specify the list of IPv4 addresses to expose the service externally.

Must be a string in the ipv4-address pattern. For information on the ipv4-address pattern, see the Input Pattern Types section.

-Or-

Must be a string in the ipv6-address pattern. For information on the ipv6-address pattern, see the Input Pattern Types section.

external-ip *ipv6_address*

Specify the list of IPv6 addresses to expose the service externally.

Must be a string in the ipv6-address pattern. For information on the ipv6-address pattern, see the Input Pattern Types section.

external-port *port_number*

Specify the port number to be exposed for external traffic.

Must be an integer in the range of 0-65535.

Default Value: 8882.

go-debug *go_debug_variables*

Specify the comma-separated list of GO Debug variables.

Must be a string.

go-max-procs *max_cpu*

Specify the maximum number CPU cores used simultaneously.

Must be an integer in the range of 1-32.

Default Value: 4.

memory-limit *max_memory*

Specify the maximum memory allowed in megabytes.

Must be an integer in the range of 512-131072.

Default Value: 131072.

memory-request *min_memory*

Specify the minimum memory required in megabytes.

Must be an integer in the range of 256-65536.

Default Value: 256.

patch-url *patch_url*

Specify the URL to download the datastore application patch.

Must be a string.

replica *ha_instances*

Specify the number of HA instances to be created.

Must be an integer in the range of 1-16.

Default Value: 2.

Usage Guidelines Use this command to configure endpoint pod parameters.

cdl datastore endpoint settings

This command is deprecated. Configures the CD endpoint parameters.

Command Modes Exec > Global Configuration

Syntax Description **settings** **index-timeout-ms** *index_timeout* **slot-timeout-ms** *slot_timeout*
remote-ep-connection-count *remote_ep_connection_count*

index-timeout-ms *index_timeout*

This keyword is deprecated. Specify the timeout period for index operations in milliseconds.

Must be an integer in the range of 25-5000.

Default Value: 200.

remote-ep-connection-count *remote_ep_connection_count*

This keyword is deprecated. Specify the number of connections to remote site CDL endpoint for GRPC replication.

Must be an integer in the range of 1-20.

Default Value: 4.

slot-timeout-ms *slot_timeout*

This keyword is deprecated. Specify the timeout period for slot operations in milliseconds.

Must be an integer in the range of 25-5000.

Default Value: 375.

Usage Guidelines

This command is deprecated. Use this command to configure CDL endpoint parameters.

cdl datastore features index-overwrite-detection unique-keys-prefix

Configures CDL Overload Protection configuration.

Command Modes

Exec > Global Configuration

Syntax Description

```
index-overwrite-detection unique-keys-prefix { action index_overwrite_action
| key-prefix key_prefix | max-tps max_tps | queue-size queue_size }
```

action *index_overwrite_action*

Specify the action to be performed on index overwrite detection.

Must be one of the following:

- **delete-record**
- **log-record**
- **notify-record**

key-prefix *key_prefix*

Specify the prefix pattern for unique keys for overwrite detection.

Must be a string.

Usage Guidelines

Index keys overwrite detection configuration. Use this command to configure the list of uniqueKey prefixes for index overwrite detection, and the action to be performed on detection.

cdl datastore features instance-aware-notification

Configures Geo Replication instance awareness configuration.

Command Modes

Exec > Global Configuration

Syntax Description

```
features instance-aware-notification enable { false | true }
```

enable { false | true }

Specify whether to enable the instance state check for slices.

Must be one of the following:

- **false**
- **true**

Default Value: false.

Usage Guidelines Use this command to configure Geo Replication instance awareness configuration.

cdl datastore features instance-aware-notification system-id

Configures the mapping of systemID to sliceNames.

Command Modes Exec > Global Configuration

Syntax Description **system-id is** *system_id* **slice-names** *slice_names*

id *system_id*

Specify the system/site ID.

Must be an integer.

slice-names *slice_names*

Specify the list of slice names associated with the systemID.

Usage Guidelines Use this command to configure the mapping of systemID to sliceNames.

cdl datastore features overload-protection

Configures CDL Overload Protection configuration.

Command Modes Exec > Global Configuration

Syntax Description **cdl datastore overload-protection disable { false | true } |**
critical-alert-threshold *critical_alert_threshold* | **hard-limit-percentage**
hard_limit_percentage | **index-record-capacity-count** *index_record_capacity_count* |
major-alert-threshold *major_alert_threshold* | **slot-record-capacity-bytes**
slot_record_capacity | **slot-record-capacity-count** *slot_record_capacity_count*

enable { false | true }

Specify to disable or enable Overload Protection.

Must be one of the following:

- **false**
- **true**

Default Value: false.

index-max-record-count *index_max_record_count*

Specify the maximum number of records that can be stored in the Index shard.

Must be an integer.

Must be an integer from the following: 100 - 1000 or 100000 - 100000000.

Default Value: 60000000.

enable { false | true }

Specify to disable or enable Overload Protection.

Must be one of the following:

- **false**
- **true**

Default Value: false.

index-max-record-count *index_max_record_count*

Specify the maximum number of records that can be stored in the Index shard.

Must be an integer.

slot-max-record-count *slot_max_record_count*

Specify the maximum number of records that can be stored in the Slot shard.

Must be an integer.

Must be an integer from the following: 100 or 100000 - 10000000.

Default Value: 2500000.

critical-alert-threshold *critical_alert_threshold*

Specify the threshold in percentage at which CDL must trigger an alert `cdlOverloaded-critical-major-alert-threshold` must be less than `critical-alert-threshold`.

Must be an integer in the range of 40-100.

Default Value: 90.

enable { false | true }

Specify to disable or enable Overload Protection.

Must be one of the following:

- **false**

- **true**

Default Value: false.

hard-limit-percentage *hard_limit_percentage*

Specify the extra percentage of capacity in addition to the soft limit.

Must be an integer in the range of 0-10.

Default Value: 5.

index-max-record-count *index_max_record_count*

Specify the maximum number of records that can be stored in the Index shard.

Must be an integer.

major-alert-threshold *major_alert_threshold*

Specify the threshold in percentage at which CDL must trigger an alert `cdlOverloaded-major-major-alert-threshold` must be less than `critical-alert-threshold`.

Must be an integer in the range of 40-100.

Default Value: 80.

slot-max-record-count *slot_max_record_count*

Specify the maximum number of records that can be stored in the Slot shard.

Must be an integer.

slot-max-size *slot_max_size*

Specify the maximum size of the Slot shard in megabytes.

Must be an integer in the range of 1024-98304.

Default Value: 16384.

Usage Guidelines Use this command to configure CDL Overload Protection configuration.

cdl datastore features slot-stale-session-detection

Configures the stale slot records detection parameters.

Command Modes Exec > Global Configuration

Syntax Description `slot-stale-session-detection disable-auto-deletion { false | true }`

disable-auto-deletion { false | true }

Specify to disable or enable stale slot record auto deletion.

Must be one of the following:

- **false**
- **true**

Default Value: false.

notification-retry-count *notification_retry_count*

Specify the minimum number of timer expiry notification retries sent to an application without receiving an update from the application. If no update is received even after "notification-retry-count" number of times, CDL proceeds to check if slot record is stale.

Must be an integer in the range of 3-10.

Default Value: 3.

Usage Guidelines Use this command to configure the stale slot records detection parameters.

cdl datastore index

Configures index pod parameters.

Command Modes Exec > Global Configuration

Syntax Description `cdl datastore datastore_name index { map index_partitions | write-factor write_factor | go-debug go_debug | patch-url patch_url | replica ha_instances }`

cpu-limit *max_cpu*

Specify the maximum CPU allowed in milli CPU.

Must be an integer in the range of 200-32000.

Default Value: 4000.

cpu-request *min_cpu*

Specify the minimum CPU required in milli CPU.

Must be an integer in the range of 100-16000.

Default Value: 2000.

go-debug *go_debug_variables*

Specify the comma-separated list of GO Debug variables.

Must be a string.

go-max-procs *max_cpu*

Specify the maximum number CPU cores used simultaneously.

Must be an integer in the range of 1-32.

Default Value: 4.

map *index_partitions*

Specify the number of partitions to be created for index.

Must be an integer in the range of 1-1024.

Default Value: 1.

memory-limit *max_memory*

Specify the maximum memory allowed in megabytes.

Must be an integer in the range of 512-131072.

Default Value: 131072.

memory-request *min_memory*

Specify the minimum memory required in megabytes.

Must be an integer in the range of 256-65536.

Default Value: 256.

patch-url *patch_url*

Specify the URL to download the datastore application patch.

Must be a string.

replica *replicas*

Specify the number of HA instances to be created.

Must be an integer in the range of 1-2.

Default Value: 2.

write-factor *write_factor*

This keyword has been deprecated. Specify the number of copies to be written into the index before a successful response.

Must be an integer in the range of 0-2.

Default Value: 2.

Usage Guidelines Use this command to configure the index pod parameters.

cdl datastore metrics

Configures the verbosity level of metrics.

Command Modes Exec > Global Configuration

Syntax Description `cdl datastore datastore_name { metrics { verbosity { debug | production } } }`

Usage Guidelines Use this command to configure the verbosity level of metrics.

cdl datastore overload-protection

This command is deprecated. Configures CDL Overload Protection configuration.

Command Modes Exec > Global Configuration

Syntax Description `cdl datastore overload-protection disable { false | true } | critical-alert-threshold critical_alert_threshold | hard-limit-percentage hard_limit_percentage | index-record-capacity-count index_record_capacity_count | major-alert-threshold major_alert_threshold | slot-record-capacity-bytes slot_record_capacity | slot-record-capacity-count slot_record_capacity_count`

disable { false | true }

This keyword is deprecated. Specify to disable or enable Overload Protection.

Must be one of the following:

- **false**
- **true**

Default Value: false.

Usage Guidelines This command is deprecated. Use this command to configure CDL Overload Protection configuration.

cdl datastore slot

Configures slot pod parameters.

Command Modes Exec > Global Configuration

Syntax Description `cdl datastore datastore_name slot { map slot_partitions | notification { host host_name | limit notification_tps | port notification_server_port } | write-factor write_factor }`

cpu-limit *max_cpu*

Specify the maximum CPU allowed in milli CPU.

Must be an integer in the range of 200-32000.

Default Value: 4000.

cpu-request *min_cpu*

Specify the minimum CPU required in milli CPU.

Must be an integer in the range of 100-16000.

Default Value: 2000.

go-debug *go_debug_variables*

Specify the comma-separated list of GO Debug variables.

Must be a string.

go-max-procs *max_cpu*

Specify the maximum number CPU cores used simultaneously.

Must be an integer in the range of 1-32.

Default Value: 4.

map *slot_partitions*

Specify the number of partitions to be created for slot.

Must be an integer in the range of 1-1024.

Default Value: 1.

memory-limit *max_memory*

Specify the maximum memory allowed in megabytes.

Must be an integer in the range of 512-131072.

Default Value: 131072.

memory-request *min_memory*

Specify the minimum memory required in megabytes.

Must be an integer in the range of 256-65536.

Default Value: 256.

patch-url *patch_url*

Specify the URL to download the datastore application patch.

Must be a string.

replica *ha_instances*

Specify the number of HA instances to be created.

Must be an integer in the range of 1-4.

Default Value: 2.

write-factor *write_factor*

Specify the number of copies to be written into slot before a successful response.

Must be an integer in the range of 1-4.

Default Value: 2.

Usage Guidelines Use this command to configure slot pod parameters.

cdl datastore slot metrics

Configures the CDL slot metrics parameters.

Command Modes Exec > Global Configuration

Syntax Description `metrics report-idle-session-type { false | true }`

report-idle-session-type { false | true }

Specify whether to report idle session types.

Must be one of the following:

- **false**
- **true**

Default Value: false.

Usage Guidelines Use this command to configure the CDL slot metrics parameters.

cdl datastore slot notification

Configures notification configuration towards application.

Command Modes Exec > Global Configuration

Syntax Description `cdl datastore slot notification { connection-count connection_count | dynamic-provisioning { false | true } | host host_name | limit notification_tps | max-concurrent-bulk-notifications max_bulk_task_notifications | notification-retry-seconds notification_retry_interval | port notification_server_port | remote-system-id remote_system_id }`

connection-count *connection_count*

Specify the Notification connections to the application notification endpoint.

Must be an integer in the range of 1-20.

Default Value: 4.

dynamic-provisioning { false | true }

Specify whether notification endpoint is provisioned by application or to use static config. If true, dynamic provisioning is enabled.

Must be one of the following:

- **false**
- **true**

Default Value: false.

host *host_name*

Specify the notification server's host name.

Must be a string.

Default Value: datastore-notification-ep.

include-conflict-data { *false* | *true* }

Specify to enable the setting to receive the original data and the data from the request along with the DB conflict notification.

Must be one of the following:

- **false**
- **true**

Default Value: false.

limit *notification_tps*

Specify the notification server's transactions per second (TPS).

Must be an integer in the range of 1-2000.

Default Value: 200.

max-concurrent-bulk-notifications *max_notifications*

Specify the maximum number of bulk task notifications that CDL can process concurrently.

Must be an integer in the range of 1-20.

Default Value: 1.

notification-retry-seconds *notification_retry_interval*

Specify the number of seconds after which the CDL will retry notification to application, if the application does not update next eval time post timer expiry notification.

Must be an integer in the range of 10-100.

Default Value: 10.

port *notification_server_port*

Specify the notification server's port number.

Must be an integer in the range of 0-65535.

Default Value: 8890.

remote-system-id *remote_system_id*

Specify the list of remote System IDs. Notification will be sent from the CDL if a record has owner system ID which matches one of the remote IDs.

Usage Guidelines

Use this command to configure notification configuration towards application.

cdl kafka

Configures the Kafka for Geo Replication.

Command Modes

Exec > Global Configuration

Syntax Description

```
cdl kafka { enable-JMX-metrics { false | true } | enable-persistence {
false | true } | external-ip { ipv4_address port_number } | replica ha_instances
| retention-size retention_size | retention data_retention_period | storage
data_storage_size | log-level { debug | error | fatal | info | trace | warn
} }
```

enable-JMX-metrics { false | true }

Specify to enable or disable JMX metrics.

Must be one of the following:

- **false**
- **true**

Default Value: true.

enable-persistence { false | true }

Specify to enable or disable persistence for the Kafka data.

Must be one of the following:

- **false**
- **true**

Default Value: false.

log-level

Specify the Kafka logging level. The logging levels include debug, error, fatal, info, trace, and warn. The changes specified here will trigger rolling restart of Zookeeper pods.

Must be one of the following:

- **debug**
- **error**
- **fatal**

- **info**
- **trace**
- **warn**

partition *number_of_partitions*

Specify the number of partitions per topic.

Must be an integer in the range of 1-1024.

Default Value: 1.

replica *ha_instances*

Specify the number of HA instances to be created.

Must be an integer in the range of 1-16.

Default Value: 2.

Default Value: 2.

retention-size *retention_size*

This keyword is deprecated. Specify the data retention size in MB.

Must be an integer in the range of 1024-1048576.

Default Value: 5120.

retention-time *data_retention_period*

Specify the data retention period in hours.

Must be an integer in the range of 1-168.

Default Value: 12.

storage *data_storage_size*

Specify the Kafka data storage size in GB.

Must be an integer in the range of 1-64.

Default Value: 20.

to-earliest

Specify to enable parallel pod management policy.

Must be one of the following:

- **false**
- **true**

Default Value: false.

total-retention-size *total_retention_size*

Specify the data total retention size in MB.

Must be an integer in the range of 10240-1048576.

Default Value: 65536.

Usage Guidelines Use this command to configure Kafka for Geo replication.

cdl kafka actions remote-index-sync start

Initiates syncing index with its remote peers.

Command Modes Exec

Syntax Description `actions remote-index-sync start`

Usage Guidelines Reconciles indexes with the remote peer. Use this command to start syncing index with its remote peers.

cdl kafka actions remote-index-sync start map-id

Triggers remote sync for the list of index map IDs.

Command Modes Exec

Syntax Description `actions remote-index-sync map-id map index_name`

map *index_map*

Specify the trigger remote index sync for the index map.

Must be an integer.

Usage Guidelines Use this command to trigger remote sync for the list of index map IDs. By default, it will trigger all index maps.

You can configure a maximum of five elements with this command.

cdl kafka actions remote-index-sync start slice-name

Triggers remote sync for the list of slice names.

Command Modes Exec

Syntax Description `actions remote-index-sync slice-name slice slice_name`

slice *slice_name*

Specify to trigger remote index sync for the slice name.

Must be a string.

Usage Guidelines Use this command to trigger remote sync for the list of slice names. By default, it will trigger all slice names.

cdl kafka actions remote-index-sync status

Displays the index instances that are still syncing with the remote peers.

Command Modes Exec

Syntax Description `cdl show status`

Usage Guidelines Use this command to view the comma-separated list of index instances that are still syncing with the remote peers.

cdl kafka clear sessions

Clears all datastore sessions data.

Command Modes Exec

Syntax Description `cdl clear sessions [db-name db_name | filter { condition query_expression key query_string } | map-id map_id]`

db-name *db_name*

Specify the name of the DB to query.

Must be a string of 1-16 characters.

Default Value: session.

map-id *map_id*

Specify the map ID to query data from a map.

Must be an integer in the range of 0-1024.

Default Value: 0.

slice-name *slice_name*

Specify the slice name to query.

Must be a string of 1-16 characters.

Default Value: session.

Usage Guidelines Use this command to clear all datastore session data. Note that this may result in the permanent deletion of some or all records.

cdl kafka clear sessions filter

Configures the filter parameters used for queries.

Command Modes Exec

Syntax Description `filter { condition query_expression key query_string } | map-id map_id`

condition *query_expression*

Specify the query expression.

Must be one of the following:

- ends-with
- match
- starts-with

key *query_string*

Specify the query string.

Must be a string of 0-512 characters.

Usage Guidelines Use this command to configure the filter parameters for queries.

cdl kafka describe

Describes the Kafka consumer groups. By default, it describes all the Kafka consumer groups.

Command Modes Exec

Syntax Description `cdl kafka describe { consumer-groups name group_name | topics { include-internal-topics | name kafka-topic_name } }`

Usage Guidelines Use this command to describe the kafka consumer groups.

cdl kafka describe consumer-groups

Describes the Kafka consumer groups.

Command Modes Exec

Syntax Description `cdl kafka describe consumer-groups name group_name`

group name *group_name*

Specify the Kafka consumer group name.

Must be a string in the pattern [a-zA-Z0-9._-]+.

Usage Guidelines

Use this command to list all the Kafka consumer groups. The default behaviour is to describe all the consumer groups.

cdl kafka describe topics

Describes the Kafka topics. By default, it describes all the topics that are not internal.

Command Modes

Exec

Syntax Description

```
cdl kafka describe topics { include-internal-topics | name kafka_topic_name
}
```

include-internal-topics

Specify whether to include internal Kafka topics.

topic *kafka-topic_name*

Specify the Kafka topic name.

Must be a string in the pattern [a-zA-Z0-9._-]+.

Usage Guidelines

Use this command to describe the Kafka topics.

cdl kafka external-ip

Configures the IPs that are exposed to the external servers.

Command Modes

Exec > Global Configuration

Syntax Description

```
cdl kafka external-ip ipv4_address port port_number ssl-port ssl_port_number
```

ip *IPv4_address*

Specify the IPv4 address for the kafka server to listen for client traffic.

Must be a string in the ipv4-address pattern. For information on the ipv4-address pattern, see the Input Pattern Types section.

port *port_number*

Specify the port address to expose for external traffic.

Must be an integer in the range of 0-65535.

ssl-port *ssl-port_address*

Specify the ssl-port address to expose for external traffic.

Must be an integer in the range of 0-65535.

Usage Guidelines Use this command to configure the IPs that are exposed to the external servers.

cdl kafka label-config

Configures label configuration for Kafka pods.

Command Modes Exec > Global Configuration

Syntax Description `cdl kafka label-config key label_key value label_value`

key *label_key*

Specify the label key.

Must be a string.

value *label_value*

Specify the label value.

Must be a string.

Usage Guidelines Use this command to configure label configuration for Kafka pods.

cdl kafka list

Shows all the kafka topics. By default, the list does not contain the internal topics.

Command Modes Exec

Syntax Description `cdl kafka list { consumer-groups name group_name | topics { include-internal-topics | name kafka-topic_name } }`

Usage Guidelines Use this command to show all the Kafka topics.

cdl kafka list consumer-groups

Shows all the Kafka consumer groups.

Command Modes Exec

Syntax Description `cdl kafka list consumer-groups name group_name`

Usage Guidelines Use this command to show all the Kafka consumer groups.

cdl kafka list topics

Shows all the Kafka topics.

Command Modes Exec

Syntax Description `cdl kafka list topics include-internal-topics`

include-internal-topics

Specify whether to include internal Kafka topics.

Usage Guidelines Use this command to show all the Kafka topics.

cdl kafka log-compaction

Configures Kafka log compaction parameters.

Command Modes Exec > Global Configuration

Syntax Description `cdl kafka log-compaction disable { false | true } delete-retention-ms delete_retention_period`

delete-retention-ms delete_retention_period

Specify the period of time in milliseconds to retain deletes for log compacted topics.

Must be an integer in the range of 60000-172800000.

Default Value: 43200000.

disable { false | true }

Specify to disable or enable Kafka log compaction.

Must be one of the following:

- **false**
- **true**

Default Value: false.

Usage Guidelines Use this command to configure Kafka log compaction parameters.

cdl kafka mirror-maker

Configures the Kafka MirrorMaker.

Command Modes	Exec > Global Configuration
Syntax Description	<pre>cdl mirror-maker { enable-JMX-metrics { false true } log-level { debug error fatal info trace warn } }</pre> <p>enable-JMX-metrics { false true }</p> <p>Specify to enable or disable JMX metrics</p> <p>Must be one of the following:</p> <ul style="list-style-type: none"> • false • true <p>Default Value: true.</p> <p>log-level { debug error fatal info trace warn }</p> <p>Specify the Kafka MirrorMaker logging level. NOTE: The changes trigger restart of the MirrorMaker pod</p> <p>Must be one of the following:</p> <ul style="list-style-type: none"> • debug • error • fatal • info • trace • warn
Usage Guidelines	Use this command to configure the Kafka MirrorMaker.

cdl kafka reset-offsets

Resets the offsets of all the topics of the Kafka consumer groups. The offset of a consumer group can be set only if the group is inactive.

Command Modes	Exec
Syntax Description	<pre>cdl kafka reset-offsets { consumer-group name <i>group_name</i> to-earliest name <i>group_name</i> to-latest name <i>group_name</i> }</pre>
Usage Guidelines	Use this command to reset the offsets of all the topics of the Kafka consumer groups.

cdl kafka reset-offsets consumer-group

Resets Kafka consumer group offset of all the topics.

Command Modes	Exec
Syntax Description	<pre>cdl kafka reset-offsets consumer-group name <i>group_name</i></pre> <p>consumer-group name <i>group_name</i> Specify the Kafka consumer group. Must be a string in the pattern [a-zA-Z0-9._-]+.</p> <p>to-earliest Specify to reset offsets to earliest.</p> <p>to-latest Specify to reset offsets to latest.</p>
Usage Guidelines	Use this command to reset Kafka consumer group offset of all the topics.

cdl kafka reset-offsets-dry-run

Shows results of reset-offsets without executing changes on consumer groups.

Command Modes	Exec
Syntax Description	<pre>cdl kafka reset-offsets-dry-run { consumer-group name <i>group_name</i> to-earliest name <i>group_name</i> to-latest name <i>group_name</i> }</pre>
Usage Guidelines	Use this command to show results of reset-offsets without executing changes on consumer groups.

cdl kafka reset-offsets-dry-run consumer-group

Resets Kafka consumer group offset of all the topics.

Command Modes	Exec
Syntax Description	<pre>cdl kafka reset-offsets-dry-run consumer-group name <i>group_name</i></pre> <p>consumer-group name <i>group_name</i> Specify the Kafka consumer group. Must be a string in the pattern [a-zA-Z0-9._-]+.</p> <p>to-earliest Specify to reset offsets to earliest.</p>

to-earliest

Specify to reset offsets to latest.

Usage Guidelines

Use this command to reset Kafka consumer group offset of all the topics.

cdl kafka show indexes

Displays session index information.

Command Modes

Exec

Syntax Description

```
show indexes all-replicas { false | true } | db-name db_name | key key |
limit max_indexes | map-id map_id ]
```

all-replicas { false | true }

Specify whether to fetch index from all the replicas.

Must be one of the following:

- **false**
- **true**

Default Value: false.

db-name *db_name*

Specify the name of the DB to query.

Must be a string of 1-16 characters.

Default Value: session.

key *key*

Specify the key - query value (uk/pk).

Must be a string of 0-512 characters.

limit *max_indexes*

Specify the maximum number of indexes to display.

Must be an integer in the range of 1-1000.

map-id *map_id*

Specify the map ID to query data from a map where the index is stored.

Must be an integer in the range of 0-1024.

slice-name *slice_name*

Specify the slice name to query.

Must be a string of 1-16 characters.

Default Value: session.

Usage Guidelines

Use this command to view session index information.

cdl kafka show sessions count detailed

Displays detailed session count information.

Command Modes

Exec

Syntax Description

```
cdl show sessions count detailed [ db-name db_name | filter { condition
query_expression key query_string } | limit number_of_records | map-id map_id ]
```

db-name *db_name*

Specify the name of the DB to query.

Must be a string of 1-16 characters.

Default Value: session.

map-id *map_id*

Specify the map ID to query data from a map.

Must be an integer in the range of 0-1024.

Default Value: 0.

slice-name *slice_name*

Specify the slice name to query.

Must be a string of 1-16 characters.

Default Value: session.

Usage Guidelines

Use this command to view detailed session count information.

cdl kafka show sessions count detailed filter

Configures the filter parameters used for queries.

Command Modes

Exec

Syntax Description

```
filter { condition query_expression key query_string } | map-id map_id
```

condition *query_expression*

Specify the query expression.

Must be one of the following:

- **ends-with**
- **match**
- **starts-with**

key *query_string*

Specify the query string.

Must be a string of 0-512 characters.

Usage Guidelines

Use this command to configure the filter parameters for queries.

cdl kafka show sessions count summary

Displays summary session count information.

Command Modes

Exec

Syntax Description

```
cdl show sessions count summary [ db-name db_name | filter { condition
query_expression key query_string } | limit number_of_records | map-id map_id ]
```

db-name *db_name*

Specify the name of the DB to query.

Must be a string of 1-16 characters.

Default Value: session.

map-id *map_id*

Specify the map ID to query data from a map.

Must be an integer in the range of 0-1024.

Default Value: 0.

slice-name *slice_name*

Specify the slice name to query.

Must be a string of 1-16 characters.

Default Value: session.

Usage Guidelines

Use this command to view summary session count information.

cdl kafka show sessions count summary filter

Configures the filter parameters used for queries.

Command Modes	Exec
Syntax Description	<p>filter { condition <i>query_expression</i> key <i>query_string</i> } map-id <i>map_id</i></p> <p>condition <i>query_expression</i></p> <p>Specify the query expression.</p> <p>Must be one of the following:</p> <ul style="list-style-type: none"> • ends-with • match • starts-with <p>key <i>query_string</i></p> <p>Specify the query string.</p> <p>Must be a string of 0-512 characters.</p>
Usage Guidelines	Use this command to configure the filter parameters for queries.

cdl kafka show sessions detailed

Displays session details with data.

Syntax Description	<p>cdl show detailed [db-name <i>db_name</i> filter { condition <i>query_expression</i> key <i>query_string</i> } limit <i>number_of_records</i> map-id <i>map_id</i>]</p>
Command Modes	Exec
	<p>db-name <i>db_name</i></p> <p>Specify the name of the DB to query.</p> <p>Must be a string of 1-16 characters.</p> <p>Default Value: session.</p> <p>limit <i>number_of_records</i></p> <p>Specify the maximum number of records to display.</p> <p>Must be an integer in the range of 1-500.</p> <p>Default Value: 100.</p> <p>map-id <i>map_id</i></p> <p>Specify the map ID to query data from a map.</p> <p>Must be an integer in the range of 0-1024.</p> <p>Default Value: 0.</p>

max-data-size-kb *max_data_size*

Specify the maximum data size to display per record. If the data size is greater than this parameter, the output is truncated.

Must be an integer in the range of 1-32.

Default Value: 4.

next-eval-end-time *next_eval_end_time*

Sessions that have next-eval-time less than this time.

next-eval-start-time *next_eval_start_time*

Sessions that have next-eval-time greater than this time.

purge-on-eval { false | true }

Sessions that have purge-on-eval flag set to false or true.

Must be one of the following:

- false
- true

slice-name *slice_name*

Specify the slice name to query.

Must be a string of 1-16 characters.

Default Value: session.

Usage Guidelines Use this command to view session details with data.

cdl kafka show sessions detailed filter

Configures the filter parameters used for queries.

Command Modes Exec

Syntax Description **filter** { **condition** *query_expression* **key** *query_string* } | **map-id** *map_id* | **limit** *number_of_records* | **purge-on-eval** { **false** | **true** } | **next-eval-start-time** *next_eval_start_time* | **next-eval-end-time** *next_eval_end_time*

condition *query_expression*

Specify the query expression.

Must be one of the following:

- ends-with
- match

- **starts-with**

key *query_string*

Specify the query string.

Must be a string of 0-512 characters.

Usage Guidelines Use this command to configure the filter parameters used for queries.

cdl kafka show sessions summary

Displays session details without data.

Syntax Description `cdl show sessions summary [db-name db_name | filter { condition query_expression key query_string } | limit number_of_records | map-id map_id]`

Command Modes Exec

db-name *db_name*

Specify the name of the DB to query.

Must be a string of 1-16 characters.

Default Value: session.

limit *number_of_records*

Specify the maximum number of records to display.

Must be an integer in the range of 1-500.

Default Value: 100.

map-id *map_id*

Specify the map ID to query data from a map.

Must be an integer in the range of 0-1024.

Default Value: 0.

max-data-size-kb *max_data_size*

Specify the maximum data size to display per record. If the data size is greater than this parameter, the output is truncated.

Must be an integer in the range of 1-32.

Default Value: 4.

next-eval-end-time *next_eval_end_time*

Sessions that have next-eval-time less than this time.

next-eval-start-time *next_eval_start_time*

Sessions that have next-eval-time greater than this time.

purge-on-eval { false | true }

Sessions that have purge-on-eval flag set to false or true.

Must be one of the following:

- false
- true

slice-name *slice_name*

Specify the slice name to query.

Must be a string of 1-16 characters.

Default Value: session.

Usage Guidelines

Use this command to view session details without data.

cdl kafka show sessions summary filter

Configures the filter parameters used for queries.

Command Modes

Exec

Syntax Description

```
filter { condition query_expression key query_string } | map-id map_id | limit
number_of_records | purge-on-eval { false | true } | next-eval-start-time
next_eval_start_time | next-eval-end-time next_eval_end_time
```

condition *query_expression*

Specify the query expression.

Must be one of the following:

- ends-with
- match
- starts-with

key *query_string*

Specify the query string.

Must be a string of 0-512 characters.

Usage Guidelines

Use this command to configure the filter parameters used for queries.

cdl kafka show status

Displays CDL status information.

Command Modes Exec

Syntax Description `cdl show status [db-name db_name]`

db-name db_name

Specify the name of the database.

Must be a string.

Default Value: session.

Usage Guidelines Use this command to view CDL status information.

cdl kafka ssl-settings

Configures Kafka SSL settings.

Command Modes Exec > Global Configuration

Syntax Description `cdl kafka ssl-settings enable-ssl { false | true }
disable-host-name-verification { false | true }`

disable-host-name-verification { false | true }

Specify to disable or enable host name verification during SSL communication between Kafka brokers.

Must be one of the following:

- false
- true

Default Value: false.

enable-ssl { false | true }

Specify to disable or enable SSL communication between Kafka brokers.

Must be one of the following:

- false
- true

Default Value: false.

Validation Condition: The parameter value can not be changed when system mode is set to running.

Usage Guidelines Use this command to configure Kafka SSL settings.

cdl label-config

Configures the list of label configuration for CDL pods.

Command Modes Exec > Global Configuration

Syntax Description **cdl label-config** *label_config_name*

label_config_name

Specify the label configuration name.

Must be a string.

Usage Guidelines Use this command to configure the list of label configuration for CDL pods.

cdl label-config endpoint

Configures endpoint label configuration.

Command Modes Exec > Global Configuration

Syntax Description **endpoint key** *label_key* **value** *label_value*

key label_key

Specify the label key.

Must be a string.

value label_value

Specify the label value.

Must be a string.

Usage Guidelines Use this command to configure the endpoint label configuration.

cdl label-config index map

Configures map-level label configuration for CDL index.

Command Modes Exec > Global Configuration

Syntax Description **index map map-id** *map_id*

key *label_key*

Specify the label key.

Must be a string.

map-id *map_id*

Specify the map-based label configuration.

Must be an integer.

value *label_value*

Specify the label value.

Must be a string.

Usage Guidelines Use this command to configure map-level label configuration for CDL index.

cdl label-config slot map

Configures map-level label configuration for CDL slot.

Command Modes Exec > Global Configuration

Syntax Description `slot map map-id map_id`

key *label_key*

Specify the label key.

Must be a string.

map-id *map_id*

Specify the map-based label configuration.

Must be an integer.

value *label_value*

Specify the label value.

Must be a string.

Usage Guidelines Use this command to configure map-level label configuration for CDL slot.

cdl logging

Configures the CDL logging parameters.

Command Modes Exec > Global Configuration

Syntax Description

```
cdl logging default-log-level default_log_level error-stack { false | true }  
sync-logging { false | true }
```

default-log-level *default_log_level*

Specify the default logging level for the CDL modules.

Must be one of the following:

- **debug**
- **error**
- **info**
- **off**
- **trace**
- **warn**

Default Value: warn.

error-stack { false | true }

Specify to enable or disable the stack trace along with error logs. To enable, set to true.;

Must be one of the following:

- **false**
- **true**

Default Value: false.

json-logging { false | true }

Specify to enable or disable logging in json format.

Must be one of the following:

- **false**
- **true**

Default Value: false.

sync-logging { false | true }

Specify to enable or disable synchronous logging. To enable, set to true.

Must be one of the following:

- **false**
- **true**

Default Value: false.

Usage Guidelines Use this command to configure CDL logging parameters.

cdl logging logger

Configures logs for each module.

Command Modes Exec > Global Configuration

Syntax Description `cdl logging logger cdl_module_name level default_log_level`

level *default_log_level*

Specify the logging level for each module.

Must be one of the following:

- **debug**
- **error**
- **info**
- **off**
- **trace**
- **warn**

cdl_module_name

Specify the name of the CDL module.

Must be a string.

Usage Guidelines Use this command to configure logs for each module.

cdl remote-site

Configures the list of remote sites.

Command Modes Exec > Global Configuration

Syntax Description `cdl remote-site system_id { db-endpoint { host host_name | port port_number } | kafka-server { host_name port_number } }`

system_id

Specify the remote system or site identity.

Must be an integer.

Usage Guidelines Use this command to configure the list of remote sites.

cdl remote-site db-endpoint

Configures the list of remote DB endpoints.

Command Modes Exec > Global Configuration

Syntax Description `cdl remote-site system_id db-endpoint { host host_name | port port_number }`

host *host_name*

Specify the remote DB endpoint's hostname or IP address.

Must be a string of 1-64 characters.

port *port_number*

Specify the remote DB endpoint's port number.

Must be an integer in the range of 0-65535.

Usage Guidelines Use this command to configure the list of remote DB endpoints.

cdl remote-site kafka-server

Configures the list of Kafka servers.

Command Modes Exec > Global Configuration

Syntax Description `cdl remote-site system_id kafka-server host_name port_number ssl-port ssl_port_number`

ssl-port *ssl_port_number*

Specify the SSL port exposed by remote Kafka Broker.

Must be an integer in the range of 0-65535.

host_name

Specify the host name or IP address of the remote Kafka server.

Must be a string of 1-64 characters.

port_number

Specify the remote Kafka server's port number.

Must be an integer in the range of 0-65535.

Usage Guidelines Use this command to configure the list of Kafka servers.

cdl ssl-config

Configures SSL configuration parameters.

Command Modes Exec > Global Configuration

Syntax Description `cdl ssl-config enable { false | true }`

enable { false | true };

Specify to enable or disable SSL configuration.

Must be one of the following:

- **false**
- **true**

Default Value: false.

Usage Guidelines Use this command to configure SSL configuration parameters.

cdl ssl-config certs

Configures database endpoint specific SSL configuration parameters.

Command Modes Exec > Global Configuration

Syntax Description `certs ip fqdn_ip_address ssl-key ssl_key_file ssl-ca ssl_ca ssl-crt ssl_cert_file`

host hostname_ip_address

Specify the hostname or IP address of external host.

Must be a string.

ssl-ca ssl_ca

Specify the CA certificate file content.

Must be a string.

ssl-crt ssl_cert_file

Specify the SSL Cert file.

Must be a string.

ssl-key ssl_key_file

Specify the SSL key file.

Must be a string.

Usage Guidelines Use this command to configure database endpoint specific SSL configuration parameters.

cdl zookeeper

Configures the Zookeeper for Geo Replication.

Command Modes Exec > Global Configuration

Syntax Description `cdl zookeeper { data-storage-size data_storage_size | { enable-JMX-metrics | enable-persistence } { false | true } | log-storage-size log_storage_size | replica ha_instances | log-level { debug | error | fatal | info | trace | warn } }`

data-storage-size *data_storage_size*

Specify the Zookeeper data storage size in GB.

Must be an integer in the range of 1-64.

Default Value: 20.

enable-JMX-metrics { false | true }

Specify to enable or disable JMX metrics.

Must be one of the following:

- **false**
- **true**

Default Value: true.

enable-persistence { false | true }

Specify to enable or disable persistence for Zookeeper data.

Must be one of the following:

- **false**
- **true**

Default Value: false.

log-level

Specify the Zookeeper logging level. The logging levels include debug, error, fatal, info, trace, and warn. The changes specified here will trigger rolling restart of Zookeeper pods.

Must be one of the following:

- **debug**
- **error**

- **fatal**
- **info**
- **trace**
- **warn**

log-storage-size *log_storage_size*

Specify the Zookeeper data log storage size in GB.

Must be an integer in the range of 1-64.

Default Value: 20.

replica *ha_instances*

Specify the number of HA instances to be created.

Must be an integer from the following: 1, 3.

Default Value: 3.

Validation Condition: The parameter value can not be changed when the system mode is set to running.

to-earliest

Specify to enable parallel pod management policy.

Must be one of the following:

- **false**
- **true**

Default Value: false.

Usage Guidelines

Use this command to configure the Zookeeper for Geo replication.