

CUDB Subscription Reallocation

FACILITY DESCRIPTION

Copyright

© Ericsson AB 2016, 2017. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

All trademarks mentioned herein are the property of their respective owners. These are shown in the document Trademark Information.



Contents

1	Introduction	1
1.1	Scope	1
1.2	Revision Information	1
1.3	Target Groups	2
1.4	Prerequisites	2
1.5	Typographic Conventions	2
2	Overview	3
2.1	Prerequisites	3
2.2	Architecture	3
2.3	Description	3
2.4	Dependencies and Interactions	7
3	Operation and Maintenance	11
3.1	Configuration	11
3.2	Fault Management	11
3.3	Performance Management	12
3.4	Security	12
3.5	Logging	12
	Glossary	13
	Reference List	15





1 Introduction

This document provides a description for the Administrative Data Redistribution feature.

1.1 Scope

The purpose of this document is to describe the function and operational conditions of the optional Administrative Data Redistribution feature, supported by the Ericsson Centralized User Database (CUDB).

1.2 Revision Information

Rev. A

This document is based on 12/155 34-HDA 104 03/9 with the following change:

- Section 2.3.1 on page 3: Updated with virtualization information.

Rev. B

Other than editorial changes, this document has been revised as follows:

- Updated virtualization terminology throughout the document.

Rev. C

Other than editorial changes, this document has been revised as follows:

- Section 2.4.8 on page 8: Added new section.

Rev. D

Other than editorial changes, this document has been revised as follows:

- Section 2.3.2 on page 4: Updated the reallocation logic description.
- Section 2.3.2.1 on page 4: Updated terminology and procedure.
- Section 2.3.3.1 on page 6: Updated reallocation requirements and conditions.
- Section 2.3.3.2 on page 7: Updated procedure.



1.3 Target Groups

This document is intended for system administrators and users working with the Administrative Data Redistribution feature.

1.4 Prerequisites

Users of this document must have a basic knowledge of CUDB and the Administrative Data Redistribution feature.

1.5 Typographic Conventions

Typographic conventions can be found in the following document:

- *Typographic Conventions*



2 Overview

CUDB supports administrative redistribution (also known as reallocation) of the data stored in the CUDB system. The feature allows moving stored data from one Data Store Unit Group (DSG) to other DSGs. In case of the proper handling of origin and target DSGs, reallocation can also be used for higher level movements, such as moving data between nodes, sites, or geographical zones.

Refer to *CUDB Data Distribution*, Reference [1] for more information on data distribution within CUDB.

2.1 Prerequisites

This section is not applicable to this feature.

2.2 Architecture

This section is not applicable to this feature.

2.3 Description

The following subsections describe the following topics of the reallocation service in more detail:

- The purpose of reallocation (see Section 2.3.1 on page 3)
- The logic of reallocation (see Section 2.3.2 on page 4)
- The types of reallocation (see Section 2.3.3 on page 5)

2.3.1 Purpose of Reallocation

The following scenarios may require reallocation:

- Unbalanced distribution of load is detected among the DSGs.
Note: In case of unbalanced load distribution, occupation balancing is recommended even if the occupancy level is not critical.
- The subscription information must be moved to new nodes in the system.
- A new, empty DSG is added to an existing CUDB system. Reallocation of DS entries is recommended in this case to level the DSGs, especially if the new DSG is added to provide additional capacity.



- An existing DSG is removed from an operating CUDB system. In this case, all entries must be moved to another DSG before the removal of the selected DSG.
- Subscribers are moved from one geographical zone to another one.
- DS entries are moved to less occupied DSGs to allocate space for new applications, services, or both.
- The subscriber profile is expanded, resulting in the movement of affected DS entries to other less occupied DSGs.

2.3.2 Logic of Reallocation

The reallocation process logic is based on the movement of Distribution Entry (DE) blocks, that is, on blocks of distributed data below a set of DEs, according to the DSG memory levels observed in the system to available and eligible DSGs. The memory level of a DSG is determined by the replica with the highest occupation. Specific destination DSG(s) can be disabled for provisioning of DEs, which means that it will not be eligible, before or during reallocation operation, and then reallocation will not be performed to such DSG(s) unless it is forced.

2.3.2.1 Reallocation of DE Blocks

Reallocation is performed on DE blocks, that is on blocks of distributed data below a set of DEs. Whenever a DE block is reallocated, the entries belonging to the DEs in the block are moved.

Each block states the maximum number of DEs whose distributed data can be moved in each reallocation iteration. Depending on the circumstances, the number of DEs whose distributed data can be moved are between 0 and the value of the `reallocationBlock` object, configured with the `reallocationBlockSize` parameter. In case there is any error with the distributed data of any of the DEs, the number of the DEs whose distributed data is moved is not a multiple of this parameter. See Section 3.1 on page 11 for more information on configuring reallocation.

During reallocation, the DE blocks are locked while the block is moved, thus no update traffic is allowed on them. Block lockdown is indicated with LDAP fault code 80, with the `Distribution entry is locked` diagnostic message. Due to the block lockdown, the value of `reallocationBlockSize` must be chosen carefully: a lower value means lower DE lockdown time and less traffic failure, but also longer reallocation time; at the same time, a higher value results in faster reallocation, but also in more traffic failure and longer DE lockdown time.

During reallocation, the first block of the DEs is moved to the destination DSG. Subsequent DE blocks are then moved to the destination DSG as long as its



memory level does not reach `stop` level or if destination DSG is not disabled for provisioning of DEs in the meantime.

2.3.2.2

DSG Memory Levels

CUDB supports obtaining information on the absolute occupancy percentage and relative memory level of all DSGs in the system. The memory level of a DSG corresponds to the memory level of the most occupied replica. The possible results in the relative memory level are the following:

- Below the `Eligible` level
- Below the `Stop` level
- Below the `Warning` level
- Above the `Warning` level
- Above the `Full` level

The DSG memory levels relevant to reallocation are as follows:

- `Full`

`Full` is a system wide parameter. If the occupancy level of a DSG is above the value of the parameter, the DSG reached full occupancy. Due to the adverse performance caused by full occupancy, a major alarm is raised if a DSG reaches this level. Once at `Full` level, addition of new data to the affected DSG is disabled, but modification of existing data is still permitted. If a DSG reaches full occupancy, it is strongly recommended to perform reallocation immediately to reduce the occupancy level.

- `Warning`

`Warning` is defined separately for each DSG, and indicates high occupancy level. Once at `Warning` level, an alarm is raised to warn of the high occupancy of the affected DSG. If a DSG reaches high occupancy, it is recommended to perform reallocation to reduce the occupancy level. Normally, if all DSGs of the CUDB system are below the `Warning` level, no reallocation is required. This memory level is lower than `Full`.

- `Eligible`

`Eligible` is defined separately for each DSG, and indicates that the DSG is available as the destination or target of reallocation. This memory level is lower than `Warning`.

- `Stop`

`Stop` is used to stop data redistribution for some types of reallocation. The level is calculated as the mid-point between `Warning` and `Eligible`.



2.3.3 Reallocation Types

Depending on the purpose, reallocation can be executed in two different ways, as described in the following sections.

2.3.3.1 Reallocation by Moving a Specified Percentage of DS Entries from a Specific Source DSG

Reallocation can be executed between a specified source DSG and a specified destination DSG, or to any available and eligible destination DSGs by moving a specific percentage of the DS entries of the specified source DSG. The reallocation requirement is that the memory level of the destination DSG(s) must be below `Eligible` level and the destination DSG(s) must not be disabled for provisioning of DEs unless reallocation is forced to such DSG(s). Reallocation will then attempt to move the specified portion of DS entries of the specified source DSG to the specified destination DSG, or to any available and eligible destination DSGs (if none were specified as parameter).

In case any available and eligible DSGs are used, the system analyzes the occupancy level of the DSGs located in the same geographical zone as the source DSG before the reallocation. If no geographical zones are configured, the system analyzes the occupancy of all DSGs. The analysis produces a list of the DSGs whose occupancy level is below the `Eligible` threshold. The DSG with the lowest occupancy level on the list is then chosen as the first destination DSG.

Reallocation of the DE blocks between the source and destination DSG(s) stops if one of the following conditions are met:

- The specified percentage of DS entries from the source DSG is moved away.
- The destination DSG (if one is specified), goes above the `Stop` memory level. In this case, further reallocation may still be possible if additional DSGs are available in the same geographical zone.
- The destination DSG(s) go(es) above the `Stop` memory level. In this case, reallocation is terminated, and the operation is reported as a `Partial success`, indicating that the amount of the moved DS entries from the source DSG has not reached the specified percentage value.
- If one destination DSG is specified, which is blocked for provisioning of DEs, and reallocation is not forced to such DSG.
- If no destination DSG is specified, one or more destination DSGs are blocked for provisioning of DEs, no other DSGs are available or eligible, and reallocation is not forced over blocked DSG(s).

The outcome of reallocation may be reported as follows:

- `Success`, if the requested percentage of DS entries of the source DSG has been moved away.



- `Partial success`, if the `Stop` memory level is reached in all available destination DSGs before reaching the requested percentage of moved DS entries in the source DSG.
- `Partial success`, if the `Stop` memory level is not reached in all the destination DSGs, but the requested percentage cannot be reached, as there are no more DS entries in the source DSG that could be moved.
- `Failure` if reallocation cannot move any DS entries from the source DSG.

2.3.3.2 Reallocation of a List of DS Entries to a Specific Destination DSG

Reallocation can also be executed by selecting a list of individual DS entries hosted either on the same or different DSG, then reallocating them to a specified destination DSG. The DSGs may be located either in the same or in a different geographical zone.

The destination DSG must not be disabled for provisioning of DEs, unless it is forced, and the occupancy of the destination DSG must be below the `Warning` memory level to execute this type of reallocation.

The list of DS entries is supplied in a file that contains the primary identity pointing to each DS entry to be reallocated (such as IMSI, MSISDN, and so on).

The outcome of reallocation may be reported as follows:

- `Success` if all entries are reallocated.
- `Partial success` if the occupancy of the destination DSG reaches the `Warning` memory level.
- `Failed` if no entries were successfully moved.

2.4 Dependencies and Interactions

This section describes the dependencies and interactions of the reallocation process towards other CUDB services.

2.4.1 LDAP Access

When reallocation is executed, affected DEs are locked while their distributed data is being moved. This prevents the modification or deletion of the distributed data entries below the DEs while reallocation is performed. Add, Modify, and Delete operations requested for locked DEs during the process are rejected with `LDAP fault code 80`.

When reallocation is executed to empty a DSG, it is recommended to block provisioning in the affected DSG before starting reallocation.



Refer to *CUDB LDAP Data Access*, Reference [2] for more information on Lightweight Directory Access Protocol (LDAP) access.

2.4.2 Reconciliation

Reallocation cannot be executed while reconciliation is ongoing. If reallocation is executed while reconciliation is running, the system returns an error.

Refer to *CUDB Data Storage Handling*, Reference [3] for more information on reallocation and reconciliation.

2.4.3 Backup

Although possible, execution of data backup is not recommended while reallocation is running. Backups are rather recommended to be executed after reallocation is completed.

Refer to *CUDB Backup and Restore Procedures*, Reference [4] for more information on backup and restore procedures.

2.4.4 Import/Export

Before executing an import/export operation, it is recommended to check that there is no reallocation in progress. Executing import/export operations during reallocation can lead to inconsistencies in the system.

Refer to *CUDB Import and Export Procedures*, Reference [5] for more information on import/export operations.

2.4.5 Software Upgrade

Reallocation can not be executed if a software upgrade is running. Likewise, software upgrades can not be initiated while a reallocation operation is running.

2.4.6 High Load in DS Alarm

Reallocation is a processing-intensive procedure that increases cluster load in the source and destination DSGs. If the affected clusters were already under load, the additional load caused by reallocation can result in raising the `High Load in DS` alarm.

2.4.7 Data Repair

Data Repair will fail for all distributed data entries below the reallocated DEs while reallocation is performed, similarly to LDAP access (see Section 2.4.1 on page 7). The concerned entries will be recorded in the `unrepaired` log.



2.4.8 Data Storage

After executing reallocation, the report of memory occupation in source DSG cluster may not be accurate without subsequent defragmentation. Refer to the *Running Defragmentation* section of *CUDB System Administrator Guide*, Reference [6] for information on the defragmentation procedure.





3 Operation and Maintenance

This section provides operation and maintenance related information on reallocation.

3.1 Configuration

This section provides basic information on configuring reallocation-related settings.

3.1.1 Configuring Reallocation

Reallocation is initiated with the `cudbReallocate` command. The command applies globally to the CUDB system, and provides a set of options that allows to run different reallocation types. Refer to *CUDB Node Commands and Parameters*, Reference [7] for more information on `cudbReallocate`.

3.1.2 Configuring DSG Occupancy Levels

The following DSG occupancy level parameters can be configured during run time:

- `Warning` level (configured for each DSG individually)
- `Eligible` level (configured for each DSG individually)
- The granularity of the number of DEs whose distributed data will be reallocated during a reallocation operation (configured globally on the CUDB system with the `reallocationBlock` parameter)

Note: The `Full` occupancy level is preconfigured by Ericsson personnel.

Modification of the occupancy levels does not affect already running operations, but takes effect for subsequent reallocation operations.

For further information on configuring occupancy levels, refer to *CUDB Node Configuration Data Model Description*, Reference [8].

3.2 Fault Management

This section is not applicable to this feature.



3.3 Performance Management

This section is not applicable to this feature.

3.4 Security

This section is not applicable to this feature.

3.5 Logging

During the processing of the `cudbReallocate` command, the following events are logged:

- Reallocation started for each source DSG and destination DSG.
- Reallocation finished for each source DSG and destination DSG.
- DS entries successfully reallocated.
- Reallocation of DS entries failed, reason is also stated.
- Reallocation finished due to master change in the source DSG, destination DSG, or Processing Layer Database (PLDB).

During the modification of the configuration, the following events can be logged:

- `memoryWarningThreshold` must be higher than `memoryEligibleThreshold`
- `memoryWarningThreshold` must be lower than `memoryFullThreshold`

Refer to *CUDB Node Logging Events*, Reference [9] for more details.



Glossary

For the terms, definitions, acronyms, and abbreviations used in this document, refer to *CUDB Glossary of Terms and Acronyms*, Reference [10].





Reference List

CUDB Documents

- [1] *CUDB Data Distribution*
- [2] *CUDB LDAP Data Access*
- [3] *CUDB Data Storage Handling*
- [4] *CUDB Backup and Restore Procedures*
- [5] *CUDB Import and Export Procedures*
- [6] *CUDB System Administrator Guide*
- [7] *CUDB Node Commands and Parameters*
- [8] *CUDB Node Configuration Data Model Description*
- [9] *CUDB Node Logging Events*
- [10] *CUDB Glossary of Terms and Acronyms*