

CUDB Application Integration Guide

User Guide

Copyright

© Ericsson AB 2016-2018. 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	Purpose and Scope	1
1.2	Revision Information	1
1.3	Prerequisites	2
1.4	Typographic Conventions	2
2	Overview	3
3	Data Management	5
3.1	LDAP Views	5
3.2	LDAP DIT	5
3.3	LDAP Schemas	5
3.4	Identities	6
3.5	Data Distribution	6
3.6	Storage Space Considerations	7
3.7	BLOB Configuration and Storage Options	8
4	LDAP Search Indexes	10
4.1	Candidate Queries for Search Index Optimization	10
4.2	Configuration	12
4.3	Storage Space Considerations	12
5	LDAP Users and Access Control	13
6	Notifications	14
7	Application Counters	15
8	Licensing	16
9	CUDB Subtree Search Optimization	17
9.1	CUDB Subtree Optimization Configuration	17
	Glossary	18
	Reference List	19





1 Introduction

This document provides information required by application Front Ends (FEs) to integrate with the Ericsson Centralized User Database (CUDB).

1.1 Purpose and Scope

This document provides the information required by applications FEs to integrate in CUDB and information regarding the services offered by CUDB to applications.

The intended audience for this document is network solution architects, application product managers and application engineers in charge of carrying out the integration of a specific application with CUDB.

1.2 Revision Information

Rev. A

This document is based on 7/1553-HDA 104 03/9 and contains the following changes:

- Content rearrangement throughout the document.
- [Overview](#) on page 3: Updated LDAP Data Views information.
- [Data Distribution](#) on page 6: Updated description of PLDB and the logic of accessing PLDB replicas by application FEs.
- [BLOB Configuration and Storage Options](#) on page 8: Updated with disk storage system information.
- [Licensing](#) on page 16: Added new section on Electronic SW Licenses.

Rev. B

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

- Terminology updated throughout the document.
- [LDAP Search Indexes](#) on page 10: Revised.

Rev. C

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



- [LDAP Users and Access Control](#) on page 13: Added reference to CUDB LDAP Interwork Description.

Rev. D

Editorial changes only.

Rev. E

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

- [Licensing](#) on page 16: Updated licensing information.

Rev. F

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

- [CUDB Subtree Search Optimization](#) on page 17: Added section regarding Subtree Search Optimization.
- [CUDB Subtree Optimization Configuration](#) on page 17: Added section regarding CUDB Subtree Optimization Configuration.

Rev. G

Editorial changes.

Rev. H

Editorial changes only.

1.3 Prerequisites

A general knowledge of CUDB is required to fully understand the concepts presented in this document. For general information on CUDB, refer to [CUDB Technical Product Description](#).

1.4 Typographic Conventions

Typographic Conventions can be found in the following document:

- [Typographic Conventions](#)



2 Overview

CUDB is a highly available, geographically redundant, real-time, distributed Subscriber-Centric Database exposed as a Lightweight Directory Access Protocol (LDAP) directory to applications (such as HLR, AuC, and HSS). CUDB is the central entity for data storage in Ericsson's offering providing the data layered architecture, that separates applications logic and data storage. With this architecture, separate dimensioning of processing and storing resources can be achieved by decoupling both logically and physically application/business logic and the subscriber data. Furthermore, it allows the unification of applications (such as HLR, AuC, HSS-EPC, and HSS-IMS) specific data under a common subscriber profile which is stored in a common database. With this architecture, application FEs, or applications from here on, hold the application business logic while CUDB provides the common database supporting the storage of the common subscriber profile for applications. For more details on CUDB functions, refer to [CUDB Technical Product Description](#).

CUDB provides a basic LDAP Directory Information Tree (DIT) extended by applications to store and query their data. The most basic information an application must provide to be integrated with the CUDB system is the data to be stored and managed by the application. The data is defined in terms of LDAP schemas that contain the types of information stored in CUDB. Due to the distributed nature of the CUDB database system, it is important for the application to understand how data are distributed and replicated in CUDB to organize its data more efficiently.

To optimize the storage space of CUDB, applications must take into consideration certain aspects of the way CUDB maps LDAP schemas into its internal structures.

CUDB provides mechanisms to optimize LDAP search queries by means of index definition.

Access to the data by applications using the CUDB LDAP interface requires the configuration of LDAP users, typically one per application, to be able to provide different types of access to different applications.

The LDAP Data Views function supports accessing stored data through customizable views. The function makes it possible for applications to access the data stored in CUDB through a custom tree structure (that is, a custom DIT) and a custom schema. Several LDAP data views can be defined to accommodate different kinds of application FEs.

LDAP users can be configured to either access CUDB using the "native" view or core DIT or to use one of the defined LDAP data views.

Note: The LDAP Data Views function can only be used if the Application Facilitator Value Package is available.

CUDB offers the possibility of defining custom notifications that may be sent to application FEs through the Simple Object Access Protocol (SOAP) protocol. The



data that would fire the custom notification and the set of extra conditions required for the notification to be sent together with the data to be included in the notification message are configurable.

CUDB also provides applications with the possibility of defining their own set of counters related to the data stored.

All these possibilities and the information to be provided by the applications are outlined in the following sections.



3 Data Management

This section contains details about the data management. For more information, refer to [CUDB LDAP Data Access](#).

3.1 LDAP Views

CUDB maintains views to support flexible data organization. If an LDAP user has an LDAP view assigned to it, then it can access the data through that view. Each view has its LDAP schema loaded into the CUDB, where entries are composed by mapping attributes from the default view, resulting in the building of a custom DIT. Each view has a mapping file, which describes how the view entries are composed from the attributes contained by the default view.

The LDAP Data Views function allows access to the data in different data structures through the LDAP interface.

All the attributes that are reachable through views should be already defined in the default view.

Once a user with an assigned view establishes a new LDAP bind, all the following LDAP operations are performed against the LDAP view selected for this user. LDAP data views are not available for export and import procedures, or for notifications. Notifications must be configured according to core DIT distinguished names and attributes, and data in SOAP messages is, likewise, identified by its core DIT distinguished names and attributes. Users having no view assigned will see the PL and DS data through the default LDAP tree. For users that have a view set through configuration, the given LDAP schema of that view will apply.

For more information on LDAP views, refer to [CUDB LDAP Data Views](#).

3.2 LDAP DIT

CUDB is exposed as an LDAP DIT to applications. CUDB provides a basic LDAP DIT to be extended by applications to store and query their data.

Normally, applications place their data under the predefined CUDB branches, but they can also create their own branches under the DIT root entry. For more details on the CUDB LDAP basic DIT, refer to [CUDB LDAP Data Access](#) and [CUDB LDAP Interwork Description](#).

3.3 LDAP Schemas

CUDB provides a typical LDAP *core* schema and a CUDB-specific LDAP schema with object classes and attributes that are used in the predefined entries that



build the LDAP basic DIT. These attributes and the object classes can also be used by applications for their own entries.

Applications can also provide their own schema files that define specific object classes and attributes to be used by the applications data. Schemas containing object classes and attributes common to several applications are also supported, provided that object classes and attribute definitions are not repeated in different LDAP schemas.

CUDB provides tools that automatically and transparently perform a translation between LDAP schema elements and the internal database structures used by CUDB to support the LDAP schemas required. The schemas provided by applications must comply with some requirements to be usable by these tools. Refer to [CUDB Application Schema Update](#) for more details.

LDAP schemas are typically provided during CUDB system installation, but new LDAP schemas can also be included once the system is up and running. With certain limitations, CUDB also offers the option of updating existing LDAP schemas, once an application has been integrated in CUDB. For more details on schema update, refer to [CUDB Application Schema Update](#).

For more information about LDAP schemas, refer to [CUDB LDAP Data Access](#).

3.4 Identities

Identities are special types of entries in CUDB that are applied to locate subscribers using network identifiers (such as MSISDN, IMSI, and so on) by means of the LDAP alias mechanism. Refer to [CUDB LDAP Interwork Description](#) for more details about identities.

Identities must be defined at installation or schema update time. A container for each type of identity is created automatically under a CUDB specific branch at installation and schema update time under a CUDB-specific branch. As identities are normally shared by different applications, it is required that a common LDAP schema containing the attributes and object classes used by identities is provided for all applications. Refer to [CUDB Application Schema Update](#) for more details about identities definition.

For more information about identities, refer to [CUDB LDAP Data Access](#).

3.5 Data Distribution

CUDB contains two types of storage layers in which applications place their data: the Processing Layer Database (PLDB) and the Data Store Unit Groups (DSGs).

- The PLDB is an indexing database replicated in certain nodes of all CUDB sites for the data stored in the DSGs. It also stores the complete set of subscriber identities. The PLDB may as well be used to store certain data which does not have massive size (for example common subscriber profiles) and is not related to one specific subscriber.



Note: The PLDB does not offer the level of scalability provided by the DSGs. Therefore, it is recommended to study carefully what is stored in the PLDB, so that its capacity is not exceeded in a specific deployment.

- The DSGs are partitions of the subscribers database. Each DSG holds the data for a group of subscribers. Subscribers are distributed at provisioning time in the DSGs available in the system. Depending on the level of redundancy of the system one or more replicas of each DSG, called Data Store (DS) Units are present in the CUDB system. Each DS is assigned to a CUDB node.

Note: To retrieve the data for a specific subscriber, a hop may be needed from the CUDB node where the request arrived to the node that contains the DS master replica where the data are stored.

Data that require extra performance in retrieval and is not massive in size is a good candidate to store in the PLDB. This is because the application FEs always access to a node that hosts a PLDB replica, and therefore no hops to other CUDB nodes are required when accessing PLDB data.

CUDB determines automatically if a certain LDAP entry is stored in the PLDB or in DSGs based on the position the entry has in the LDAP DIT. For details on which LDAP branches are stored in the PLDB and which LDAP branches are stored in DSGs, refer to [CUDB LDAP Interwork Description](#).

For more information on how data is distributed in CUDB, refer to [CUDB LDAP Data Access](#).

3.6 Storage Space Considerations

As stated previously, CUDB provides tools that automatically and transparently perform a translation between LDAP schema elements and the internal database structures used by CUDB to support the LDAP schemas required. Refer to [CUDB Application Schema Update](#) for further details.

For optimizing the storage space, the following information is provided regarding the mapping performed by CUDB from LDAP types into internal database types:

- For each entry in the LDAP DIT (of type other than identity) CUDB holds a row in a master table that contains the Distinguished Name (DN) of the entry among other data. Therefore, the longer the DN, the more space is taken by the entry in the database.
- For each identity entry, CUDB holds a row in a table in the PLDB for each identity type holding the values for the identity attributes (but not the complete DN).
- For each object class defined in any schema file, CUDB uses one table that holds one row for each LDAP entry containing that object class. If none of the attributes of the object class has a value for the entry, no row is present in the object class table. Therefore, it is recommended, to optimize the space



occupied by data, to group attributes used together by the application (when one of these attributes is present all of them in the object class are present) in the same object classes.

- Each object class table contains a key column for each row and one column for each single-attribute belonging to the object class. Multi-value attributes are stored in separate tables (one for each pair of object classes or attributes) and CUDB uses one row of these latter tables for each multi-value attribute value.

The mapping of LDAP attribute types into internal database types performed by CUDB tools is shown in [Table 1](#).

Table 1 Mapping of LDAP Attribute Types into Internal Database Types Performed by CUDB

LDAP attribute type	Internal Database type
Bit String	BIT (LEN)
Boolean	VARCHAR
Country String	CHAR(2)
Directory String	VARCHAR CHARACTER SET UTF8 ⁽¹⁾
GeneralizedTime	VARCHAR
NumericString	VARCHAR
DN	VARCHAR(512) CHARACTER SET UTF8
Integer	TINYINT/SMALLINT/MEDIUMINT/ BIGINT/INT
OctetString	VARBINARY
OctetString{M} where M<255	VARBINARY(255)
OctetString{M} where M <= 4 KB	VARBINARY(M)
OctetString{M} where 4 KB < M <= 64 KB	BLOB(M)
OctetString{M} where 64 KB < M <= 16 MB	MEDIUMBLOB(M)
OctetString{M} where 16 MB < M <= 4 GB	LOB(M)
Other	VARCHAR

(1) A UTF-8 encoding of a Unicode character set uses one to three bytes per character.

3.7 BLOB Configuration and Storage Options

CUDB offers the possibility of storing Binary Large Objects (BLOBs). Attributes of BLOB type are defined at installation time together with the storage support option, Random Access Memory (RAM) or the disk storage system, for them on a



per LDAP attribute basis. Refer to [CUDB Binary Large Object Attributes Management](#) for more details on how BLOBs are configured.

The decision on whether a certain BLOB attribute is to be kept on the disk storage system or in RAM must take into consideration both performance and storage capacity requirements. BLOB attributes on the disk storage system have higher access time than BLOBs in memory, but leave more main memory available for other data to be stored.



4 LDAP Search Indexes

CUDB allows the definition of indexes on LDAP attributes to optimize, that is, speed up, search requests containing indexed attributes in the search filter.

Search indexes are used by CUDB when an LDAP search is received with a filter that includes one or more indexed attributes, and certain conditions are fulfilled by the filter. Search indexes are used by CUDB for optimizing queries, so that entries complying with a specific filter are selected without the need of retrieving all the possible candidates (according to the base DN and scope parameters of the search) from the database to check if the filter is fulfilled.

All searches using indexed attributes benefit from these optimizations, but the optimization is especially relevant for distributed searches which take a lot of time and resources to execute. Distributed searches using indexed attributes run much faster and use up much less resources than equivalent searches using non-indexed attributes. For information about LDAP distributed searches, refer to [CUDB LDAP Data Access](#).

Note: The optimizations with indexes are only applicable for one level and subtree search operations.

Search indexes are not applicable to identity entries, even if these entries contain an indexed attribute.

Consider that LDAP queries that specify `search` or `always alias dereferencing` are not optimized as much as queries not specifying these options. This is because alias dereferencing in this fashion always requires extra search operations that cannot take advantage of the search indexes.

4.1 Candidate Queries for Search Index Optimization

In order to make full use of the optimization provided by search indexes, the following conditions must be fulfilled.

- Condition 1: The search must be subtree or one level.
- Condition 2: The search filter must contain at least one indexed attribute.
- Condition 3: The indexed attribute in the filter is inside an `equalityMatch`, `greaterOrEqual`, or `lessOrEqual` filter.
- Condition 4: The `equalityMatch`, `greaterOrEqual`, or `lessOrEqual` part of the filter is either by itself in the filter, or it is part of an outermost AND term.

Search subtree and onelevel operations with filters that do not fulfill all the conditions above can still benefit from some, although lower level, optimization if the filter contains at least one indexed attribute in it.



The following examples show a set of LDAP filters for LDAP subtree or onelevel search operations for a CUDB system in which attributes IMSI and IMSICH0 are configured as indexed attributes.

Example 1

```
' (IMSICH0=3333333) '
```

The search is optimized using the IMSICH0 attribute.

Example 2

```
' ( ! (IMSICH0=999) '
```

The search is not optimized by indexed attributes. Condition 4 is not fulfilled. The outermost term is NOT, not AND.

Example 3

```
' ( | (IMSICH0=333) (IMSICH0=999) '
```

The search is not optimized by indexed attributes. Condition 4 is not fulfilled. The outermost term is OR, not AND.

Example 4

```
' (IMSI>=3333) '
```

The search is optimized using the IMSI attribute.

Example 5

```
' (&(IMSI>=3333) (IMSI<=4000)) '
```

The search is optimized using the IMSI attribute.

Example 6

```
' ( ! (&(IMSI>=3333) (IMSI<=4000)) ) '
```

The search is not optimized by indexed attributes. Condition 4 is not fulfilled. Even though there is an AND term inside, the outermost term is NOT.

Example 7

```
' (&(IMSI>=3333) (IMSICH0=4000)) '
```

The search is optimized using the IMSI attribute (IMSICH0 is also an indexed attribute that could be candidate to optimize the search, but IMSI is used as it appears first in the filter).

Example 8

```
' (&(CSP=3) (IMSICH0=4000)) '
```



The search is optimized using the IMSICH0 attribute.

Example 9

```
' ( | (CSP=3) (&(IMSICH0=4000) (IMSI=4000)) ) '
```

The search is not optimized by indexed attributes. Condition 3 is not fulfilled. Even though there is an AND term inside, the outermost term is OR.

Example 10

```
' (IMSI=333*) '
```

The search is not optimized by indexed attributes. Condition 4 is not fulfilled. The filter is a Substring one.

4.2 Configuration

To configure Search Indexes, refer to [CUDB System Administrator Guide](#).

Note: These indexes must be defined before any entry with the indexed attributes is populated in CUDB. Otherwise, search operations may not work correctly.

4.3 Storage Space Considerations

Adding search indexes has an impact in the space taken by each LDAP entry in CUDB. For each search index a new column is added to the CUDB table that contains the DN's for each entry. If the entry has a value for the indexed attribute, the value is stored in this table and in the table that holds all the attributes of the entry. Attributes with value take double space if they are indexed. Depending on the type of the indexed attribute, a small overhead may be taken to hold the length, even if the entry does not have a value for the specific attribute.



5 LDAP Users and Access Control

To enable the applications to access the data stored in CUDB through the LDAP interface, one or several LDAP users must be defined with the corresponding credentials, configuration parameters, and Access Control Rules (ACR). Application FEs may bind to the CUDB LDAP interface with one of these users, then send the corresponding LDAP operations to CUDB with optional LDAP controls. Refer to Section 6.4 LDAP Controls, [CUDB LDAP Interwork Description](#). User groups can optionally be defined to share Access Control Rules between LDAP users.

For more details on LDAP User Management, refer to [CUDB LDAP Data Access](#) and [CUDB System Administrator Guide](#).



6 Notifications

CUDB offers the possibility of defining custom notifications to be sent to application servers when changes are detected in the data stored in the DSG. The notifications protocol is SOAP-based. For information on the specific notifications protocol, refer to [CUDB SOAP Interwork Description](#).

CUDB allows the configuration of which LDAP attributes are monitored to send notifications in case they change. CUDB also allows the specification of a condition set on other related data to decide if the notification must be sent. The specific data to be sent in the notification (which could be the monitored data and/or any other data) can also be configured. For information on how to configure notifications in CUDB, refer to [CUDB Node Configuration Data Model Description](#).

For more information on this feature, refer to [CUDB Notifications](#). In case the LDAP Data Views function is activated, see [LDAP Views](#) on page 5 for more information about notifications related to LDAP views.



7 Application Counters

CUDB provides a framework that allows applications to define their own set of counters to be generated by CUDB. The framework is based on the definition of scripts that calculate the desired counters by accessing the data stored in CUDB and output them in XML files.

For details on how to configure these counters, refer to [CUDB Application Counters](#).



8 Licensing

Contact Ericsson personnel to handle the licenses required for the integration of a new application.



9 CUDB Subtree Search Optimization

In CUDB, LDAP subtree search requests can be optimized by using specific information about the LDAP entries to be returned for specific subtree searches. The system can search more efficiently by accessing LDAP entries directly without scanning and filtering the database. Application FE queries do not have to be changed.

9.1 CUDB Subtree Optimization Configuration

CUDB subtree optimized searches are configured by YAML files on a per LDAP user basis. To configure subtree search optimization, refer to the Optimized Subtree Searches section of *CUDB System Administrator Guide*.



Glossary

For the terms, definitions, acronyms and abbreviations used in this document, refer to CUDB Glossary of Terms and Acronyms



Reference List

CUDB Documents

1. CUDB Technical Product Description
2. CUDB LDAP Data Access
3. CUDB LDAP Data Views
4. CUDB LDAP Interwork Description
5. CUDB Application Schema Update
6. CUDB Binary Large Object Attributes Management
7. CUDB System Administrator Guide
8. CUDB SOAP Interwork Description
9. CUDB Node Configuration Data Model Description
10. CUDB Notifications
11. CUDB Application Counters
12. CUDB LDAP Interwork Description
13. CUDB Glossary of Terms and Acronyms