CUDB Performance Guide

USER GUIDE

4/1553-HDA 104 03/10 Uen A

Contents

1 Introduction

This document describes the performance management solution provided by Ericsson Centralized User Database (CUDB).

1.1 Document Purpose and Scope

This document provides an overview of performance management in CUDB, describes available performance data and its generation, and how it can be collected and used to measure the performance of a CUDB node.

1.2 Revision Information

Rev. A	This document is based on 4/1553-HDA 104 03/9 with revised content and structure.

1.3 Target Groups

This document is intended for CUDB system operators who will be monitoring the performance of CUDB nodes and for solution architects and system integrators who will be integrating CUDBs performance management solution with a management system.

1.4 Prerequisites

The reader of this document should have general knowledge of CUDB. Knowledge of LDAP data access mechanisms and CUDB architecture is recommended for proper understanding of the CUDB performance data.

1.5 Typographic Conventions

Typographic conventions can be found in the following document:

- *Typographic Conventions*

2 Counters in CUDB

2.1 Overview

For each CUDB node, a set of counter groups is provided, containing performance data for the following:

- Individual Lightweight Directory Access Protocol (LDAP) servers
- Overall CUDB node performance
- Application groups
- Database clusters
- Simple Object Access Protocol (SOAP) Notifications

More details about the information provided by CUDB counters can be found in *CUDB Counters List*, Reference [1].

Note: As part of the integration of different application Front Ends (FEs), CUDB also provides the Application Counters Framework. The framework makes it possible for application FEs to have CUDB gather and publish performance management information about their application data stored in CUDB (on behalf of the application FEs). For more information about this framework please refer to *CUDB Application Counters*, Reference [3].

2.2 Counter Generation and Publishing

CUDB counters are generated and published independently on each CUDB node, and are available only on that node. They are not replicated to the rest of the CUDB system.

The generation of counter value samples and publishing of counter data are independent processes, with different execution periods:

- Generation period for cluster memory counters (`memoryUsage`) is 5 minutes
- Generation period for the rest of CUDB's own counters is 1 minute
- Publishing period for all counters is 15 minutes

Counters are published in 3GPP XML format and can be found in the following output location:

`/home/cudb/oam/performanceMgmt/output/`

The file format is described in *ESA XML Interface for Performance Management*.

Depending on counter type, the files contain the following information:

For `gauge` counters:

- The value of the last generated sample
- Maximum value in the publishing period
- Minimum value in the publishing period

For accumulated counters:

- The value of the last generated sample
- Delta value, compared with the value of the first sample of the publishing period

  Note: Drops of counter values for certain accumulated counters may happen in case of an LDAP FE restart. In that case, the delta value is not valid and waiting for the next counter publishing in necessary to get a valid delta value.

Attention!

Files are kept in the specified location for one day.

Counter users collect CUDB counter values by copying the generated files from the output location. It is recommended to retrieve output files with the `cudbadmin` user through SFTP protocol. Refer to *CUDB Users and Passwords*, Reference [5] for more information on user credentials.

2.3 Configuring Counter Output Files Names

The filenames of these counter output files are based on the following format:

`A<date>.<starttime>-<stoptime>-<jobname>_<networkElementName>.xml`

The variables in the above file name are the following:

<date>	The date of the measurement in format YYYYMMDD.
<starttime>	The start time of the measurement in format HHMM.
<stoptime>	The stop time of the measurement in format HHMM.

<jobname>	The job name of the measurement.
<networkElement Name>[1]	A string used as unique identity representing the node that runs the ESA.

(1) ESA refers to this variable as `uniqueId`.

networkElementName can be configured.

Refer to *ESA Performance Management*, Reference [7] for a complete description of the file names.

The ***<networkElementName>*** parameter is set through CUDB configuration CLI, by setting the value of the ***networkElementName*** configuration attribute. For more information, refer to the "Class CudbLocalNode" table of *CUDB Node Configuration Data Model Description*, Reference [4].

For more information on all the steps required to change and check the value of the ***<networkElementName>*** attribute, refer to the Object Model Modification Procedure section of *CUDB Node Configuration Data Model Description*, Reference [4].

After the ***<networkElementName>*** attribute is changed, restart the Performance Management Agent with the `cudbPmJobReload` command.

2.4 Effects of Structure and Configuration on CUDB Counters

In order to properly understand and interpret counter values, important aspects of CUDB data access, architecture, and features need to be taken into account. The relationship of the previous factors with CUDB counters is described in the following sections, as well as some general considerations.

2.4.1 Master Distribution

Due to the supported combinations of `readModeInDS` and `readModeInPL` configuration parameters, master DS replicas will receive much higher amounts of traffic compared to slave replicas within the same DSG.

This will be reflected in the following counter values:

- `intendedLdapRequests, DSn`
- `processedLdapRequests, DSn`

Master PLDB replicas may receive higher amounts of traffic compared to the slave replicas during provisioning. This will be reflected in the following counter values:

- `intendedLdapRequests, Pldb`

- `processedLdapRequests, Pldb`

If a node hosts multiple master replicas, the values of the following counters may be higher compared to nodes with fewer master replicas:

- `ldapTpsAtFrontEndn`
- `receivedLdapReqsTotal`
- `processedLdapReqsLocalNode`
- `notificationsSent`

For more information on `readModeInDS` and `readModeInPL`, refer to *CUDB Node Configuration Data Model Description*, Reference [4] and *CUDB LDAP Data Access*, Reference [2].

2.4.2 Distribution of Subscriber Profiles

Higher memory occupation in a DSG will typically result in its master replica receiving more traffic. In terms of CUDB counters, this means that DSG master replicas with higher `memoryUsage, Dsn` counter values may also have higher values than master replicas of other DSGs for the following counters:

- `intendedLdapRequests, DSn`
- `processedLdapRequests, DSn`

A higher active/inactive subscriber ratio in a DSG will also result in its master replica receiving more traffic. Such master replicas may have higher values of the same counters as listed above, compared to master replicas of other DSGs in the system.

2.4.3 Application FE Connections

The CUDB nodes that are the primary targets for Application FE connections will receive most of the traffic intended for a CUDB System. Depending on the master distribution in the system, such traffic may either end at the primarily affected nodes or be proxied to other nodes in the system.

If the CUDB nodes connected to Application FEs don't host many master replicas, they will have a high number of proxied requests, resulting in a higher value of `processedLDAPReqsRemoteNodes` than other nodes of the system.

If there are no nodes in the system with a high concentration of master replicas, nodes with Application FE connections will have higher values than other nodes in the system for the following application counters:

- `ldapTpsAtFrontEndn`
- `receivedLdapReqsTotal`

- `processedLdapReqsLocalNode`

Otherwise, nodes with a concentration of master replicas are expected to have the highest values for the listed counters.

2.4.4 Network Issues

Increased network latency can result in a higher number of failed proxied reques ts, such as in the increased value of `nonProcessedLdapReqsRemoteNodes`.

Network issues in communication with Notification end points can result in failed SOAP notifications, such as an increase of `notificationsFailed` counter values.

2.4.5 Overload Protection and Load Regulation

Incidents in the core network or on UDC solution level can cause high traffic and trigger the overload protection and load regulation mechanisms, resulting in an increased value of the dropped requests counters:

- `droppedLdapReqsLocalLdapLayer`
- `droppedLdapReqsLocalClusters`
- `droppedLdapRequests, Pldb`
- `droppedLdapRequests, Dsn`
- `nonProcessedLdapReqsRemoteNodes`
- `droppedAndFailedLdapReqsAppGrpn`
- `droppedAndFailedLdapReqsAppGrpn`

2.4.6 General Considerations

CUDB maintenance operations can impact local redundancy of a CUDB node or cause high network, storage, and processing load, resulting in an increase of dropped or failed requests.

Infrastructure problems or maintenance can impact the capacity and availability of network, storage, and processing resources, resulting in an increase of dropped or failed requests.

Glossary

For the terms, definitions, acronyms, and abbreviations used in this document, refer to *CUDB Glossary of Terms and Acronyms*, Reference [6].

Reference List

CUDB Documents

[1] *CUDB Counters List*

[2] *CUDB LDAP Data Access*

[3] *CUDB Application Counters*

[4] *CUDB Node Configuration Data Model Description*

[5] *CUDB Users and Passwords*, 3/006 51-HDA 104 03/9

[6] *CUDB Glossary of Terms and Acronyms*

Other Ericsson Documents

[7] *ESA Performance Management*