

Storage Engine, Execution of Selective Replica Check Failed, DS, Major

Ericsson Centralized User Database

Operating Instructions

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 | Alarm Description | 1 |
| 1.2 | Prerequisites | 2 |
| 2 | Procedure | 4 |
| | Glossary | 5 |
| | Reference List | 6 |



Storage Engine, Execution of Selective Replica Check Failed, DS, Major



1 Introduction

This instruction concerns alarm handling for the Storage Engine, Execution of Selective Replica Check Failed, DS, Major alarm.

1.1 Alarm Description

The alarm is raised when Selective Replica Check was not able to retrieve all applicable entries when analyzing operational logs for the database cluster .

For further information about Selective Replica Check, refer to [CUDB Data Storage Handling](#).

The alarm is issued in the following situations:

- When Selective Replica Check has encountered time gaps in the in the operational logs .

The possible alarm causes and the corresponding fault reasons, fault locations and impacts are described in [Table 1](#) .

Table 1 Alarm Causes

| Alarm Cause | Description | Fault Reason | Fault Location | Impact |
|------------------------------------|---|--|--|---|
| Time gaps in the operational logs. | Certain events in the CUDB can cause failures when writing in the operational logs. | If a process that is responsible for writing in an operational log, crashes or is in any way prevented from writing, not all relevant information is provided in the operational logs. | Operational logs on the former master replica. | There is a possibility that all the changes to the data in the database cluster were not detected by Selective Replica Check. |

The following are the consequences for the node if the alarm is raised:

- Possible data durability issues in the affected Data Store Unit Group (DSG). Eventually, this can have a service impact for certain subscribers.

The alarm attributes are listed and explained in [Table 2](#):

Table 2 Alarm Attributes

| Attribute Name | Attribute Value |
|------------------|---|
| Auto Cease | No |
| Module | STORAGE - ENGINE |
| Error Code | 27 |
| Timestamp First | Date and time when the alarm was raised for the first time. |
| Repeated Counter | Number which indicates how many times the alarm was raised. |



| Attribute Name | Attribute Value |
|------------------------------|--|
| Timestamp Last | Date and time of the most recent alarm raised. |
| Resource ID | .1.3.6.1.4.1.193.169.1.1.27.<DS>.<Timestamp> |
| Alarm Model Description | Execution of Selective Replica Check Failed, Storage Engine. |
| Alarm Active Description | Storage Engine (DS-group #<DS>): Selective Replica Check failed, task <Task ID>. It was not possible to retrieve all applicable entries. |
| ITU Alarm Event Type | processingErrorAlarm (4) |
| ITU Alarm Probable Cause | databaseInconsistency (160) |
| ITU Alarm Perceived Severity | (4) - Major |
| Originating Source IP | Node ID where the alarm was raised. |
| Sequence Number | Number which indicates the order in which alarms were raised. |

In [Table 2](#), the indicated variables are as follows:

- <DS> is the DSG this cluster belongs to.
- <Timestamp> is the Unix epoch in seconds representing the time of the incident, that is the timestamp which is used to determine where, in the operational logs of the former master Selective Replica Check starts looking for modified data.
- <Task ID> is an identifier for the Selective Replica Check processes for an individual DS.

For more information about attributes description, refer to [CUDB Node Fault Management Configuration Guide](#).

1.2 Prerequisites

This section provides information on the documents, tools, and conditions that apply to the procedure.

1.2.1 Documents

Before starting this procedure, ensure that you have read the following documents:

- [CUDB Node Fault Management Configuration Guide](#), regarding alarm configuration.
- [System Safety Information](#)
- [Personal Health and Safety Information](#)



1.2.2 Tools
Not applicable.

1.2.3 Conditions
Not applicable.



2 Procedure

Do the following:

Steps

1. If the Storage Engine, Unable to Synchronize Cluster in DS, Major alarm is raised when the Self-Ordered Backup and Restore function is not enabled or fails to restore the replication automatically, follow the procedure described in [Storage Engine, Unable to Synchronize Cluster in DS, Major](#).
2. Cease the alarm manually.

Results

Note: The procedure is to fix data inconsistency among master and slave replicas, but it cannot guarantee the full repair of system data.



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 Data Storage Handling
2. CUDB Node Fault Management Configuration Guide
3. Storage Engine, Unable to Synchronize Cluster in DS, Major
4. CUDB Glossary of Terms and Acronyms

Other Ericsson Documents

1. System Safety Information
2. Personal Health and Safety Information