

# Distributed Storage Alarm

Cloud Execution Environment

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

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Alarm Description	1
1.2	Prerequisites	2
<b>2</b>	<b>Procedure</b>	<b>3</b>
2.1	Actions	3
2.2	Advanced Log Collection	3
<b>3</b>	<b>Additional Information</b>	<b>5</b>





# 1 Introduction

This instruction concerns alarm handling for ScaleIO® related alarm events in the Cloud Execution Environment (CEE), in case a managed ScaleIO solution is used. In case of unmanaged ScaleIO, fault management is handled outside of CEE, so in CEE no ScaleIO alarms are received. For more information on managed and unmanaged ScaleIO, refer to section [Introduction](#) in the [ScaleIO Architecture Description](#).

The below sections explain how to identify fault location and get detailed alarm information. The troubleshooting procedure uses tools provided by ScaleIO.

**Note:** This document contains generic instructions for all ScaleIO alarms. There are no separate OPIs for each individual ScaleIO alarm.

## 1.1 Alarm Description

Distributed storage alarms are issued for ScaleIO alarm events by the Managed Object (MO) `DistributedStorage`.

The severity of the alarm is `CRITICAL`, `MAJOR`, or `MINOR`.

For each ScaleIO alarm a Watchmen alarm is mapped in CEE. Identify the relevant ScaleIO alarm with the `Additional Text` field in the alarm message, see Table 1.

ScaleIO provides the following severities: `ALERT_LOW`, `ALERT_MEDIUM`, and `ALERT_HIGH`. For ScaleIO alarm integration into CEE, and to enable correct alarm handling, alarms are mapped in the following way:

ScaleIO Severity Level	CEE Severity Level
<code>ALERT_LOW</code>	<code>MINOR</code>
<code>ALERT_MEDIUM</code>	<code>MAJOR</code>
<code>ALERT_HIGH</code>	<code>CRITICAL</code>

For more information on the different ScaleIO alarms, see Section 3 on page 5.

The following are the consequences for the system if the alarm is not solved:

- Data or data protection can be lost.
- Redundancy can be lost.
- System performance can be reduced.
- Storage capacity can be degraded.
- ScaleIO licenses expire.



The alarm attributes are listed in Table 1.

Table 1 Alarm Attributes

Attribute Name	Attribute Value
Major Type	193
Minor Type	2031719-2031816
Managed Object Class	AlarmType
Managed Object Instance	Region=<name_of_the_region>, Equipment=1, DistributedStorage=<storage_system_name>
Specific Problem	Distributed Storage Alarm
Event Type	<alarm_specific_event_type>
Probable Cause	<alarm_specific_probable_cause>
Additional Text	ScaleIO Alarm: <scaleio_alarm_name> <sup>(1)</sup>
Severity	CRITICAL (3), MAJOR (4), or MINOR (5)

(1) The alarm name is based on the ScaleIO REST alarm message, listed in Table 2, Section 3 on page 5. For example, the additional text for ScaleIO alarm "LICENSE\_EXPIRED" is "ScaleIO Alarm: License Expired".

See Table 2 for detailed alarm information and recommended actions.

## 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:

- Dell EMC ScaleIO Version 2.x Security Configuration Guide
- Part 4 Reference in Dell EMC ScaleIO Version 2.x User Guide

### 1.2.2 Tools

No tools are required.

### 1.2.3 Conditions

Before starting this procedure, ensure that the following conditions are met:



- The virtual IP address of the Meta Data Manager (MDM) cluster is known.
- Information about how to connect to the MDM master and how to use the `scli` commands is available.
- Administrator user role information is available.
- `superuser` privileges are available.

## 2 Procedure

This section describes the procedure to follow when this alarm is received.

### 2.1 Actions

Follow the below procedure:

1. Run the log collection script for MDM, SDS (ScaleIO Data Server), or SDC (ScaleIO Data Client):

```
/opt/emc/scaleio/<mdm|sds|sdc>/diag/get_info.sh -d  
<output_file.zip>
```

For example, for MDM the command is:

```
/opt/emc/scaleio/mdm/diag/get_info.sh -d <output_file.zip>
```

2. Fetch the `<output_file.zip>` file.
3. Identify the received alarm from the output and check the ScaleIO log files to get an overview about the fault. Perform the recommended action in Table 2.
4. If the alarm ceases, exit this procedure.

If the alarm remains, continue with the advanced log collection steps described in Section 2.2 on page 3.

### 2.2 Advanced Log Collection

This section describes how to collect detailed ScaleIO system log files.

1. Retrieve the ScaleIO component logs for each component.

For manual log collection, continue with Step 2.

For automated log collection, continue with Step 3.



2. For manual log collection follow the below procedure:

- a. Issue the below command:

```
/opt/emc/scaleio/<scaleio_component>/diag/get_info.sh -f
```

- b. If the selected node is the Master MDM, use the flags -u <mdm\_user>, and -p <mdm\_password> instead of -f.

If the selected node contains more than one ScaleIO component, running any script gathers logs for all components on that node.

- c. Verify that you get an output similar to the following, which shows that the process of log collection has completed successfully:

```
Archive /tmp/scaleio-getinfo/getInfoDump.tgz created successfully
```

When the log collection process completes, a ZIP file containing the logs of all ScaleIO components is created in the node.

To check the options for get\_info.sh, use the below command:

```
get_info.sh --help
```

3. For automated log collection follow the below procedure:

- a. Log in to the Installation Manager (IM):

— From Internet Explorer, browse to `https://<im_server_ip>`, where <im\_server\_ip> is the IP address of the Gateway (GW) server, with the Gateway or IM package installed on it.

— From the IM welcome screen, enter the IM credentials.

- b. From the IM main menu, select **Maintain**.

- c. Enter the login credentials and click **Retrieve system topology**. The **Maintenance operation screen** displays the system topology.

- d. Click **Collect logs**.

- e. Enter the MDM admin password and select one of the following options:

— **Copy repositories:** Besides full logs, it includes the MDM repository from installed components. Usual size for the repository is ~306 MB.

— **Collect exceptions only:** Instead of full logs it collects the ScaleIO core dumps only. By default, ScaleIO core dumps are disabled. Contact customer support if ScaleIO core dumps need to be enabled.





- **Last version:** Instead of full logs it collects the most recently created log file from each component.
  - Click **Collect Logs** to start the Get Info operation, which is the log collection function of the IM. This operation may take some time to complete, depending on your system topology. Once started, this operation can be rolled back.
- f. Select the **Monitor** tab to display log collection progress.
  - g. When the Get Info operation completes, click **Download logs** to download the log files. A ZIP file, containing all ScaleIO component logs, is downloaded.
  - h. Click **Mark operation completed** to clear the log files from the IM and enable it to be available for other operations.
4. Use the collected information and perform the recommended action in Table 2 to solve the fault.
  5. If the alarm ceases, exit this procedure.  
  
If the alarm remains, collect troubleshooting data as described in the [Data Collection Guideline](#) and contact the next level of maintenance support.
  6. The job is completed.

## 3 Additional Information

Table 2 lists alarms that can be received from the ScaleIO system and provides recommended actions for each alarm.

**Note:** Use the GUI only for alarm monitoring. For acting on alarms and performing the recommended actions described below, use CLI.

Table 2 ScaleIO Alarms and Recommended Action

ScaleIO Alarm Message in GUI	ScaleIO Alarm Message in REST	CEE Severity Level	Recommended Action
License expired	LICENSE_EXPIRED	CRITICAL	To resume operational mode, contact the next level of maintenance support for license renewal. If you have already renewed your license, install it.



Table 2 ScaleIO Alarms and Recommended Action

ScaleIO Alarm Message in GUI	ScaleIO Alarm Message in REST	CEE Severity Level	Recommended Action
The system's license will expire in n days	LICENSE_ABOUT_TO_EXPIRE	MAJOR or MINOR according to time left and limits	Contact the next level of maintenance support for license renewal. If you have already renewed your license, install it.
ScaleIO is using a trial license	TRIAL_LICENSE_USED	MINOR	Purchase a license and install it.
Oscillating failures reported	OBJECT_HAS_OSCILLATING_FAILURES	MINOR	Check oscillating failures of the component and take action accordingly. If the oscillating failure does not indicate a problem, change the settings of the oscillating failure window to suppress this alarm.
There are oscillating network failures	OBJECT_HAS_OSCILLATING_NETWORK_FAILURES	MINOR	Check the oscillating failure report that can be accessed from one of the management interfaces. Check if there is a problem with network links, fix it, and restart the counters.
No valid MDM credentials are configured in ScaleIO Gateway	GW_CONFIGURATION_INVALID_MDM_CREDENTIALS	CRITICAL	Configure the MDM credentials in the ScaleIO gateway using SioGWTool.
MDM credentials are not configured in the ScaleIO Gateway	MDM_CREDENTIALS_ARE_NOT_CONFIGURED	CRITICAL	Configure MDM credentials on the ScaleIO gateway using SioGWTool.
The MDM user configured in ScaleIO Gateway requires a password change	GW_USER_REQUIRES_PW_CHANGE	CRITICAL	Configure MDM credentials on the ScaleIO gateway using SioGWTool.
System upgrade is in progress	UPGRADE_IN_PROGRESS	MAJOR	Monitor the upgrade process, and check that it is completed successfully.
ScaleIO Gateway version is too old	GW_TOO_OLD	CRITICAL	Upgrade the ScaleIO gateway to the same version as the rest of your system.



Table 2 ScaleIO Alarms and Recommended Action

ScaleIO Alarm Message in GUI	ScaleIO Alarm Message in REST	CEE Severity Level	Recommended Action
The MDM is not operating in clustered mode	MDM_NOT_CLUSTERED	CRITICAL	MDM cluster was manually set to SINGLE mode. Confirm that this is an expected operation. Working in SINGLE mode is not recommended. Prepare the cluster modules if needed, and return to CLUSTER mode.
MDM fails over frequently	MDM_FAILS_OVER_FREQUENTLY	CRITICAL or MAJOR or MINOR  according to disconnect count and hard coded values (2/3/10)	The MDMs frequently swap ownership. Check the IP network connections and configuration.
Forward rebuild cannot proceed	FWD_REBUILD_STUCK	CRITICAL or MINOR	Check the system for lack of spare capacity and failed capacity, and either fix the problem or add capacity if necessary.
Backward rebuild cannot proceed	BKWD_REBUILD_STUCK	CRITICAL or MINOR	Check the system for lack of spare capacity and failed capacity, and either fix the problem or add capacity if necessary.
Rebalance cannot proceed	REBALANCE_STUCK	CRITICAL or MAJOR or MINOR	Add a physical disk; if this is not possible, reduce the spare policy while maintaining enough spare to sustain a rebuild, if necessary.
The MDM cluster is degraded, and data is not protected	CLUSTER_DEGRADED	CRITICAL or MAJOR	Check that all MDM cluster nodes are functioning correctly, and fix and replace faulty nodes, if necessary, to return to full protection.
	MDM_CONNECTION_LOST	CRITICAL	Check the connection to the MDM.



Table 2 ScaleIO Alarms and Recommended Action

ScaleIO Alarm Message in GUI	ScaleIO Alarm Message in REST	CEE Severity Level	Recommended Action
Inactive Protection Domain	PD_INACTIVE	MINOR	Protection Domain was inactivated by a user command. Confirm that this is an expected operation. This is usually done for maintenance. When maintenance is complete, reactivate the Protection Domain.
Storage Pool has failed capacity	STORAGE_POOL_HAS_FAILED_CAPACITY	CRITICAL	For the given storage pool, for some blocks, both primary and secondary copies are inaccessible. Check and fix the state of all devices in the storage pool and all holding devices of the server in the storage pool.
Storage Pool has degraded capacity	STORAGE_POOL_HAS_DEGRADED_CAPACITY	MAJOR	For the given storage pool, for some blocks, one of the two copies is inaccessible. Check if a server is offline or if there is another server hardware-related issue. Check if a storage device is down.
Capacity utilization above critical threshold	CAPACITY_UTILIZATION_ABOVE_CRITICAL_THRESHOLD	CRITICAL	Due to thinly provisioned volumes or snapshot usage, the capacity utilization of the storage pool is reaching a critical threshold. Remove snapshots, if possible, or add physical storage.
Capacity utilization above high threshold	CAPACITY_UTILIZATION_ABOVE_HIGH_THRESHOLD	MAJOR or MINOR	Due to thinly provisioned volumes or snapshot usage, the capacity utilization of the storage pool is reaching a high threshold. Remove snapshots, if possible, or add physical storage.
Failure recovery capacity is below the threshold	FAILURE_RECOVERY_CAPACITY_BELOW_THRESHOLD	MAJOR	The capacity available for recovery in a degraded storage event is lower than the predefined threshold. Replace failed hardware or add more physical storage.
Configured spare capacity is smaller than largest fault unit	CONFIGURED_SPARE_CAPACITY_SMALLER_THAN_LARGEST_FAULT_UNIT <sup>(1)</sup>	MINOR	Increase the spare percentage configured in the system for the storage pool, so that the capacity reserved for failure recovery is larger than the largest fault unit in the storage pool.



Table 2 ScaleIO Alarms and Recommended Action

ScaleIO Alarm Message in GUI	ScaleIO Alarm Message in REST	CEE Severity Level	Recommended Action
The Storage Pool relies too heavily (over 50%) on capacity from a single SDS or Fault Set. Balance capacity over other SDSs or Fault Sets.	STORAGE_POOL_UNBALANCED	MAJOR	Move some physical disks from the large SDS to the others, or add disks to the smaller SDS in order to approximate the capacity of the large SDS as much as possible.
Storage Pool does not meet the minimum requirement of 3 fault units	NOT_ENOUGH_FAULT_UNITS_IN_SP	MAJOR	Add more SDSs to the storage pool to meet the minimum requirement of three hosts.
There are cluster certificates pending approval. For more information, open <b>System Settings&gt;Certificates</b> .	UNTRUSTED_CERTIFICATE	MAJOR	Check the certificates and proceed at your own risk. For more information, open <b>System Settings&gt;Certificates</b> .
Master MDM Certificate is about to expire	CERTIFICATE_ABOUT_TO_EXPIRE	CRITICAL	Install a valid SSL certificate on the MDM before the old one expires.
Master MDM Certificate has expired	MDM_CERTIFICATE_EXPIRED	CRITICAL	Install a valid SSL certificate on the host.
Secure connection disabled on MDM	MDM_SECURE_CONNECTION_DISABLED	CRITICAL	Enable secure connections on the MDM to protect your login information.
The self-signed certificate presented by the Master MDM is not trusted	MDM_SELF_SIGNED_CERTIFICATE_NOT_TRUSTED	CRITICAL	Check the certificate, and proceed at your own risk.
MDM does not support secure connections	MDM_SECURE_CONNECTION_NOT_SUPPORTED	CRITICAL	Check MDM cluster nodes.



Table 2 ScaleIO Alarms and Recommended Action

ScaleIO Alarm Message in GUI	ScaleIO Alarm Message in REST	CEE Severity Level	Recommended Action
The validity period of the certificate presented by the Master MDM starts in the future	MDM_CERTIFICATE_NOT_YET_VALID	CRITICAL	The time and date on the computer where the certificate was created is not consistent with the time and date set in the ScaleIO system. Replace the certificate or fix the system time.
The Certificate Authority that signed the Master MDM's certificate is not trusted	MDM_CA_SIGNED_CERTIFICATE_CA_NOT_TRUSTED	CRITICAL	Proceed at your own risk.
SDS is disconnected	SDS_DISCONNECTED	MAJOR	The SDS service can be down or unreachable over the network. Verify that the SDS service is up and running and that the network is properly connected.
SDS disconnects frequently	SDS_DISCONNECTS_FREQUENTLY	MAJOR or MINOR  according to disconnect count and hard-coded	The SDS connection is fluctuating due to an unstable network connection. Check the SDS data network connection for packet drops, and try to disconnect one of the ports to see if the SDS disconnection issue is resolved by using only one port. If this does not resolve the issue, switch to the other port. If there is still an issue, it can be due to a faulty NIC, faulty switch ports, or a faulty switch. If there is no issue with another switch, the issue was switch-related. Otherwise, the issue can be due to a faulty NIC, which requires NIC replacement.



Table 2 ScaleIO Alarms and Recommended Action

ScaleIO Alarm Message in GUI	ScaleIO Alarm Message in REST	CEE Severity Level	Recommended Action
Memory allocation for RAM Read Cache failed on SDS	SDS_RMCACHE_MEMORY_ALLOCATION_FAILED	MINOR	The system failed to allocate memory to the SDS RAM read cache. For 32 GB RAM or less, up to 50% of the memory can be allocated for caching. From 32 GB or more, up to 75% of the memory can be allocated for caching. Reduce the configured RAM read cache memory to match the allocation conditions.
DRL mode: Hardened	DRL_MODE_NON_VOLATILE	MINOR	DRL mode is configured to Hardened instead of Volatile. Both modes are configurable.
RFcache card I/O error	RFCACHE_CARD_IO_ERROR	MINOR	Disable caching on the device and check the health of the device, because it can be faulty. If necessary, replace the device.
RFcache skipped due to heavy load	RFCACHE_CACHE_SKIPPED_DUE_TO_HEAVY_LOAD	MINOR	Read flash cache is working under a heavy load, and therefore has skipped some IOs. This is a temporary error which can resolve itself. If it persists, try to balance the storage pool contents across more SDSs, or add more cache cards.
RFcache IO stuck error	RFCACHE_IO_STUCK_ERROR	MINOR	IO has become stuck on the cache device. Disable caching on the device and check the health of the device, because it can be faulty. If necessary, replace the device.
RFcache resources are low	RFCACHE_LOW_RESOURCES	MINOR	There is not enough RAM available on the server for read flash cache optimal operation. Increase the amount of available RAM.
RFcache driver path is invalid	RFCACHE_INVALID_DRIVER_PATH	MINOR	The read flash cache driver xcachefs is either not installed, or was installed in the wrong location. Install the driver, and contact the next level of maintenance support if the problem persists.
RFcache source configuration is inconsistent	RFCACHE_INCONSISTENT_SOURCE_CONFIGURATION	MINOR	Check RFcache state of all disks in the pool and adjust them so that all disks have the same caching state.



Table 2 ScaleIO Alarms and Recommended Action

ScaleIO Alarm Message in GUI	ScaleIO Alarm Message in REST	CEE Severity Level	Recommended Action
RFcache source configuration is inconsistent	RFCACHE_INCONSISTENT_CACHE_CONFIGURATION	MINOR	Query the system to determine what is not consistent in the configurations of the read flash cache driver and the SDS where the cache device is located.
RFcache device does not exist	RFCACHE_DEVICE_DOES_NOT_EXIST	MINOR	You tried to add a cache device that does not exist. Check and fix read flash cache configuration.
RFcache API mismatch	RFCACHE_API_ERROR_MISMATCH	MINOR	The read flash cache xcache driver version and SDS version do not match. Upgrade them to the same version. If the problem persists, contact the next level of maintenance support.
SDS is in Maintenance Mode	SDS_IN_MAINTENANCE	MINOR	The SDS is currently in maintenance mode. Exit maintenance mode once maintenance is complete. If a non-Disruptive Upgrade (NDU) is in progress, ignore this warning.
Device failed	DEVICE_FAILED	MAJOR	The SDS device cannot be opened, read from or written to. Validate the device state. Check the cause of the error, and determine if it is a human error or a system malfunction. Check hardware if needed.
Device test is done and device is pending activation	DEVICE_PENDING_ACTIVATION	MINOR	The SDS device has been added and tested. Activate the SDS device.
Device has fixed read errors	FIXED_READ_ERROR_COUNT_ABOVE_MINOR_THRESHOLD	MAJOR if counter > 0	Read from the SDS device failed. Data was corrected from the other copy. No action is required, but note that the device can be faulty.
Device has fixed read errors	FIXED_READ_ERROR_COUNT_ABOVE_CRITICAL_THRESHOLD	CRITICAL if counter >= 5	SDS device read failed more than five times. Replace the physical device.





Table 2 ScaleIO Alarms and Recommended Action

ScaleIO Alarm Message in GUI	ScaleIO Alarm Message in REST	CEE Severity Level	Recommended Action
The SDC is either down or unreachable over the network	SDC_DISCONNECTED	MAJOR	Verify that the SDC service is up and running and that the network is properly configured and connected.
No more SDCs can be defined on this system; the maximum has been reached	SDC_MAX_COUNT	MAJOR	The maximum number of SDCs in the system has been reached.

(1) Due to SNMP related character limits in Watchmen, Additional Text for this alarm is "ScaleIO Alarm: Spare Capacity Smaller Than Largest Fault Unit".