

Troubleshoot, Manage and Monitor Hard Disk Drive Medium Errors

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[Handle HDD Medium Errors](#)

[HDD Role](#)

[Grown Defects \(G-list\) HDD Level](#)

[RAID Controller Role](#)

[Patrol Read](#)

[Consistency Check](#)

[Conditions when a RAID Controller Cannot Repair a Medium Error](#)

[Related Information](#)

Introduction

This document describes different types of disk errors, how to classify them, and tools you can use to identify them.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used


The information in this document is based on hard disks in Unified Computing System (UCS).

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Background Information

The document also outlines the Hard Disk Drive's (HDD) and Redundant Array of Independent Disks

(RAID) controller's role when you identify medium errors on the drives.

 **Note:** Medium errors are also referred to as media errors

Handle HDD Medium Errors

What causes HDD medium errors?

The most common cause of medium errors is poor signal amplitude that results in

- Unreliable Logical Bus Address (LBA) read location. Sometimes recoverable with multiple retries.
- Transient conditions, high fly writes caused by soft particles.
- Transient conditions caused by temporary shock, vibration, or acoustic events that result in off-track writes.
- Poor error map function in HDD manufacture that results in **padding** the current primary defect locations.

How does the HDD detect the medium error?

Step 1. The HDD periodically performs Background Media Scans to detect errors.

Step 2. The HDD tries to read from the media and for some reason is unable to retrieve the data that was written.

Step 3. When the HDD is unable to retrieve data that was written it invokes the HDD recovery code which will try various error recovery steps to successfully read the data from the media.

Step 4. If all the recovery steps fail the drive will generate a **03/11/0x** error back to the host and the LBA(s) will be placed on the **pending defect list**.

How does the Raid controller detect medium errors?

- The RAID controller will encounter medium errors while **Patrol Reads, Consistency Checks, Normal Reads, Rebuilds, and Read / Modify / Write operations**.
- Based on the RAID configuration, the controller might be able to handle the medium error reported by the HDD and no further action will be required.
- In some cases, the controller will not be able to handle the medium error and will pass the error to the host to handle the error.

When does the Operating System (OS) see medium errors?

- If the HDD reports a medium error and the RAID controller cannot handle the recovery, then the host will be notified of the error.
- This notification is no longer just an advisory message that would inform the system that the event has occurred, it is a request for the OS to act because the HDD and RAID controller was not able to recover from the medium error.
- If the OS has the required context to correctly resolve the medium error, it must be handled by the OS
- If disks are in Just a Bunch Of Disk (JBOD), the OS will see errors as they are not corrected by the controller. This is common in HyperFlex (HX)/ Virtual Storage Area Network (VSAN) environments.

HDD Role


Grown Defects (G-list) HDD Level


While a drive is in operation, the head might come across a sector with a weakened magnetic read level. The data is still readable but might fall below the preferred threshold for qualified good sector read levels. This disk drive would consider this a sector that could and would **sector spare** this data to a new location available in the **known good reserve** list. Once the data is moved, the old sector address is added to the **Grown Defects** list, never to be used again. This process is a **recoverable** media error. The drive will give a SMART trigger once a majority of its known-good spare sectors are exhausted.

RAID Controller Role

Patrol Read

- Patrol Read is a user-definable option that performs drive reads in the background and maps out any bad areas of the drive.
- Patrol Read checks for physical disk errors that could lead to drive failure. These checks usually include an attempt at corrective action. Patrol read can be enabled or disabled with automatic or manual activation.
- A Patrol Read periodically verifies all sectors of physical disks that are connected to a controller, that include the system reserved area in the RAID configured drives. Patrol Read works for all RAID levels and all hot spare drives.
- This process starts only when the RAID controller is idle for a defined length of time and no other background tasks are active, though it can continue to run at the same time as heavy Input/Output (I/O) processes.
- You cannot conduct patrol reads on drives configured in JBOD.

 **Note:**Latent Semantic Indexing (LSI) recommends that you leave the patrol read frequency and other patrol read settings at the default values to achieve the best system performance. If you decide to change the values, record the original default value here so you can restore them later.

 **Note:** Patrol Read does not report on its progress as it runs. The patrol read status is reported in the event log only.

Patrol Read options are as shown in the image:

Table 47: Set Patrol Read Options

Convention	MegaCli -AdpPR -Dsbl EnblAuto EnblMan Start Stop Info SSDPatrolReadEnbl SSDPatrolReadDsbl {-SetStartTime yyyyymmdd hh} maxConcurrentPD -aN -a0,1,2 -aALL
Description	<p>Sets Patrol Read options on a single controller, multiple controllers, or all controllers:</p> <ul style="list-style-type: none"> -Dsbl: Disables Patrol Read for the selected controller(s). -EnblAuto: Enables Patrol Read automatically for the selected controller(s). This means Patrol Read will start automatically after the controller initialization is complete. -EnblMan: Enables Patrol Read manually for the selected controller(s). This means that Patrol Read does not start automatically; it has to be started manually by selecting the Start command. -Start: Starts Patrol Read for the selected controller(s). -Stop: Stops Patrol Read for the selected controller(s). -Info: Displays the following Patrol Read information for the selected controller(s): <ul style="list-style-type: none"> • Patrol Read operation mode • Patrol Read execution delay value • Patrol Read status <p>SSDPatrolReadEnbl: Enable the patrol read operation (media scan) on a SSD. SSDPatrolReadDsbl: Disable the patrol read operation (media scan) on a SSD. SetStartTime yyyyymmdd hh: Set the start time for the patrol read in year/month/day format. maxConcurrentPD: Sets the maximum number of concurrent drives that patrol read runs on.</p>

MegaCli Examples

To see information about the patrol read state and the delay between patrol read runs:

```
# MegaCli64 -AdpPR -Info -aALL
```

To find out the current patrol read rate, execute:

```
# MegaCli64 -AdpGetProp PatrolReadRate -aALL
```

To disable automatic patrol read:

```
# MegaCli64 -AdpPR -Dsbl -aALL
```

To enable automatic patrol read:

```
#MegaCli64 -AdpPR -EnblAuto -aALL
```

To start a manual patrol read scan:

```
# MegaCli64 -AdpPR -Start -aALL
```

To stop a patrol read scan:

```
# MegaCli64 -AdpPR -Stop -aALL
```

Consistency Check

- In RAID, the Consistency Check verifies the correctness of redundant data in an array. For example, in a system with parity, to check consistency means to compute the parity of the data drives and compare the results to the contents of the parity drive.
- JBOD does not support consistency check.
- RAID 0 does not support consistency check.
- RAID 1 uses a data compare not parity.
- RAID 6 computes parity for 2 parity drives and verifies both.



Note: It is recommended you run a consistency check at least once a month.

Consistency Check management options are as shown in the image:

Table 74: Manage Consistency Check

Convention	MegaCli -LDCC -Start -Abort -ShowProg -ProgDsply -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
Description	Allows you to select the following actions for a data CC: -Start: Starts a CC on the virtual drive(s), then displays the progress (optional) and time remaining. -Aabort: Aborts an ongoing CC on the virtual drive(s). -ShowProg: Displays a snapshot of an ongoing CC. -ProgDsply: Displays ongoing CC progress. The progress displays until at least one CC is completed or a key is pressed.

Consistency Check scheduling options are as shown in the image:

Table 75: Schedule Consistency Check

Convention	MegaCli -AdpCcSched -Dsbl -Info {-ModeConc -ModeSeq [-ExcludeLD -LN -L0,1,2] [-SetStartTime yyyyymmdd hh] [-SetDelay val] } -aN -a0,1,2 -aALL
Description	Schedules check consistency on the virtual drive of the selected adapter. Dsbl: Disables a scheduled CC for the given adapter(s). Info: Gets information about a scheduled CC for the given adapter(s). ModeConc: The scheduled CC on all of the virtual drives runs concurrently for the given adapter(s). ModeSeq: The scheduled CC on all of the virtual drives runs sequentially for the given adapter(s) ExcludeLd: Specify the virtual drive numbers not included in the scheduled CC. The new list will overwrite the existing list stored on the controller. This is optional. StartTime: Sets the next start time. The date is in the format of yyyyymmdd in decimal digits and followed by a decimal number for the hour between 0 ~ 23 inclusively. This is optional. SetDelay: Sets the execution delay between executions for the given adapter(s). This is optional. Values: The value is the length of delay in hours. A value of 0 means continuous execution.

MegaCli Examples

To see the next scheduled Consistency Check time:
#MegaCli64 -AdpCcSched -Info -aALL

To change scheduled Consistency Check time:

```
#MegaCli64 -AdpCCSched -SetStartTime 20171028 02 -aALL
```

To disable Consistency Check:

```
#MegaCli64 -AdpCcSched -Dsbl -aALL
```

Conditions when a RAID Controller Cannot Repair a Medium Error

- In JBOD
 - The host OS is responsible for medium errors.
- In RAID 0
 - There is no redundancy, so the controller cannot provide the HDD with the data to write to the LBA.
- In RAID 1
 - When the controller cannot tell which mirror copy contains the correct data. This will only occur if both LBAs can be read, but the data does not match.
- RAID 5
 - If there are 2 or more errors in the same stripe. Most likely to occur when after a rebuild of an array is initiated. The drive that is rebuilt is one error, and a medium error on any other drive rebuild would be the second error. The controller would not be able to reconstruct the data needed to rebuild the LBA on the replacement drive.
- RAID 6
 - If there are 3 or more errors in the same stripe. Most likely to occur when while an array is rebuilt. The drive that is rebuilt is one error, and a medium error on any two other drives while the rebuild is in progress would be second and third errors, or a medium error and a second drive failure. The controller would not be able to reconstruct the data needed to rebuild the LBAs on the drives with the errors.

Related Information

- [MegaRaid® SAS Software User Guide](#)
- [Technical Support & Documentation - Cisco Systems](#)