Manage the RAID system from event log

Tim Chung

Version 1.0 (JAN, 2010)

QSAN Technology, Inc.
http://www.QsanTechnology.com
White Paper# **QWP201001-ALL**

Introduction

Event log records the information about the health of whole RAID system. User should always keep an eye on the event log and know well about the status of RAID system. This document will describe how to analyze the event log, furthermore, to understand what is wrong with the RAID system. At last, user can learn how to deal with the error to prevent it from being a disaster to the whole RAID system.

Contents

Part 1: Basic knowledge of event log

1. For more detail of event list, please refer to the “Event notifications” section in user manual.
2. The event log history is stored within the first four hard drives. At most there are four copies. Please make sure that at least one hard drive is installed at the first four disk slots.
3. The maximum size of event logs stored in hard drive is 16MB. When there is no one installed at the first four slots, event logs will be stored in memory within size 4MB and will be gone when system is shutdown. The logs will rotate automatically. When event logs are full, the newest one will overwrite the oldest.
4. When exporting event logs, all event levels (info, warning, and error) will be exported no matter what the event filter is configured.
5. Disk roaming (no matter online or offline) will clean all event history stored in the disks. Please make sure to save the event logs before performing disk roaming.
6. Event log can not be restored after cleaning.

Part 2: Case study

Case 1: RAID system shutdown abnormally.

Example of event log:

Section A. Log of normal shutdown/reboot
INFO: Wed, 23 Dec 2009 19:14:32 CST System reboot from 192.168.10.62 via Web UI
INFO: Wed, 23 Dec 2009 19:15:22 CST ECC memory is installed
INFO: Wed, 23 Dec 2009 19:15:29 CST Battery backup feature is disabled.
INFO: Wed, 23 Dec 2009 19:15:31 CST The global cache is ok.

Section B. Log of abnormal shutdown/reboot
INFO: Wed, 23 Dec 2009 19:16:28 CST admin login from 192.168.10.62 via Web UI
// There is no system shutdown/reboot exists before the “ECC memory is installed” message.
INFO: Wed, 23 Dec 2009 19:17:00 CST ECC memory is installed
INFO: Wed, 23 Dec 2009 19:17:01 CST Battery backup feature is disabled.
INFO: Wed, 23 Dec 2009 19:17:03 CST The global cache is ok.

Issue:

The “ECC / Non-ECC memory is installed” is the first event logged by the RAID

system after booting. If there is neither “System reboot” nor “System shutdown” events displayed before “ECC / Non-ECC memory is installed” event, it means that the RAID system is not reboot or shutdown properly with the standard procedure. The abnormal shutdown will let the dirty data in memory (system cache) lose because it has not been flushed into the disks. This will damage the data integrity.

Solution:

User should always shutdown or reboot the RAID system correctly by executing the shutdown or reboot command from LCM, Web UI or console UI. Complete shutdown procedure will prevent the data from damaged or corrupted.

Case 2: Dirty data in cache is flushed into the disks after abnormal shutdown.

Example of event log:

INFO: Thu, 01 Oct 2009 14:23:32 ECC memory is installed
INFO: Thu, 01 Oct 2009 14:23:42 Battery backup feature is enabled.
// The battery backup feature is enabled only when the BBM is installed already.
INFO: Thu, 01 Oct 2009 14:23:42 Battery backup module is detected
INFO: Thu, 01 Oct 2009 14:23:45 Battery backup module is good
INFO: Thu, 01 Oct 2009 14:23:58 Abnormal shutdown detected, start flushing battery-backed data (256 KB).
INFO: Thu, 01 Oct 2009 14:23:58 Abnormal shutdown detected, flushing battery-backed data finished

Issue:

When the RAID system is shutdown or rebooted abnormally or fails by power outage, the dirty data in memory (system cache) will be lost. It will damage the data integrity.

Solution:

During the power outage, the BBM (Battery backup module) can provide enough power to preserve the dirty data in memory for a period of time (72 hours with full battery). When next booting, the RAID system will flush the dirty data into the disks.

Case 3: ECC-bit error in memory.

Example of event log:

WARNING: Fri, 11 Sep 2009 19:17:53 Single-bit ECC error is detected at 0x373a49e8
ERROR: Fri, 11 Sep 2009 19:17:53 Multi-bit ECC error is detected at 0x373a4db0

Issue:

The ECC (Error-correcting code) is used to detect the bit error in RAM. You can refer to the wiki page “Errors and error correction” for reference. For “Single-bit ECC error”, it can be corrected by extra parity bit without breaking the data integrity. The RAID system is still working. For “Multi-bit ECC error”, it is a critical error that two or more

bits with errors are detected and can not be corrected. The “Multi-bit ECC error” may cause the data corrupted and makes the RAID system act abnormally.

Solution:

When user finds ECC error event, no matter Single-bit or Multi-bit errors, it’s better to replace a new RAM module within certification list. If ECC error still exists, the problem may come from both RAM module and controller board. Please send both of them back for RMA.

Case 4: Recoverable Read/Write Error.

Example of event log:

Section A. Recoverable read error
ERROR: Thu, 17 Sep 2009 11:13:05 HKT Read error occurred at LBA 0x3e51af80-0x3e51afff of PD 6.
// There is a read error block happened in PD6.
INFO: Thu, 17 Sep 2009 11:13:05 HKT Rewrite at LBA 0x0368779900 of UDV udv2 starts.
// The error block starts rewriting data into the same LBA; the error block belongs to UDV udv2.
ERROR: Thu, 17 Sep 2009 11:13:05 HKT Recoverable read error occurred at LBA 0x3e51af80-0x3e51afff of UDV udv2.
// A recoverable read error belongs to UDV udv2.
INFO: Thu, 17 Sep 2009 11:13:05 HKTRewrite at LBA 0x0368779900 of UDV udv2 completes.
// The recovered data is rewritten successfully.

Section B. Recoverable write error.
ERROR: Thu, 17 Sep 2009 14:33:18 HKTWrite error occurred at LBA 0x0c447c00-0x0c447c7f of PD 14.
// There is a write error block happened in PD14.
WARNING: Thu, 17 Sep 2009 14:33:18 HKTUDV udv2 is in degraded mode.
// This error block belongs to UDV udv2 and makes it become degraded.
ERROR: Thu, 17 Sep 2009 14:33:18 HKTRecoverable write error occurred at LBA 0x0c447c00-0x0c447c7f of UDV udv2.
// A recoverable write error belongs to UDV udv2.

Issue:

When a read error block happens on a redundant RAID group, such as RAID 3/5/6, the RAID system can recover the data by calculating the parity data stored in other member disks of this RAID group. Then the data will be rewritten into the same LBA. On the other hand, when a write error block happens on a redundant RAID group, if the member disks with write error blocks do not exceed the tolerated-disk number of this RIAD group, then it will be defined as a recoverable write error. Because the RAID system still can read the complete data by calculating the parity data stored in other member disks of this RAID group.

Solution:
If the recoverable read/write error happens occasionally in different disks of the RAID system, it won’t have immediate danger to the data integrity. User should pay attention on the disk health of the RAID system. If the recoverable read/write error happens a lot on the same disk or on the same logical volume, then user has to

replace a new disk to prevent data lost.

Case 5: Unrecoverable Read/Write Error.

Example of event log:

Section A. Unrecoverable read error.
ERROR: Fri, 25 Dec 2009 14:16:08 CST Read error occurred at LBA 0x01ed4380-0x01ed43ff of PD 6.
// There is a read error block happened in PD6.
ERROR: Fri, 25 Dec 2009 14:16:08 CST Unrecoverable read error occurred at LBA 0x01ed4380-0x01ed43ff of UDV 5.
// An unrecoverable read error belongs to UDV 5.
ERROR: Fri, 25 Dec 2009 14:16:08 CST UDV 5 is failed.
// This unrecoverable read error makes UDV 5 fail.

Section B. Unrecoverable write error.
ERROR: Wed, 18 Nov 2009 12:27:44 EST Write error occurred at LBA 0x2dcb2b80-0x2dcb2bff of PD 3.
// There is a write error block happened in PD 3.
ERROR: Wed, 18 Nov 2009 12:27:44 EST Unrecoverable write error occurred at LBA 0x2dcb2b80-0x2dcb2bff of UDV 5.
// An unrecoverable write error belongs to UDV 5.
ERROR: Wed, 18 Nov 2009 12:27:44 EST UDV 5 is failed.
// This unrecoverable write error makes UDV 5 fail.

Issue:

When a read/write error happens on a non-redundant RAID group (such as RAID 0) or on a redundant RAID group in degraded mode (such as RAID 3/5/6), the data can not be recovered because there is no parity data. In this case, the read/write error becomes unrecoverable and makes the VD/UDV with read/write error fail. But other VDs/UDVs in the same RAID group, they do no be affected by this unrecoverable read/write error, and still stay at the degraded mode as the RAID group does.

Solution:

When user meets the recoverable read/write error, the most important thing is to backup the data in the remaining good VDs/UDVs in the same RAID group if possible. Then user has to replace the faulty disk(s) to make the RAID group start rebuilding. User should keep in mind that it is important to backup valuable data always.

Case 6: Disk is disabled.

Example of event log:

ERROR: Fri, 25 Dec 2009 10:49:03 Disk 6 is disabled

Issue:

When a disk encounters too many unrecoverable read/write errors, this disk will be disabled by RAID system. The health status of this disk will be displayed as **Fail** in

web UI.

Solution:

User should always replace the failed disk to a new one as soon as possible and makes the RAID system start data rebuilding.

Case 7: Disk gets no response.

Example of event log:

ERROR:	Sun, 20 Sep 2009 15:06:14	Disk 14 gets no response

Issue:

When a disk gets no response, it means that the disk has no response to I/O commends from RAID system for a period of time. Because the RAID system still has to continue the process of data accessing from the host, the RAID system has to give up retrying the disk and set it as fail.

Solution:

Most of the times, the disk with no response can be online again after RAID system reboots or user hot-plug it again. At the first time of this situation, user can hot-plug it and continue to use it as a member disk of the original RAID group after rebuilding. But if the same disk has no response 2 times or more, it's better to replace it to a new disk.

Case 8: SCSI Bus reset.

Example of event log:

INFO:	Tue, 27 Oct 2009 21:23:34 MDT	Received SCSI Bus Reset event at the SCSI Bus 1

Issue:

The bus reset event happens seldom in SCSI model of the RAID system. A few bus reset events can be ignore when system booting, user removes or connects the SCSI cables. But if a lot of bus reset events appear, it means that the electrical signal on SCSI bus is not good. To many bus resets may cause I/O failure and performance drops.

Solution:

To eliminate the bus reset issue, user has to reconnect all the SCSI devices including SCSI HBA, SCSI cable and terminator. Make sure every device is connected firmly. If the bus reset events still happen seriously, user may have to replace the related SCSI devices for cross-test or run the system with lower speed.

Case 9: Temperature warning and overheated.

Example of event log:

WARNING: Mon, 28 Dec 2009 15:43:04 CST System temperature(Core Processor : 80) is above normal range
ERROR: Mon, 28 Dec 2009 15:45:42 CST System Overheated(Core Processor : 85)!!!
// The temperature is over the range.
ERROR: Mon, 28 Dec 2009 16:42:32 CST System Overheated(Core Processor : 85)!!! The system will auto-shutdown immediately
// The RAID system will be shutdown automatically if the feature “Auto shutdown” is enabled.

Issue:

There are two event levels about temperature. One is warning level, and the other is error level. When the temperature is over the range, the RAID system will issue a warning event of temperature. In addition, there is another set of critical temperature values defined in RAID system. If the temperature is over the range of critical value, the RAID system will issue an error event of temperature, and the RAID system will be malfunction or even be damaged if the temperature problem is getting worse.

Solution:

User should always run the RAID system in a cooling environment. Make sure that the air flow of FAN modules in RAID system is unhindered. The **QSAN** RAID system provides a function “Auto shutdown” which can monitor and shutdown the RAID system automatically when the RAID system reaches the critical temperature.

Case 10: Voltage warning and out of range.

Example of event log:

INFO: Mon, 28 Dec 2009 16:54:46 CST Onboard +1.2V Voltage = 1.17V
ERROR: Mon, 28 Dec 2009 17:13:33 CST System voltages(Onboard +1.2V Voltage = 1.17V) failed!!!
// The voltage is over the range.
ERROR: Mon, 28 Dec 2009 16:55:25 CST System voltages(Onboard +1.2V Voltage = 1.17V) failed!!! The system will auto-shutdown immediately
// The RAID system will be shutdown automatically if the feature of “Auto shutdown” is enabled..

Issue:

There are two different event levels about voltage. One is info level, and the other is error level. When the voltage is higher or lower the range which is defined in “Hardware monitor” page, the RAID system will issue a info event of voltage. In addition, there is another set of critical voltage settings defined in RAID system. These values are different with various chassis. If the voltage is higher or lower the range of critical value, the RAID system will issue an error event of voltage, and the RAID system will malfunction or even be damaged if the voltage problem is getting worse.

Solution:

Usually, the problem of abnormal voltage comes from the faulty PSU. It can be fixed by placement PSU. In some area, the power supply is unstable, it may cause the PSU bad earlier. So a UPS is needed for protection of surges, lightning, and other power disturbances. The **QSAN** RAID system provides a function “Auto shutdown” which can monitor and shutdown the RAID system automatically when the RAID system reaches the critical voltage.

Case 11: S.M.A.R.T threshold exceed condition / failure to get S.M.A.R.T information.

Example of event log:

WARNING: Fri, 25 Dec 2009 10:42:34 Disk 6: S.M.A.R.T. Threshold Exceed Condition occurred for attribute reallocated sector count
// The threshold of specific S.M.A.R.T attribute (which is “reallocated sector count“ here) on disk 6 is reached.
WARNING: Mon, 13 Jul 2009 14:57:30 GMT Disk 2: Failure to get S.M.A.R.T information
// The S.M.A.R.T information of disk 2 is not available now.

Issue:

The S.M.A.R.T. values are reported by hard drive itself. The RAID system reads the S.M.A.R.T. information only and displays those in “S.M.A.R.T.” page. The system will not use the data to judge whether the hard drive is healthy or not. When the threshold of an S.M.A.R.T. attribute is reached, the RAID system will report a warning message in event log. If the RAID system fails to read the S.M.A.R.T. value of a hard drive, it will also report a warning message in event log. For more information about S.M.A.R.T., please refer to the wiki page “S.M.A.R.T.”.

Solution:

Generally speaking, when the threshold of attributes of “Read error rate”, “Reallocated sector count”, and “Spin up retires” is reached, it means that the hard drive is unhealthy. But it doesn’t mean the hard drive will fail immediately. We suggest replacing a new one. If not, user has to take the risk of losing data.

Case 12: Unable to log on iSCSI target when session limit is reached.

Example of event log:

WARNING: Fri, 14 Nov 2008 22:41:33 CST iSCSI login from iqn.1986-03.com.sun:01:baeaf6effffff.4917db75 (192.168.1.77:33062) was rejected, reason [out of resources].

Issue:

This message means the maximum number of session is reached and the iSCSI target can not make any more connections. The maximum session number of iSCSI RAID system is different by various models. Please check the specification of each model.

Solution:

User can see how many sessions are connecting now in “Session” page. Please make sure whether the session number reaches the limit or not.

References

- Dynamic random access memory - Errors and error correction
  http://en.wikipedia.org/wiki/Dynamic_random_access_memory#Errors_and_error_correction

  S.M.A.R.T.
  http://en.wikipedia.org/wiki/S.M.A.R.T