



Maintenance Best Practices for Direct-Attached SCSI Solutions

Revision 1.1

July 20, 2005

(This page left intentionally blank)

1.0	Revision History.....	4
2.0	Introduction.	4
3.0	Overview of Steps Ensuring RAID Best Practices.	5
	CABLING PRACTICES	5
	MAINTENANCE OF ARRAYS.....	5
	RECOVERY OF ARRAYS.	5
4.0	Cabling Practices	6
	4.1 Only properly qualified cables should be used.....	6
	4.2 Cables should be properly secured to the PV22x and controller.....	6
	4.3 Cables should not be excessively bent.....	6
	4.4 Cables should not have exposed shielding or cuts in the outer jacket material	6
	4.5 Cable length should be appropriate for the installation.....	6
5.0	Maintenance of Arrays.....	7
	CONSISTENCY CHECKS FOR RAID ARRAYS	7
	5.1 Setup for automated scheduling of consistency checks on Windows systems.....	9
	5.2 Upgrade firmware, drivers and Array Manager concurrently to the latest versions.....	15
	5.3 Recondition the battery on the RAID/ROMB Controllers (if present).....	15
	5.3.1 Reconditioning the Battery using Array Manager.....	16
	5.4 Monitor System Event Logs and Array Manager Event Logs.	17
	5.5 Backup and Recovery of Data.....	18
6.0	Recovery of Arrays.	18
	6.1 Capture system logs and details surrounding array failures to assist in the recovery.	18
	6.2 Write down the exact steps or circumstances that caused the system to get in the failed state.	19
	Appendix.....	20
	PERC 3\SC\DC\QC AND PERC 4\SC\DC\DI RAID REBUILD BEHAVIOR	20
	1.1 Pre-May 2003 behavior	20
	1.2 Post May 2003 behavior	22
	1.3 Array Manager Online help files	24
	1.3.1 Rebuild Completes with Errors on PERC Subsystem 1 Controller.	24
	PV220s/PV221s MODE SETTING VERIFICATION	25
	2.0 Joined Bus Mode	25
	2.1 Split Bus Mode	25
	2.2 Cluster Bus Mode	26
	CUSTOMER USING LINUX DELL MGR UTILITY.....	27
	2.2.1 Using DellMon.....	27
	2.2.2 Customer using Ctrl+M utility.....	29

1.0 Revision History.

Date	Revision	Explanation of Changes
02/25/04	0.1	Initial draft
04/19/04	0.2	Adding Core Team Input.
05/11/04	0.3	Format rewrite
05/20/04	0.4	Format Changes, adding cabling best practice, and bus mode comments
05/25/04	0.5	Adding Core Team Input.
05/25/04	1.0	Official Release.

2.0 Introduction.

Objective and Scope: This document is intended to provide insight into the best practices for routine maintenance of direct attached SCSI disk systems. This document is not intended for addressing or recommending the type or size of arrays for specific applications the user may encounter. These maintenance best practices are recommended to all Dell Enterprise customers to avoid failures, downtime, and data loss. The use of these practices will ensure a better customer experience through maintaining the integrity of data and minimizing the costs of downtime.

The practices which will be discussed are as follows:

- Cabling Practices
 - Maintenance of Arrays
 - Maintenance of software/firmware stack
 - Reconditioning of RAID controller battery
 - Recovery of Arrays
-

3.0 Overview of Steps Ensuring RAID Best Practices.

Cabling Practices

- Assure that properly qualified cables are being used
- Assure that SCSI cables are properly secured to the PV22x and controller.
- Verify cables are not excessively bent
- Inspect cabling for cuts, exposed shielding
- Cable length should be appropriate for installation

Maintenance of Arrays.

- Run regular consistency checks on the system.
- Upgrade firmware, drivers and *Array Manager* in sync with releases.
- Recondition the battery on the RAID/ROMB Controllers (if present).
- Monitor System Event Logs and *Array Manager* Event Logs
- Establish Best Practices for Backup and recovery of data

Recovery of Arrays.

- Capture system logs and details surrounding array failures to assist in the recovery.
- Write down the exact steps or circumstances that caused the system to get in the failed state

4.0 Cabling Practices

4.1 Only properly qualified cables should be used

The SCSI standard has been through several iterations with speed increases at each change in the standard. The result is that there are several different SCSI cables that may look identical but have differing capabilities. Part numbers should be carefully verified to ascertain whether each cable is capable of operating properly at the transmission speed of the system.

4.2 Cables should be properly secured to the PV22x and controller.

High density connectors make it more difficult to see the proper alignment and fit of the VHDCC connectors' attachment point at the RAID controller. Check all connections and press them firmly during seating and securing to ensure that they are sitting squarely on the connector. If the connector doesn't seat easily and squarely, or requires too much effort to get in place, it may be that there are damaged connector pins. Inspect the connectors for bent or damaged pins and replace the cables as needed to correct the defect.

4.3 Cables should not be excessively bent.

Cables should have a smooth bend radius at the connector and should not kink. Some cable management techniques encourage a more aggressive bend than is appropriate for these cables. Also, keep in mind that the torque applied to the cable is translated onto the RAID card and its seating in the PCI slot on the server. If sufficient torque is applied to the RAID card by the cable then partial or complete unseating of the RAID controller may occur.

4.4 Cables should not have exposed shielding or cuts in the outer jacket material

Cables with exposed shielding or damaged jacket material should be replaced.

4.5 Cable length should be appropriate for the installation.

Excessively long cables may unnecessarily attenuate SCSI signals. Choose cable lengths that reach properly without stressing the connectors but that are not longer than needed.

5.0 Maintenance of Arrays.

RAID arrays are an industry standard for the protection of important data through redundancy. This redundancy may take the form of parity calculations that are dispersed throughout the array. It could also be the simple mirroring of data to maintain a complete copy that does not require parity calculations to reconstruct the missing elements.

Hard drive media defects and other drive quality issues have steadily improved over time, even as drive sizes have grown substantially. However, hard drives are not expected to be totally free of flaws. In addition, normal wear on a drive may result in an increase in media defects, or "grown defects," over time. The data block containing the defect becomes unusable and must be "remapped" to another location on the drive.

If a bad block is encountered during a normal write operation, the controller marks that block as bad and the block is added to the "grown defects list" in the drive's NVRAM. That write operation is not complete until the data is properly written in a remapped location. When a bad block is encountered during a normal read operation, the controller will reconstruct the missing data from parity operations and remap the data to the new location.

A condition known as a double fault occurs when a RAID controller encounters a bad block on a drive in a RAID volume and then encounters an additional bad block on another hard drive in the same data stripe. This double fault scenario can also occur while rebuilding a degraded logical drive, leaving the controller with insufficient parity information to reconstruct the data stripe. The end result is a rebuild failure with the loss of any data in that stripe, assuming the stripe is in the user data area.

Background Patrol Read

The June/July 2005 firmware and driver updates to PERC 3, PERC 4, and PERC 4e controllers (not including PERC 3/Di or PERC 4/im) introduce a new feature called Background Patrol Read. This is a background process that does a read verify across the physical disk in order to find bad blocks. Background Patrol Read is a media level feature and does not verify data inside the block or stripe. Background Patrol Read is the most efficient feature to address media errors.

Run Modes

Because this is a background process, Background Patrol Read is secondary to data I/O and it efficiently throttles to minimize performance impact

1. Auto

For the majority of the time, this is the recommended mode to use. Auto mode allows continual checking so that media errors can be detected and data remapped throughout the product's life.

2. Manual

Manual mode is useful for completing a single scan more quickly than using Auto mode. It accomplishes this by using different throttling parameters. This "run-once" mode is useful after making hardware changes to get the hard drives up to optimal condition more quickly.

NOTE:(Your management solution should be updated at the same time as the controller firmware and driver)

Consistency Check

Consistency check is a data level check that both verifies data inside the block or stripe, as well as checks for bad blocks. Consistency check will perform read operations on both the user data areas of the logical drive as well as the currently unused areas not containing any user data. This check finds and repairs stripes where data and parity are not matched. This allows the array to be accessed more efficiently because the RAID controller will not have to spend extra time checking when writing data. In addition, like Background Patrol Read, any bad blocks that are encountered will be entered in the grown defects list so they are not used in the future. The data in the damaged block is reconstructed from parity operations and written to a new location on the hard drive. The data stripe is completely recovered and now in an optimal state.

To avoid downtime and ensure data integrity, consistency checks should be part of the routine maintenance of any RAID system. This is especially important for RAID systems that do not have the Background Patrol Read feature. For these systems, it is recommended to schedule consistency checks at least once a month. To help maintain a regular interval, consistency checks can be scheduled, as indicated in the instructions in the next section.


For systems that do have the Background Patrol Read feature, Consistency Checks should still be run at a regular interval. The specific frequency is dependent on such factors as the size of the array, the type of data (mission critical, production, near line, archive, test, etc.), and other factors associated with the operations of your IT infrastructure, in order to minimize impact and maintain protection.

The amount of data in a given disk group has grown dramatically over the last few years as data storage needs have mushroomed in the business world. Additionally, Service Level Agreements continue to become more aggressive and data access is more mission critical in almost every business environment. These factors combine to create tension between business needs and technical realities. Backup and restore windows have shrunk almost inversely to the growth of data stores, creating a scenario whereby data recovery has become less timely due to data store sizes, while the need for the availability of that data has become more business-critical. Consideration should be given to the maximum volume size which can be useful in a given data storage situation. This strategy will aid in avoiding the challenges of mounting, handling, and maintaining very large file systems coupled with very long restore windows in the event of data corruption or loss. Purposefully limiting data store sizes also reduces the numerical probability of encountering additional bad blocks during rebuilds. A similar benefit accrues to limiting data store sizes by reducing the amount of time and system load required to perform consistency checking on a given logical drive.

Additional consideration may be given to performing intermediary backups to disk, which are then shuttled to tape as a background operation. This will decrease the backup window while also making the most recent full backup available in a faster media than going back to tape to perform a restore.

5.1 Setup for automated scheduling of consistency checks on Windows systems.


1. For systems with a Windows OS system and *Array Manager* installed, you can use the '**Scheduled Tasks**' option from the menu under the '**Accessories**' folder. Double-click on '**Add Scheduled Task**' and the following wizard will appear:



- Click '**Next**' and the following screen will appear:



- Click '**Browse**' and locate the file 'amcli.exe'. The AMCLI executable is located in the Array Manager installation directory.
- Select the file and click '**OK**'.
- Click '**Next**' and the following screen will appear:




- Enter a name for this task and select how often the task should be run. The minimum recommendation for this task is to be run at least once a month.

7. Click '**Next**' and the following screen will appear:



The screenshot shows the 'Scheduled Task Wizard' dialog box. On the left is a graphic of a clock and a pen. The main text says 'Select the time and day you want this task to start.' Below this, 'Start time:' is set to '8:47 PM'. There are two radio buttons: 'Day' (selected) with a spinner set to '1', and 'The' with dropdowns for 'first' and 'Monday'. Below that, 'Of the month(s):' has a grid of checkboxes for all months (Jan, Apr, Jul, Oct, Feb, May, Aug, Nov, Mar, Jun, Sep, Dec), all of which are checked. At the bottom are buttons for '< Back', 'Next >', and 'Cancel'.

8. Select the time at which the Consistency Check should run. Remember that there will be a system performance impact so you want to run this at a low traffic time.
9. Click '**Next**' and the following screen will appear:



The screenshot shows the 'Scheduled Task Wizard' dialog box, Step 2. On the left is the same clock and pen graphic. The main text says 'Enter the name and password of a user. The task will run as if it were started by that user.' Below this are three input fields: 'Enter the user name:' with 'AMERICAS\John_Doe' entered, 'Enter the password:', and 'Confirm password:'. A note below the fields states 'If a password is not entered, scheduled tasks might not run.' At the bottom are buttons for '< Back', 'Next >', and 'Cancel'.


10. Fill in the name and password fields appropriately so the task can be executed correctly.

11. Click '**Next**' and the following screen will appear:



12. Select the checkbox for '*Open advanced properties for this task when I click Finish.*'

13. Click '**Finish**' and the following screen will appear:



14. In the '**Run**' textbox you can type different parameters. The following is example syntax for scheduling a check consistency on virtual disk 1. '**"C:\PathName\amcli.exe" /c1**' where *PathName* is the path to the AMCLI executable. This will be the command executed by the scheduler every time it runs this event.
15. To run a consistency check the parameter is **amcli /cn** where the **c** option indicates "consistency check" and *n* is the number of a virtual disk as displayed in the Array Manager tree view.

To view the other AMCLI command options, display the AMCLI help information by entering ` **amcli** /? ` or look at the Array Manager help file.

Sample of command line options:

amcli /da where the d option indicates display and the a option indicates adapter.	Displays the status of the system controllers (adapters). The possible status values are: • None -- The controller does not have a battery. • Ready -- The battery is functioning normally and can be recharged. • Charging -- The battery is recharging. • Auto -- The battery is a self-charging battery. A battery recharge does not need to be initiated for this type of battery. Degraded -- The battery charge is too low. Initiate a recharge immediately. • Low -- The battery charge is low. Initiate a recharge soon. • Failed -- The battery needs to be replaced.
--	---

5.2 Upgrade firmware, drivers and Array Manager concurrently to the latest versions.

Enclosure firmware, RAID controller firmware, the RAID controller driver, and Array Manager are maintained in block releases. These block releases should be observed and installed/upgraded to receive maximum performance, dependability and functionality from your direct-attached SCSI solution. All the latest versions of code are available through support.dell.com.

Upgrading all the above at the same time will ensure correct functionality at all levels and that all features are available to the user.

5.3 Recondition the battery on the RAID/ROMB Controllers (if present).

RAID Controller batteries preserve the read and write cache contained with the nonvolatile memory (NVRAM) in the event of power loss. The period of time (holdover time) for which the battery is capable of maintaining the integrity of the cached information is dependent upon the charge level of the battery. A RAID controller battery should be considered to be fully discharged when installed into a system. After power is applied the controller will automatically initiate a full charge cycle of the battery.

Dell provides two distinct types of batteries on its PERC controllers:

- Lithium-Ion (Li-Ion) batteries
- Nickel-Metal Hydride (NiMH) batteries

Li-Ion batteries do not require reconditioning.


As is the case with any NiMH battery, a battery reconditioning (full discharge and recharge) is required to maintain the reliability of the battery. On Dell RAID Controllers equipped with a NiMH battery, this process should occur every six months, when the RAID controller reports that it needs to be reconditioned. To see what type battery your RAID controller has, please refer to the user guide for that product.

While the battery is being reconditioned (approximately 8 - 10 hours) the battery's holdover time is reduced to zero (during discharge) and then restored (during full charge). During this period, write-cache is disabled resulting in performance degradation.

The expected life span of any RAID battery is three years. You should replace the battery of your controller every three years to guarantee the quality and capacity of the RAID\ROMB battery.

5.3.1 Reconditioning the Battery using Array Manager.

Using *Array Manager* to recondition the battery simply requires right-clicking on the controller icon and selecting the "**Properties**" option. Once the "**Controller Properties**" dialog box (Figure 6) is displayed, click the "**Recondition**" button.




Once the reconditioning has started, the status of the battery displayed in *Array Manager* may not change. To get an accurate status reading, close *Array Manager* and then re-open it or perform a controller rescan from within *Array Manager*.


5.4 Monitor System Event Logs and Array Manager Event Logs.

System Events are generated for informational purposes, such as for record-keeping activities, or for notifying the user of events that may affect the physical security and availability of their data. If you are using a Windows-based OS and you have *Array Manager* installed, a comprehensive list of event types is provided in the help file of the application. You can view specific information on each event listed on the Events tab by double-clicking the event to display the Event Details dialog box as well. Regularly check the System Event Logs to look for any warnings or errors.

Regularly check the Array Manager Event Logs for corrected media errors. Corrected media errors are normal but an excessive number within a short time period may indicate that a drive should be proactively replaced during a maintenance cycle.



You will also be able to get the events from any Netware server with the Windows console. These events will be shown on the *Array Manager* Event Log.



The use of *Array Manager* is recommended for Windows or Netware based OS.

5.5 Backup and Recovery of Data.

The implementation of comprehensive backup and recovery strategy is recommended to guarantee the preservation of your data. The recovery plan should be tested at intervals to ensure its efficacy for your environment. During the backup window a reduction in normal system performance might be seen due to the increased workload.

6.0 Recovery of Arrays.

To avoid loss of data integrity or the recovery of lost arrays the following simple steps should be followed:

6.1 Capture system logs and details surrounding array failures to assist in the recovery.

In the Windows OS environment, the use of DSET is recommended. This tool will capture all the system description and configuration data needed in a debug or recovery effort. DSET is a small, non-intrusive tool that does not require a reboot of the system to provide basic functionality. Immediately after installation, DSET can collect information about Windows drivers, services, network settings, etc. It will also collect basic information about the system's storage such as active drives and RAID containers. DSET will also collect extended hardware information such as processors, memory, PCI cards, ESM log, BIOS/firmware versions and system health (fan/voltage levels). If Array Manager 2.5 or later is installed, DSET will gather Dell-specific storage information such as PERC controllers and their firmware, array/containers, logical disk signatures, enclosures and physical hard drives installed. DSET always collects Windows NT/2000 information such as services, tasks, drivers, and events logs. The ESM Log of supported systems can be cleared so that the amber hardware warning light on Dell PowerEdge systems can be properly reset once an event has been remedied.

You should alternatively use Array Manager for Netware-based systems to retrieve issue information. All the latest Array Manager versions for Windows and Netware offer the ability to retrieve NVRAM information for all the RAID controllers except the Perc 2\SC\DC. This NVRAM dump will contain all the configuration information as well as the contents of the NVRAM logs. To obtain this dump from Array Manager right click on the controller selected on the left panel, and select the appropriate option.

For the Linux operating system, the "Getconfig" tool can be used to retrieve issue information when possible. Getconfig is a utility that pulls the hardware configuration data from various sources on a Linux box. Retrieving the system data from a Linux box is very labor intensive and this tool fully automates the process.

6.2 Write down the exact steps or circumstances that caused the system to get in the failed state.

If possible, document all steps that occurred prior to the failure, including any changes made in the operating system, system software or hardware. The ability to backtrack and understand the activities that may have contributed to a failure will help in the attempt to recover a failed array and could help correct the conditions which helped contribute to the failure.



Appendix.

To avoid loss of data integrity or the recovery of lost arrays the following steps should be followed:

PERC 3\SC\DC\QC and PERC 4\SC\DC\Di RAID rebuild behavior

1.1 Pre-May 2003 behavior


When a media error is found during a rebuild, the PERC RAID Firmware will not complete the rebuild process. The new HDD is marked failed, and the Virtual Disk is kept in degraded state. The behavior in this scenario between a RAID 1 and a RAID 5 are different.

- RAID 1 Customer Experience: When a media error is found during a rebuild, the PERC RAID Firmware will complete the rebuild process. Errors are generated for each of the media errors (03 011 errors). The Virtual Disk is marked ready. The customer is left a ready Virtual Disk but does not know that media error were found and that some of the data may be compromised.
- RAID 5 Customer Experience: When a media error is found during a rebuild, the PERC RAID Firmware will not complete the rebuild process. The new HDD is marked failed, and the Virtual Disk is kept in degraded state. The customer is left with a degraded Virtual Disk that can not be corrected until the data is backed up, all HDDs changed, Virtual Disk recreated and the data is restored.


Array Manager log will look like this:

```
10 Information 1/15/2004 1:14:50 PM PercPro 554 PERC 4/DI Controller 0, Array
Disk 0:4 rebuild completed.
9 Warning 1/15/2004 12:59:49 PM PercPro 571 PERC 4/DI Controller 0, Array Disk
0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = 0. If this disk is
part of a redundant virtual disk, select the 'Offline' option and then replace
the disk. Then configure a hot spare and it will start the rebuild
automatically. If this disk is a hot spare, select the 'Prepare to Remove'
option and then replace the disk. If this disk is part of a non-redundant
disk, you should back-up your data immediately. If the disk fails, you will
not be able to recover the data.
8 Warning 1/15/2004 12:59:49 PM PercPro 571 PERC 4/DI Controller 0, Array Disk
0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = 0. If this disk is
part of a redundant virtual disk, select the 'Offline' option and then replace
the disk. Then configure a hot spare and it will start the rebuild
automatically. If this disk is a hot spare, select the 'Prepare to Remove'
option and then replace the disk. If this disk is part of a non-redundant
disk, you should back-up your data immediately. If the disk fails, you will
not be able to recover the data.
7 Information 1/15/2004 12:55:20 PM PercPro 527 PERC 4/DI Controller 0, Array
Disk 0:4 rebuild started.
```

When the rebuild process find the media error, Array Manager Pop-up the message "rebuild failed" as follows:



Screenshot for Array Manager showing multiple media errors (03 11) with a "rebuild completed" message:



1.2 Post May 2003 behavior

For the May release, the PERC Firmware, and Array Manager has changes introduced to improve the customer experience in the RAID rebuild scenario.

RAID 1 and RAID 5 Customer Experience:

- The new RAID FW handles media errors and completes the rebuild process, creating an *Array Manager* event for each of the media errors and completing the rebuild with a "completed with error" message.
- The *Array Manager* documentation highlights the next steps to the customers when the message is found.

Array Manager log will look like this:

```
102 Error 1/14/2004 10:00:33 AM PercPro 690 PERC 4/DI Controller 1, Array Disk
0:5 rebuild completed with error(s).
101 Warning 1/14/2004 9:48:23 AM PercPro 671 PERC 4/DI Controller 1, Array Disk
0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = 0. If this disk is part of a
non-redundant virtual disk, the data for this block cannot be recovered. The disk
will require replacement and data restore. If this disk is part of a redundant
virtual disk, the data in this block will be reallocated.
12 Information 1/14/2004 9:32:13 AM PercPro 527 PERC 4/DI Controller 1, Array
Disk 0:5 rebuild started.
11 Warning 1/14/2004 9:32:02 AM PercPro 509 PERC 4/DI Controller 1, Virtual Disk
(Virtual Disk 1) degraded.
10 Error 1/14/2004 9:32:02 AM PercPro 500 PERC 4/DI Controller 1, Array Disk 0:5
failed.
9 Warning 1/14/2004 9:32:02 AM PercPro 682 PERC 4/DI Controller 1, Array Disk
0:5 forced offline.
8 Information 1/14/2004 9:31:33 AM PercPro 550 PERC 4/DI Controller 1, Virtual
Disk (Virtual Disk 1) initialization completed.
7 Information 1/14/2004 9:31:30 AM PercPro 523 PERC 4/DI Controller 1, Virtual
Disk (Virtual Disk 1) initialization started.
6 Information 1/14/2004 9:31:29 AM PercPro 505 PERC 4/DI Controller 1, Virtual
Disk (Virtual Disk 1) created.
5 Warning 1/14/2004 9:30:49 AM PercPro 506 PERC 4/DI Controller 1, Virtual Disk
(Virtual Disk 1) deleted.
4 Information 1/14/2004 9:26:48 AM PercPro 550 PERC 4/DI Controller 1, Virtual
Disk (Virtual Disk 1) initialization completed.
3 Information 1/14/2004 9:26:45 AM PercPro 523 PERC 4/DI Controller 1, Virtual
Disk (Virtual Disk 1) initialization started.
2 Information 1/14/2004 9:26:43 AM PercPro 505 PERC 4/DI Controller 1, Virtual
Disk (Virtual Disk 1) created.
```


Screenshot for *Array Manager*:

The screenshot displays the Dell OpenManage Array Manager interface. The left sidebar shows a tree view with 'JIMG_ALPHA' selected, containing 'Arrays', 'PERC Sub...', 'PERC...', 'Ar...', 'Ba...', 'Ch...', 'Disks', 'Volumes', 'My Network Places', 'Favorites', and 'History'. The main pane is titled 'Events' and shows a table of alerts. The table has columns: Alert, Date, Time, Source, Event, and Description. The 'Alert' column shows a yellow warning icon. The 'Source' column is 'PercPro'. The 'Event' column is '571'. The 'Description' column contains detailed error messages for PERC 4/DI Controller 0, Array Disk 0:3, including Sense Key, Sense Code, and Sense Qualifier values.

Alert	Date	Time	Source	Event	Description
Information	1/12/2004	10:04:24...	PercPro	690	PERC 4/DI Controller 0, Array Disk 0:4 rebuild completed with error(s).
Warning	1/12/2004	9:51:54 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:51:54 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:51:54 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:51:44 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:51:44 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:51:34 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:51:34 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:51:24 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:51:24 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:51:24 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:51:14 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:51:14 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:51:04 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:51:04 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:50:54 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:50:54 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:50:54 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:50:44 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:50:44 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:50:34 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:50:34 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:50:24 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:50:24 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:50:24 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:50:14 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:50:14 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:50:14 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:50:04 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:50:04 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:49:54 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:49:54 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:49:54 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:49:44 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:49:44 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1
Warning	1/12/2004	9:49:34 PM	PercPro	671	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 3, Sense Code = 11, Sense Qualifier = C
Warning	1/12/2004	9:49:34 PM	PercPro	571	PERC 4/DI Controller 0, Array Disk 0:3 Sense Key = 2, Sense Code = 4, Sense Qualifier = 1

1.3 Array Manager Online help files

The online help for *Array Manager* was modified to include more information and explanation for the above message "Completed with Errors". The following is the content added to the last version of *Array Manager* (build 555):

1.3.1 Rebuild Completes with Errors on PERC Subsystem 1 Controller.

1.3.1.1 This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, PERC 4/IM, and CERC ATA100/4ch controllers.

In some situations, a rebuild may complete successfully while also reporting errors. This may occur when a portion of the disk containing redundant (parity) information is damaged. The rebuild process can restore data from the healthy portions of the disk but not from the damaged portion.

When a rebuild is able to restore all data except data from damaged portions of the disk, it will indicate successful completion while also generating event 690 <eventm14.htm>. The rebuild may also report sense key errors. In this situation, take the following actions to restore the maximum data possible:

Back up the degraded virtual disk onto a fresh (unused) tape.

- **If the backup is successful.** If the backup completes successfully then the user data on the virtual disk has not been damaged. In this case, you can continue with step 2 <troub128.htm>.
- **If the backup encounters errors.** If the backup encounters errors then the user data has been damaged and cannot be recovered from the virtual disk. In this case, the only possibility for recovery is to restore from a previous backup of the virtual disk.

Perform a Check Consistency on the virtual disk that you have backed up onto tape. Restore the virtual disk from the tape onto healthy array disks.

1.3.1.2 Receive Event 691 with "BBR Error" or "Error with Sense," or "Medium Error" Message

Event 691 is generated when a portion of a disk is damaged. This damage is discovered when the controller performs an operation that requires scanning the disk. Examples of operations that may result in event 691 are as follows:

- Consistency check
- Rebuild
- Virtual disk format
- I/O

If you receive event 691 as the result of doing a rebuild or while the virtual disk is in a degraded state, then data cannot be recovered from the damaged disk without restoring from backup. If you receive event 691 under circumstances other than a rebuild, then data recovery may be possible. The following describes each of these situations.

1.3.1.3 Event 691 Received during a Rebuild or while a Virtual Disk is Degraded

Do the following if you receive event 691 during a rebuild or while the virtual disk is in a degraded state:

1. Replace the damaged drive.
2. Create a new virtual disk and allow the virtual disk to completely resynchronize. To monitor the progress of the resynchronization, select the parent of the virtual disk in the Array Manager tree view and look at the status in the right pane. While the resynchronization is in progress, the status of the virtual disk will be Resynching.
3. Restore data to the virtual disk from backup.

1.3.1.4 Event 691 Received while Performing I/O, Consistency Check, Format, or Other Operation

If you receive event 691 while performing an operation other than a rebuild, you should replace the damaged disk immediately to avoid data loss.

Do the following:

1. Back up the degraded virtual disk to a fresh (unused) tape.
2. Replace the damaged disk.
3. Do a rebuild.

PV220s/PV221s Mode Setting Verification

After changing the operational mode setting of a PV220 it is always necessary to power cycle the unit to apply the new setting. Failure to do so will result in abnormal and unpredictable behavior.

2.0 Joined Bus Mode

1. Check that the mode setting switch is in the topmost position (joined bus mode).
2. Only one SCSI cable will be connected to the PV220s/PV221s, attaching to the master EMM module (with two green lights) with no cable connection going to the other EMM module.
3. After changing the switch to joined bus mode, power down the master server, and reboot the PV220s/PV221s.
4. After the PV220s/PV221s has been rebooted, power up the master server.
5. Failure to power cycle in this instance will likely show symptoms including drives not found, bus hangs and equipment not identified properly.
6. Check the front of the unit and verify that the Split bus mode and Cluster mode LEDs are not lit.
7. During server post, enter the PERC BIOS (ctrl-m) utility to verify the controller bus ID is 7.

2.1 Split Bus Mode

1. Check that the mode setting switch is in the middle position (split bus mode).
2. Two SCSI cables will be attached to the PV220s/PV221s, one attaching to each EMM module.


3. After changing the switch to Split Bus Mode, power down both servers and reboot the PV220s/PV221s.
4. After the PV220s/PV221s has been rebooted, power up each of the attached servers.
5. Failure to power cycle in this instance will likely show symptoms such as no post or hang on the HBA/PERC controller initialization.
6. Check the front of the unit and verify that the Joined Mode and Cluster Mode LEDs are not lit, and that the Split Mode LED is lit.
7. During the post for each server that is attached (maximum 2), enter the PERC BIOS (ctrl-m) utility and verify that the controller ID is set for 7 on each of the controllers accessing that half of the bus.

2.2 Cluster Bus Mode

1. Check that the mode setting switch is in the bottom-most position (Cluster Bus Mode).
2. Two SCSI cables will be attached – one from each of the two cluster servers. It is not permitted to operate both buses from a single controller in one server, or two controllers in one server.
3. After changing the switch to Cluster Bus Mode, power down both servers and reboot the PV220s/PV221s.
4. After the PV220s/PV221s has been rebooted, power up each of the attached servers.
5. Failure to power cycle in this instance will likely show symptoms such as no post or hang on the HBA/PERC controller initialization.
6. Check the front of the unit and verify that the Joined Mode and Split Mode LEDs are not lit, and that the Cluster Mode LED is lit.
7. During the post for the **first** of the two cluster node servers that are attached, enter the PERC BIOS (ctrl-m) utility and enable clustering on the controller and set its ID to 6. On the second cluster node server, enter the utility and enable clustering, and verify that its ID is still set to 7, which is the default. If it is not set to 7 it should be done at this time. The backplane will then use ID 6 so that ID 15 will no longer be available and any drive in slot 15 will not spin up.

Customer using Linux DellMgr utility.

The administrator will see the percentage increase from 0% to 100% and change to "INCON" as follows:



As in the case of Ctrl-M application, DellMgr does not message when media errors are found during the rebuild. The administrator can run DellMon in the Linux system to obtain further events and notifications related to the RAID subsystem.

Note: The administrator can also open a new Linux session to run other operations and then go back to the DellMgr screen to check on the end result of the rebuild.

2.2.1 Using DellMon.

DellMon is a Linux monitoring utility that Dell post in the Web site. The utility is not factory installed but the customer can download and install from the Dell external web site. As in the case of DellMgr, Dell does a best effort support.

DellMon is responsible to logging RAID events in a Linux environment. The events are added to a flat file. The following is a log file from DellMon captured in the same test case

mentioned above. The file shows the media errors (01 33) between the "rebuild started" and "rebuild completed" message.

Sample Event logs generated by DellMon:

```
[01/13/2004 (11:53:23)]:
  Adapter 1: No of Charge Cycles = 4
[01/13/2004 (11:54:14)]:
  Adapter 1: Battery Voltage GOOD.
[01/13/2004 (22:57:53)]:
  Adapter 1: Battery Fast Charging FAILED.
[01/14/2004 (00:58:51)]:
  Adapter 1: The Number of Logical Drive(s) has Changed to 2.
[01/14/2004 (01:06:50)]:
  PERCRAID Monitor Service Ver 4.00(40-LD) Aug 05, 2003 started
[01/14/2004 (01:06:50)]:
  MailFlag value set as:-
  MailFlag: 1 (0-means no mailing of events will be done
              1-means events will be mailed also).

[01/14/2004 (01:06:50)]:
  -: Initial CheckConsistency Schedule :-
  EnableFlag: 1
  Date: 01/14/2004
  DayOfWeek: 0
  HourOfDay: 0
  Week(s): 1
  ReportChkonProgInterval: 0 seconds(0 means report in-frequently)

[01/14/2004 (01:06:50)]:
  -: Initial SMART Polling Schedule :-
  SMARTPollingOn: Yes
  SMARTPollingInterval: 60 Seconds

[01/14/2004 (01:06:50)]:
  At StartUp Adapter 0 has 1 drive(s) as
  (Shown as Chnl:Target: Vendor Information)
  Ch-0:ID-1: MAXTOR ATLASU320_18_SCAB120
```

2.2.2 Customer using Ctrl+M utility.

The PERC 3\SC\DC\QC and the PERC 4\SC\DC\Di RAID controllers can also be managed with the Ctrl-M utility. This utility is accessed during the booting process by pressing Ctrl-M.

Administrator can monitor the rebuild process from Ctrl-M screen. The screen will show the percentage completed during the rebuild process until a 100% is displayed. If a media error is found during the rebuild, the Ctrl-M utility will not communicate that to the customer, but it will display a "INCON" message upon completed as shown here

