

STORAGE MANAGER UTILITY

Note: Storage Manager releases prior to rel. 2.0 are specific and dedicated to a determined type of DPT controller (EISA or PCI); from release 2.0 onwards Storage Manager can be used by both boards. The evolutions of the different releases are explained with the respective controllers in Appendix C.

Note: Information on how to use Storage Manager is provided, in a readme file on Disk 1, in the CD-ROM kit and in the related online manual provided on the CD-ROM. Read this file using a text editor.

This utility is contained on two diskettes or on the Orchestra Systema CD-ROM and is only available for systems equipped with the RAID DPT SCSI controller. Storage Manager performs several functions:

- Checks the initial hardware configuration of the SCSI controller after the addition of SCSI devices.
- Creates, if necessary, a specific Supplemental Utility diskette for the operating system.
- Allows disk arrays to be configured.
- Runs diagnostics on the subsystem.

F

Storage Manager runs under DOS or as a Client-Server application under Novell Netware 386 3.1x and 4.x. It is a graphical utility that automatically switches to text mode when a monochrome or CGA monitor is used. On the top of the screen are menus that can be selected using either the mouse or keyboard. Below the menu bar is a tool bar containing the function buttons, and below the tool bar is the main screen which contains the device icons.

A mouse simplifies the use of Storage Manager however if a mouse is not available, this utility can be controlled using the arrow and tabulation keys. The table below indicates the correspondance between the keys and buttons:

STORAGE MANAGER FUNCTION	KEY SEQUENCE
Select menu or function button	Alt + Underlined letter
Exit menu	Esc
Move to next icon	Tab
Move to previous icon	Shift + Tab
View icon information window	Enter
Scroll up or down	Up or down arrows
Scroll one page	Page Up or Page Down
Move to top or bottom	Home or End
Select drive icon	Space
Select multiple drive icons	Shift + Space

The first screen always displays the physical configuration of the SCSI subsystem. It is a graphical representation of each RAID DPT SCSI controller with the connected SCSI peripherals, along with separate icons representing each peripheral.

Selecting the Switch View button switches the physical configuration window to the SCSI subsystem's logical and physical configuration window. This view is displayed on the left-hand side of the screen and gives a graphical representation of the devices as seen by the operating system.

The SCSI devices that are not HDUs will always have the same logical and physical icon. Hard disks, however, can be seen by the operating system as individual hard disks or as members of an array group. In both cases, the hard disk or array group is displayed on the left-hand side of the screen as a Logical Storage Unit (LSU) icon.

INITIAL INSTALLATION

After configuring the GO2044 with the ECU and the GO2061, GO2098 and GO2173 with the DPT Configuration Utility, the next step is to run the Storage Manager Utility to begin the system's initial installation phase. You will need a blank diskette.

- **Diskette Kit** - Insert the diskette containing the Storage Manager Utility in drive A or B and reboot the system. The utility will be automatically activated. Select "Initial System Installation" from the "Installation" menu.

CD-ROM Kit - Insert the Orchestra Systema Boot diskette in drive A, the Orchestra Systema CD-ROM into the CD-ROM drive and reboot the system. From the main menu displayed, select "System Configuration", followed by "DPT Storage Manager" and "Create RAID Info".

- Once the physical configuration of the subsystem is displayed, check the following:
 - Check that all the SCSI devices installed are displayed with the correct ID. If any device is missing, exit the utility and check the hardware connections.
 - Check all the devices marked with a warning flag. This flag indicates that the utility has detected a problem with that device. View the device's Icon Information Window for information on the problem detected.
 - View the Icon Information Window for each SCSI device to make sure that all the parameters are correct. As far as the controller is concerned, make sure that the total cache memory installed corresponds to the value displayed.
- After checking the SCSI configuration, Storage Manager will ask you to specify the operating system that will be installed. This utility will then check whether the selected operating system will support the SCSI configuration. If there are problems, the utility will suggest possible changes to the configuration.
- **Diskette Kit** - The utility will ask you to insert the second disk which contains the drivers for the different operating systems.

CD-ROM Kit - To install the operating system reconfigure it using the SCSI peripheral drivers; select the Operating Systems option from the Orchestra Systema Main Menu and follow the instructions provided by the online documentation associated to the selected operating system.

- Storage Manager will ask you to insert a blank diskette into the drive. This utility will then format the diskette and store into it the utilities and supplemental operating system drivers. This Supplemental Utility diskette will be used during operating system installation.

Note: *Some versions of some operating systems do not have the Storage Manager Utility and Supplemental Drivers. In this case some functions are not available.*

- Storage Manager will now return to its normal operating mode, where you can create different disk array groups. This will be explained further on.

Note: *Array groups can be created or modified at any time after the installation of the operating system. However, if the boot drive is an array, then it must be created before the operating system is installed.*

- After defining the arrays, exit Storage Manager. You will be asked to save the configuration. If you choose to save the configuration, any array which was created or modified will be built at this time. This process may take several hours to complete. You can exit the utility and perform other activities on the system while the build operation continues. It is not possible, however, to access the array groups being built until the build operation has completed. You can also monitor the progress of the array build operation through the Icon Information Window.
- Storage Manager will now ask you to install the operating system and instruct you on using the Supplemental Utility diskette if one was created.

Note: *If the system is powered off during the creation of a disk array, reconstruction will start from the beginning when the system is powered on again.*

ADDING SCSI PERIPHERALS

Each time a new SCSI peripheral is connected to the RAID controller, or whenever a new RAID controller is added, you will have to run the Storage Manager Utility to check the hardware configuration.

- Run the Storage Manager Utility.
- Select "Add SCSI Devices" from the "Installation" menu. The physical configuration of the SCSI subsystem is displayed. Make sure this configuration is correct.
- Storage Manager will verify if the selected operating system supports the new SCSI configuration. If there are problems, the utility will suggest possible changes to the configuration.
- Storage Manager will now return to its normal operating mode in which you can create different disk array groups as explained further on.
- When the array has been defined, exit Storage Manager. You will be asked to save the configuration. If you choose to save the configuration, each array that was created or modified will now be built.
- Reconfigure the operating system with the new peripheral (see the section entitled Notes on Operating System Installation).

F

ARRAY GROUPS

The "Create Array Group", "Modify Array Group" and "Delete Array Group" function buttons are on the Logical Configuration View screen tool bar. These buttons allow drives to be added or removed from RAID-0, RAID-1 or RAID-5 array groups.

Any combination of RAID-0, RAID-1 or RAID-5 disk arrays can be implemented. Each single array built, in other words each single hardware array, must be connected to the same SCSI controller.

The table below lists the number of drives in every Hardware Array and Array Group of each RAID level.

RAID Level	Number of Drives in the Hardware Array	Number of Disks in the Array Group
RAID-0	2 or more	2 or more
RAID-1	2	Even number per controller
RAID-5	3 to 7	3 or more per controller

Hardware Array is the name for a group of HDUs with the same RAID level connected to the same SCSI controller. Any number of Hardware Arrays of the same level, even on different SCSI controllers, can be combined by the software driver into larger Dual-Level Array Groups. Array Group is a generic term for a group of Hardware Arrays or Dual-Level Arrays of the same type that appear as a single LSU to the operating system.

Note: *Mirroring (RAID-1) is not possible, for example, between HDUs connected to different controllers, but, in the case of the GO2061, GO2098 and GO2173 controllers, it is possible between HDUs connected to the channels of the same controller and therefore also on different PEMs.*

CREATING AN ARRAY GROUP

- From the "Logical Configuration View" screen, select the "Create Array Group" function button.
- When the "Select Array Type" window appears, specify whether or not fault-tolerance is desired and whether the Array Group is to be optimized for capacity and performance. After making your selection, the "Chosen Array Parameters" will automatically change to indicate which RAID level best suits your requirements. However, you can choose a different RAID level by selecting the Override function button. Select "Continue" to proceed after this phase.
- The "Logical Configuration View" screen will reappear with the option "Choosing drives for Array (RAID-n)", where n is the RAID level previously selected. At the top of the screen there are three new function buttons labeled "Include Drive", "Remove Drive" and "Done". Select now the drives to be included in the Array Group.
 - To select a drive to be included in the Array Group, first select the drive's corresponding icon and then select the Include Drive button. The drive icon will move from the other peripheral icons to the area at the bottom of the screen where the Array Groups are displayed. Scroll the screen to view the drive in the Array Group. To speed up this process instead of selecting one drive at a time, multiple drives can be selected while holding down the Shift key.
 - To remove a drive from the Array Group, select the corresponding drive icon and then select the Remove Drive button. The drive icon will return with the other peripheral icons. To speed up this process instead of selecting one drive at a time, multiple drives can be selected while holding down the Shift key.
 - During this process some drives may be displayed in blue. This indicates that these drives cannot be included in the Array Group unless the current configuration is changed.
- When you have finished selecting the drives to include in the Array Group, select Done. Scroll the screen to view the new display. An LSU with a black flag (Missing) is displayed under the new Array Group until the array build process is complete.
- To create the array, select the Set System Configuration option from the File menu and wait for reconstruction to begin.

MODIFYING AN ARRAY GROUP

You can change the number of drives included in the Array Group by following the procedure explained below:

- If data has been written to the Array Group, you will need to backup that data on a storage media such as a streaming tape. This is because all the data on the Array Group will be destroyed when the array is modified.
- Select the "Modify Array Group" button from the "Logical Configuration View" screen.
- A message is displayed confirming the selection of the Array Group to be modified. Select "OK" then select the LSU icon corresponding to this Array Group.
- Drives can now be added or removed from the Array Group as explained in the section entitled "Creating an Array Group".
- Select Done at the end of the modification process. An LSU with a black flag (Missing) is displayed under the modified Array Group until the new array is built.
- Once the array build process has completed, restore the backup data onto the modified Array Group.

CANCELLING AN ARRAY GROUP

- Select the "Delete Array Group" button from the "Logical Configuration View" screen.
- A message is displayed confirming the selection of the Array Group to be cancelled with a warning that all the data stored in the selected Array Group will be lost. Select the LSU icon corresponding to the desired Array Group.
- Another message is displayed warning you again that all the data stored in the selected array will be lost, in addition to the "OK" button continue with the cancellation process, and the "Cancel" button to not cancel the array. An Array Group is not physically cancelled until you exit Storage Manager and save the changes made, or until you select the "Set System Configuration" from the "File" menu. The data cannot be recovered once these operations have been performed.

NAMING AN ARRAY GROUP

F

Any Array Group or Hardware Array within an Array Group can be named at any time by selecting the "Name" function button for the LSU or Hardware Array to be named. This name will then be displayed under the array icon and used in event messages.

LOGICAL SCSI DEVICE ADDRESS

Every SCSI device or Array Group visible to the operating system is assigned a Logical Device Address by the Storage Manager. This address is used by the operating system to access the device or array, and in some cases must be entered during the operating system installation or configuration process. The logical drive address appears in parenthesis under the logical device icon in the Logical Configuration View, and consists of four fields: HBA, Bus, Device and LUN. SCSI devices that are not part of an Array Group are automatically assigned a Logical Drive Address as follows:

- **HBA** - Host Bus Adapter, represents the controller to which the logical device is attached. Bus slots are scanned, starting from the first slot, to detect the RAID controllers. As the controllers are found, they are assigned incremental numbers, starting from 1.
- **Bus** - SCSI bus for the HBA. The SCSI controller has only one SCSI channel, always assigned 0.
- **Device** - Represents the SCSI ID for that device.
- **LUN** - Represents the SCSI LUN for that device.

Arrays are automatically assigned the lowest Logical Device Address of any drive in that array.

STATUS

Storage Manager indicates the status of the hard disks and Array Groups through status flags on the device icons. Changes in these status conditions may be logged and optionally transmitted.

STATUS FLAGS

Certain status conditions of arrays or hard disk drives are indicated by a flag that changes the color of the device icon as indicated in the table below. If the status of the array or hard disk drive is optimal, no flag will be displayed.

FLAG	DRIVE	ARRAY	COMMENT
White	Verifying	Verifying	Performance may be affected
Yellow	Warning	Degraded	Warning: risk of failure
Red	Failed	Failed	Device failure has occurred
Blue		Building	Array is being built
Black	Missing	Missing	Device is not responding

HARD DISK DRIVE STATUS

Hard disks always have one of the status conditions shown below. Any change in this status is considered a level 4 event and should typically be transmitted to the system administrator.

- **Optimal** - The drive is fully functional and is running diagnostics which could degrade performance.
- **Verifying** - Diagnostics are being run on the drive, performance may be affected.
- **Warning** - Drive failure may be imminent.
- **Failed** - Nonrecoverable drive failure has occurred.
- **Missing** - The drive is physically missing or does not respond to commands on the SCSI bus.

ARRAY STATUS

Array Groups will always have one of the status conditions shown below. Any change in this status is considered a level 4 event and should typically be transmitted to the system administrator.

- **Optimal** - The Array group is fully functional, has no failed drives and is not performing operations such as a Verify or diagnostics which might degrade performance.
- **Verifying** - Diagnostics are being run on an Array Group or consistent redundancy is being verified. Performance may be affected.
- **Degraded** - A drive failure has occurred in a fault-tolerant Array Group. The data of the faulty drive cannot be read but can be recovered by the other drives in the array. Performance may be affected.
- **Failed** - A drive failure has occurred in a non fault-tolerant Array Group or two or more drives have failed in a fault-tolerant Array Group. The data on the Array Group are lost.
- **Building** - A fault-tolerant Array Group which has just been created is now being initialized, or data is being recreated onto a replacement for the failed drive.
- **Missing** - The Array Group has been created or modified but the configuration has not yet been saved and the array has not been built.

EVENTS

Events can be caused by fault conditions and some non-fault conditions. The four levels of events are listed below:

- **Level 1: Soft Error** - An error occurred on a disk drive but upon retry, the operation was successful.
- **Level 2: Recoverable Error** - An error has occurred on a hard disk drive or on a controller, but the data can be recovered because the array containing the hard disk is fault-tolerant.
- **Level 3: Unrecoverable Error** - An error has occurred on a hard disk drive or on a controller and the data cannot be recovered.
- **Level 4: Status Change** - The status of an Array Group or hard disk drive has changed.

F

EVENT LOGGING

When events occur, they are automatically stored in the controller cache memory. You can also specify that Storage Manager maintain an event log on disk (see the section entitled "Event Broadcasting").

The contents of these event logs can be displayed by selecting the "Event Log" button in any controller, hard disk drive or array "Icon Information Window". By selecting "Display Level Threshold", you can also specify to only display events of a specific level or higher. Select the "View Log" button to display the messages of the selected events.

EVENT BROADCASTING

In some operating systems you can specify that event messages be broadcast to users, groups, devices, system error log or through E-mail. The options provided vary, depending on the operating system. By selecting "Event Broadcast Control" from the "Options" menu you can configure the following options:

- **Time Interval** - A time interval is specified within which the event logs stored in the SCSI controller cache memory can be read. Within this time frame each new event message will be broadcast to each specified destination.
- **Send Events to DPT Log File** - If this option is enabled, events read from the controller cache memory are recorded in the "DPT Event Log File" on hard disk. This makes sure that events will not be lost when the system is powered off. The "DPT Event Log File" can also contain a number of events limited only by the available disk space.
- **Broadcast Threshold** - Click on the arrow buttons to determine the event levels to broadcast.
- **Broadcast to Users/Groups** - When enabled, each user and group specified in the "Users/Groups" list will have event messages broadcast to them.
- **Send to System Error File** - When enabled, event messages will be sent to the System Error File specified in this option.
- **Broadcast to Devices** - When enabled, each device specified in the "Devices" list will have event messages broadcast to it.
- **E-Mail to Users/Groups** - When enabled, each user and group specified in the "Users/Groups" list will have event messages broadcast to their E-Mail.
- **Users/Groups List** - This list specifies which users and groups must receive event messages. The names of the users and groups can be listed in any order.

DRIVE FAILURES

If a hard disk drive belonging to an Array Group fails, the relative drive icon will appear with a red flag in both the "Physical" and "Logical Configuration View" screens.

AUDIBLE ALARM

Failure of a hard disk drive in an Array Group will cause an audible alarm to sound on the SCSI controller. This alarm will go off automatically when Storage Manager is entered. Audible alarms can also be turned off by selecting "Turn Off Audible Alarms" from the "Options" menu.

FAILURES IN NON-FAULT TOLERANT ARRAYS

If a failed drive belongs to a non-fault tolerant Array Group (RAID-0), the icons for that drive, for the Hardware Array and for the corresponding LSU are displayed with a red flag indicating that the data on that Array Group has been lost.

HOT SWAPS

If the failed drive belongs to a fault-tolerant Array Group (RAID-1 or RAID-5) and that drive is not protected by a Hot Spare, the icon for the failed drive is displayed with a red flag (Failed) while those for the Hardware Array and LSU containing the failed drive are displayed with a yellow flag (Degraded).

Physically replace the failed drive without powering off the system; the data from the failed hard disk are automatically rebuilt on the new drive. During this data rebuild phase, the icon for the new drive is displayed with a white flag while those for the Hardware Array and LSU will remain yellow until the rebuild process is complete.

Note: *To replace a faulty HDU there is no need to activate Storage Manager unless the replacement HDU has already been used on other systems in RAID configurations, in which case see the note in section "Storage Manager Notes and Limitations".*

HOT SPARES

Any non-arrayed hard disk drive can be assigned as a Hot Spare. Hot Spares are not accessible to the operating systems and are used to automatically replace failed drives which belong to fault-tolerant Hardware Arrays. Hot Spares can only protect drives attached to the same SCSI controller as the Hot Spare and whose capacity is less than or equal to that of the Hot Spare. To make a drive a Hot Spare, select the "Make HotSpare" button in the "Icon Information Window" of the corresponding drive.

By selecting the "Remove HotSpare" function in the "Icon Information Window" of the Hot Spare drive will cause this drive to be assigned as a normal hard disk drive accessible by the operating system.

When an array protected by a Hot Spare has a drive failure, the SCSI controller will automatically begin to rebuild the missing data onto the Hot Spare. During this process, Storage Manager will swap the positions of the failed drive with the Hot Spare in the "Logical Configuration View". The icon for the failed drive is displayed with a red flag in the old position of the Hot Spare, and the Hot Spare will appear as a member of an Array Group with a white flag to indicate that a rebuild operation is in process. The array and LSU icons are displayed with the yellow, degraded, flags. At the end of the rebuild process, the white and yellow flags will disappear and the drive will be displayed as a normal member of the Array Group. The red flag will remain on the failed drive until that drive is replaced.

The physical replacement of the failed drive is always performed without powering the system off. At the end of this process the installed drive will become the new Hot Spare.

F

DIAGNOSTICS

Diagnostics can be performed on hard disk drives or Array Groups by selecting the "Diagnostics" button in the "Icon Information Window" of any hard disk, Hardware Array or Array Group LSU icon. Diagnostics which can be performed include device buffer tests, which send and receive data from a test buffer on the hard disk, and media read tests which test for error-free data in each hard disk sector.

Verification diagnostics can also be performed on fault-tolerant arrays to ensure the consistency of the information. In the case of RAID-1, the mirrored drive pairs are compared sector by sector to make sure that both drives contain identical data. In the case of RAID-5, parity is recalculated and checked against stored parity information.

Note: *Under normal conditions, data inconsistencies for fault-tolerant arrays should never occur. However, a power failure which interrupts an array write operation could cause inconsistency of the data stored.*

Diagnostics can be performed either for a specified number of times or in continuous loop. Diagnostics can also be performed on multiple drives or arrays simultaneously. The status of the diagnostics is displayed at the bottom of the diagnostics window, and the statuses are: "Running", "Completed", "Stopped on Error" and "Stopped by User".

BACKGROUND TASK PRIORITY

Diagnostics and Verification can be performed without stopping the operations on the drives or arrays. By selecting the "Set Background Test Priority" from the "Options" menu, background tasks such as diagnostics, array verification and array rebuild will occupy a specified percentage of the total activities performed on the device. During periods of inactivity, 100% of the device will be dedicated to background tasks. Selecting a priority for background tasks will cause these tasks to occur only if there is no adapter activity for 250 ms, thus minimizing the impact on performance.

DISPLAYING I/O OPERATIONS

The SCSI controller automatically tallies I/O operations in cache memory. These numbers can be viewed by selecting the "I/O Monitor" button in the "Icon Information Window" for every hard disk drive, Hardware Array or Array Group LSU. The reported operations include:

- Write operations with a length of less than or equal to:

1 KB	2 KB	4 KB	8 KB
16 KB	32 KB	64 KB	128 KB
256 KB	512 KB	1 MB	
- Read operations with a length of less than or equal to:

1 KB	2 KB	4 KB	8 KB
16 KB	32 KB	64 KB	128 KB
256 KB	512 KB	1 MB	
- Cache Hit and Miss
- Array access which cross stripe boundaries, in other words the size into which a write operation is divided so that it can be stored on more than one disks.

By analyzing the relative numbers of read and write operations of various lengths, the best array architecture and stripe size can be determined.

- Systems with a high percentage of write operations often the SCSI controller cache memory, and should preferably use RAID-1 instead of RAID-5 if fault tolerance is required.
- A high percentage of short read and write operations (4 KB or smaller) still remains an indication that the SCSI controller cache memory is often used. If the percentage of cache hits is low, then more cache memory will need to be added to increase the hit rate. However, this operation cannot be performed on these systems.
- Single-user environments with a high number of long read and write operations (64 KB to 256 KB) may benefit from a small stripe size and therefore with a greater division of the records to be logged. Making this selection, however, will affect performance for shorter record sizes.
- Extremely long read and writes (512 KB to 1 MB) will provide better performance with the default 128 KB stripe size.
- A high number of accesses which cross stripe boundaries in a multi-user system indicate that the stripe size has been set too small.

REMOTE OPERATION

Storage Manager consists of two primary parts. The Storage Manager GUI is a user interface management program. The GUI communicates through a message protocol to one or more systems equipped with the RAID DPT controller and on which Storage Manager is installed.

When Storage Manager is run, it automatically loads the hardware configuration of the local system's SCSI subsystem, but can also view the configuration of a remote system by selecting the "Make Connection" button from the "Communications" menu.

REMOTE OPERATION OVER A NETWORK

By selecting "Make Connection" from the "Communications" menu, Storage Manager will search the local network for any fileserver with Storage Manager installed, and displays the icons representing these servers in the "Available Connections" window. By selecting any one of these icons and entering the relative password, the SCSI subsystem of the remote server can be accessed.

CONFIGURATION FILES

The following four options are available in the File menu:

- Read System Configuration.
- Set System Configuration.
- Load Configuration File.
- Save Configuration File.

F

The "Read System Configuration" option is automatically activated when Storage Manager is accessed. It can also be activated by the "File" menu to restore the hardware configuration of the SCSI subsystem before any changes which have been made are stored.

The "Set System Configuration" stores in non-volatile memory any changes that may have been made during the SCSI subsystem configuration phase. If any Array Groups have been created or modified, selecting "Set System Configuration" will cause array build operations to begin.

The "Load Configuration File" option allows the current configuration stored in Storage Manager to be loaded and to be applied to the existing hardware.

The "Save Configuration File" option allows the current configuration to be stored in a file on diskette, for future use. With this feature you can configure the SCSI subsystem for other systems.

NOTES ON OPERATING SYSTEM INSTALLATION

This section gives some notes on the installation of the operating systems listed below. Before installing an operating system you need to physically install one or more RAID DPT SCSI controllers and the connected SCSI peripherals. Then run the Storage Manager Utility and select "Initial System Installation" from the "Installation" to load the correct configuration. Create the Supplemental Utility diskette, if necessary, and then create the required Array Groups. Each Array Group will be viewed by the operating system as a single hard disk.

Listed below are the operating systems with full driver support:

- IBM OS/2 rel. 2.1 and 3.0 (WARP).
- Microsoft Windows NT ver. 3.5 and later.
- Novell Netware ver. 3.12 and 4.1 and later.
- UnixWare rel. 2.0x.
- SCO Unix System V ver. 3.2.4.2 and later.
- Banyan Vines ver. 5.53 and later (DPT EISA controller only).
- Olivetti UNIX System V rel 4.0 ver. 2.3 and later.
- SCO Open Server 5.0

Note: OS/2 is supported from Storage Manager release 1.03 onwards.

OS/2

IBM OS/2 version 2.1 and later releases have DPT drivers included in the operating system. A Supplemental Utility diskette providing useful information on its own installation is automatically created when Storage Manager is run. This diskette with the appropriate installation drivers is also created for IBM OS/2 2.0 and 1.3 which do not have DPT drivers included in the operating system.

If OS/2 is to reside with another operating system then that operating system must be installed before OS/2.

To install OS/2 on a hard disk or Array Group, see the Installation Guide.

Note: *When using the RAID PCI board under OS/2 2.11, install the Supplemental disk as follows:*

- *Open "OS/2 System Folder" then "System Setup Folder".*
 - *Select "Device Driver Install".*
 - *Insert the Supplemental disk into the drive.*
 - *Select "Install". A message is displayed indicating the current state: an error is signalled if the controller is not found.*
- This limitation is currently solved.*

WINDOWS NT

Microsoft Windows NT has DPT drivers included in the operating system. However, if more recent drivers exist, the Supplemental Utility diskette will be automatically created when Storage Manager is run, in which case Storage Manager will provide information on how to install the diskette.

To install Windows NT on a hard disk or Array Group, see the System Guide.

NOVELL NETWARE

NetWare 386 3.12 and 4.x has DPT drivers included in the operating system. A Supplemental Utility diskette, which also provides the information on its own installation, is automatically created when Storage Manager is run. This diskette containing installation drivers is also created for NetWare 386 version 3.11, which does not have the drivers included in the operating system.

To install NetWare on a hard disk or Array Group, see the Installation Guide.

SCO 3.2.4.2 and SCO OPEN SERVER 5.0x

SCO 3.2.4.2 and 5.0x have the DPT drivers included in the operating system. However if more recent drivers exist, a Supplemental Utility diskette will be automatically created when Storage Manager is run. In this case Storage Manager will provide all the information on how to install the diskette.

This operating system is preconfigured to support an HDU with ID=0, an STU with ID=2 and a CD-ROM with ID=5.

This preconfiguration could raise problems when an HDU with ID=2 or ID=5 is installed in the system. In resilience systems, in fact, the HDUs are automatically assigned an ID according to the position they occupy on the rack, and therefore the third HDU installed, which in resilience systems will have an ID=2, will enter into conflict with the STU, while the last HDU which will have an ID=5 will create conflict with the CD-ROM.

In the case of a maximum configuration, when the rack of a resilience system is full up with HDUs (RAID-5 for example), it is impossible to install the SCO operating system from the STU and CD-ROM. In this case, only operating system installations from floppy disk will be accepted.

To install SCO on a hard disk or Array Group, see to the Installation Guide.

To use the STU exclusively for backup operations, you will need to remove the drive with the "mkdev tape" command and install the STU again this time using an ID greater than 5.

UNIXWARE

All UnixWare versions have DPT drivers included in the operating system. However, if more recent drivers exist, a Supplemental Utility diskette is automatically created when Storage Manager is run. In this case Storage Manager will provide all the information on how to install this diskette.

F

To install UnixWare on a hard disk or Array Group, see the installation Guide.

BANYAN VINES

Banyan VINES version 5.53 and later has support for DPT boards integrated into the operating system kernel. The Supplemental; Utility diskette is not created when Storage Manager is run.

To install Banyan VINES on a hard disk or Array Group, see the Installation Guide.

OLIVETTI UNIX SV 4.0

Olivetti Unix SV 4.0 and later releases have DPT drivers included in the operating system. The Supplemental Utility diskette is not created when Storage Manager is run.

To install Olivetti Unix on a hard disk or Array Group, see the Installation and Configuration Guide.

STORAGE MANAGER FOR DPT EISA CONTROLLER EVOLUTION

DATE	REV.	REASON FOR CHANGE	APPLIC.
11/94	1.00	The Storage Manager for the EISA controller is introduced and comes on the following disks: - DPT/Olivetti EISA Storage Manager Utilities disk 1/2 code 2690954 B or 2692824 K - DPT/Olivetti EISA Storage Manager Utilities disk 2/2 code 2690955 F or 2692823 W.	Factory

STORAGE MANAGER FOR DPT PCI CONTROLLER EVOLUTION

DATE	REV.	REASON FOR CHANGE	APPLIC.
03/95	1.01	Introduction of Storage Manager for PCI controller on diskettes: - DPT/Olivetti PCI Storage Manager Utilities disk 1/2 code 2691011 R - DPT/Olivetti PCI Storage Manager Utilities disk 2/2 code 2691012 V.	Factory
05/95	1.02	The "Unrecoverable Error" problem under SCO with a RAID PCI SCSI controller installed in the system is corrected. The device drivers were the same as those on the Storage Manager diskettes. The diskettes are: - DPT/Olivetti PCI Storage Manager Utilities disk 1/2 code 2691022 W - DPT/Olivetti PCI Storage Manager Utilities disk 2/2 code 2691023 S.	Factory
05/95	1.03	Support for OS/2. The diskettes are: - DPT/Olivetti PCI Storage Manager Utilities disk 1/2 code 2691029 L - DPT/Olivetti PCI Storage Manager Utilities disk 2/2 code 2691030 J.	Factory
06/95	2.00	Single Storage Manager version for EISA and PCI controllers and compatibility with Server View. Only the EISA controller is supported by the Vines 5.5x O.S.. This S.M. release must be combined with Resilience Support rel. 2.01 or later for SNX 160/RS, rel. 2.00 or later for SNX 200/400/RS. The diskettes are: - DPT/Olivetti Storage Manager Utilities disk 1/2 code 2691095 S - DPT/Olivetti Storage Manager Utilities disk 2/2 code 2691096 E.	Factory
10/95	2.01	New release that allows the generation of the DPT driver for NT 3.51 and that solves some problems in the NT environment. The following disks are available: - DPT/Olivetti Storage Manager Utilities disk 1/2 code 2691208 M - DPT/Olivetti Storage Manager Utilities disk 2/2 code 2691209 R.	Factory
11/95	2.1	New release supporting Server View 1.1. The following disks are available: - DPT/Olivetti Storage Manager Utilities disk 1/2 code 2691245 G - DPT/Olivetti Storage Manager Utilities disk 2/2 code 2691246 L.	Factory
05/96	2.2	New S.M. release aligned with Server View 2.01. The following disks are available: - DPT/Olivetti Storage Manager Utilities disk 1/3 code 2692121 P - DPT/Olivetti Storage Manager Utilities disk 2/3 code 2692122 T - DPT/Olivetti Storage Manager Utilities disk 3/3 code 2692123 X.	Factory

STORAGE MANAGER NOTES AND LIMITATIONS

- When an HDU is added to, or removed from, the system, if using a disk which has already been used on another system equipped with a RAID controller make sure that the disk is free of logical RAID markers, in other words it must have been cleared ("ZAPPED").
Since the HDU_DIA.EXE diagnostic does not guarantee the cancellation of blocks containing RAID data, there are two ways in which the disk can be cleared:
 - Boot the system from the Storage Manager diskette and press CTRL-C to prevent the utility from loading automatically. Invoke the utility by typing "dptmgr /ZAP" at the DOS prompt. The utility will display the list of drives connected and will ask which drive needs to be "ZAPPED".
 - Another way of clearing one or more hard disks is by using the CLEANHDU Utility. This utility comes on diskette and allows you to recover via software the hard disks considered as being unusable (turn to Appendix M).
- To avoid Storage Manager failures on the SNX 200/400, the system resources that are automatically allocated for the motherboard (COM1/COM2, parallel port, floppy interface, keyboard/mouse interface) cannot be disabled when the system has a 32 MB memory.
- Even though nearly all the operating systems contain DPT drivers, always use the ones generated by the Storage Manager. The only operating system that needs to use its own drivers is OS/2 version 3.0 (WARP).
- Updated information and limitations are provided in the README file of the Storage Manager Utility.
- With systems with more than two HDUs, activate Storage Manager before installing the operating system even though a RAID configuration is not required. A RAID configuration needs to be created, and then cancelled.
This limitation is currently solved.

F

RAID

RAID (Redundant Arrays of Independent Disks) is an improvement made to the usage of mass memory especially in multiuser environments. Essentially, it combines two or more hard disks into a single array of disk drives capable of providing fault-tolerant performance.

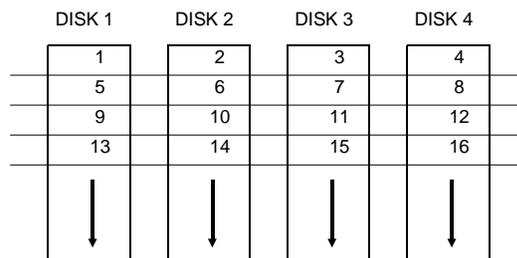
The Mean Time Between Failure (MTBF) of the array will be equivalent to the MTBF of an individual drive divided by the number of drives in the array. Thus the MTBF of an array of drives would be too low, but to increase the reliability of the data stored fault-tolerant arrays can be created.

Different types of RAID fault-tolerant architectures, the main ones ranging from RAID-1 to RAID-5, were defined, each offering different characteristics and performance. Also the RAID-0 array was defined, but it is a non-redundant array of disks. The RAID DPT SCSI controllers allow the more popular RAID-0, RAID-1 and RAID-5 disk arrays to be configured.

RAID-0 OR DISK STRIPING

RAID-0 or Disk Striping is a method by which data is evenly distributed across all the hard disk drives of the array. RAID-0 is a non fault-tolerant array. If one drive in a RAID-0 array crashes, the entire array crashes as it is impossible to reconstruct the data of the hard disk that crashed. However, this array is the only configuration that reserves the entire disk space for data storage, therefore it is available at a low cost and provides the best possible performance in terms of throughput (the performance increases proportionally to the number of disks used in the array). For this reason RAID-0 can be used in systems where performance is priority over reliability.

Striping partitions each drive's storage space into stripes whose size depends on the number of variable sectors. The stripe can be as small as one sector (512 bytes) or as large as several megabytes. These stripes are then interleaved on all drives as shown in the following figure.



The type of operating environment determines the size of the stripe that should be used. Most multiuser operating systems like UNIX, Novell NetWare and NT support overlapped I/O operations across multiple drives. To maximize the performance of the disk subsystem, the I/O load must be balanced across all the drives so that each drive can be kept busy as much as possible. In multiple drive systems without striping, the I/O load is never perfectly balanced since the drives containing data files will be frequently accessed while the other drives will only be rarely accessed.

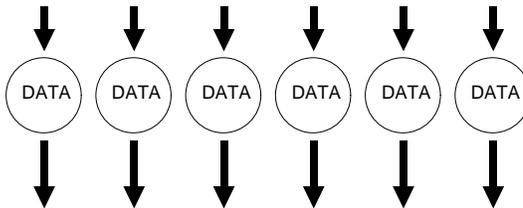
The best performance is obtained by striping the drives in the array with stripes large enough so that each record falls entirely within one stripe. In this way the records will be evenly distributed across all drives and the I/O load will be balanced, so all the drives in the array will be kept heavy during heavy load conditions. This allows each drive to work on a different I/O operation, thus maximizing the number of simultaneous I/O operations which can be performed by the array.

In single-user systems which access large records (when acquiring data or graphics), small stripes (512 bytes) must be used so that each record will span across all the drives in the array. This causes long record accesses to be performed faster since the data transfer on multiple drives occurs faster. Unfortunately, small stripes do not allow multiple overlapped I/O operations since each I/O usually involves all drives. However, no limits are imposed on operating systems that are not multi-user, multi-tasking or that do not allow overlapped disk I/O.

The inconvenience of using small stripes is that synchronized spindle drives are required in order to keep performance at top level even during short accesses. Without synchronized spindles, each drive in the array will be at different random rotational positions. And since an I/O cannot be completed until each drive has accessed its part of the record, the drive with the slowest rotational speed will determine the completion of the I/O operation. Synchronized spindles assure that every drive in the array reaches its data at the same time. The access time of the array will thus be nearly equal to the average access time of a single drive.

F

WRITES OCCUR
SIMULTANEOUSLY ON EVERY
DRIVE



READS OCCUR
SIMULTANEOUSLY ON EVERY
DRIVE

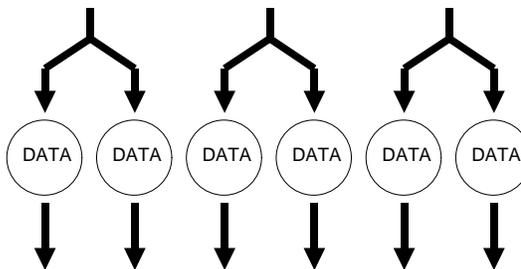
RAID-1 OR MIRRORING

RAID-1, better known as Disk Mirroring, is a fault-tolerant array consisting of a pair of disk drives which store duplicate data but which appear to the system as a single drive. RAID-1 is distinguished from other fault-tolerant RAID arrays as all the data of the array are duplicated on a backup mirrored disk.

This solution offers a high degree of reliability, since the data can only be lost if both mirrored disks fail. It is, however, costly since it requires double hard disks as the storage space available to the user is 50% of the total disk subsystem capacity.

Writes must go to both drives in a mirrored pair, so that the information on the drives is kept identical. Each individual drive, however, can perform individual read operations. RAID-1 delivers the best performance of any fault tolerant array in multi-tasking environments.

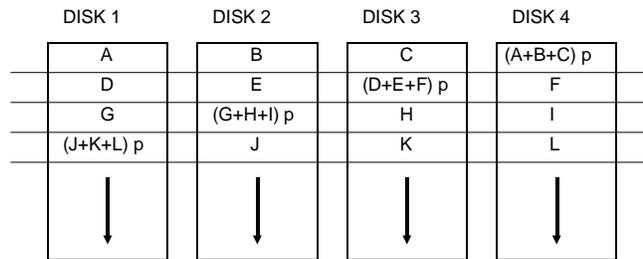
DUPLICATE DATA ARE
WRITTEN TO PAIRS OF DRIVES



READS OCCUR
SIMULTANEOUSLY ON EVERY
DRIVE

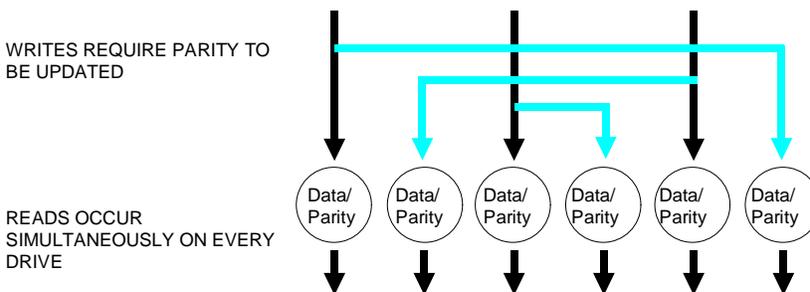
RAID-5

RAID-5, also known as Rotating Parity Array, avoids the write bottleneck present with RAID-3 and 4 in which there is a single drive dedicated to parity checking. RAID-5 uses large stripes so that each drive takes turns storing parity information for a different series of stripes.



Since there is no dedicated parity drive, all drives contain data, and read operations can be overlapped on every drive in the array thus increasing the data throughput.

Write operations will typically access a single data drive plus the parity drive for that record, so different records store their parity on different drives and write operations can be overlapped.

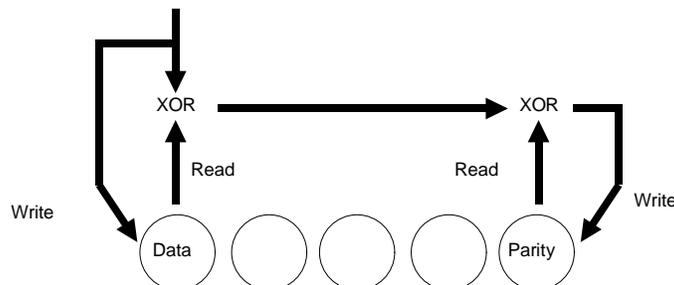


RAID-5 is a fault-tolerant array since the parity information distributed on the different disks allow the reconstruction of the lost data in the event one of the drive in the array fails. Parity is calculated by the controller for each stripe of stored data, and recorded on another disk. If a drive fails, this drive is replaced and the data lost is reconstructed by the controller. The array will continue to provide correct data even if the faulty drive has not been replaced. In fact, corrupt are automatically reconstructed by the parity information and are made available on the bus. Under these conditions, the array offers lower performance. RAID-5 arrays provide greater storage efficiency than RAID-1 arrays since for an array to become fault-tolerant, only parity information is stored and not a complete copy of all the data. The result is the possibility of combining any number of drives instead of just an even number, and to sacrifice the effective capacity of a single drive to store parity information, instead of just half its capacity.

A greater storage efficiency corresponds to a loss in performance. In fact, each write requires four I/O operations instead of the two required by the RAID-1 array.

When data is written to a RAID-5 array, the parity information must be updated in one of the two following ways. The first way is straightforward but very slow. The parity information is the XOR of the data on every drive of the array. Therefore whenever the data of any one drive is changed, the other drives in the array which store data are read and XORed to create the new parity. This requires accessing every drive in the array for each write operation.

The second way of updating parity is usually the more efficient. It consists of finding out which data bits were changed by the write operation and then change the corresponding parity bits. This is accomplished by firstly reading the old data to be overwritten. This data is then XORed with the new data which is to be written. The result is a bit mask which has a one in the position of every bit that has been changed. This bit mask is then XORed with the old parity information which is read from the parity drive of the array. In this way it is possible to determine which bits in the old parity have changed and the new updated parity is then written back to the parity drive. This results in only two reads, two writes and two XOR operations, instead of a read or write and XOR for every drive in the array.

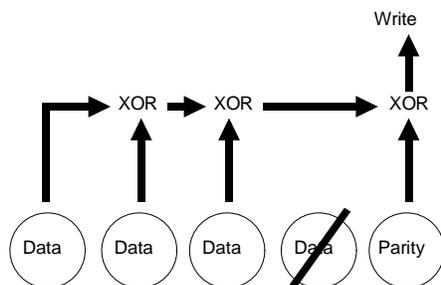


F

The disadvantage of storing parity information rather than a copy of the data is the extra time taken during write operations to regenerate parity information. This additional time reduces write performance for RAID-5 arrays over RAID-1 arrays by a factor of between 3/5 and 1/3, in other words RAID-5 writes are between 3/5 and 1/3 the speed of RAID-1 writes. For this reason, RAID-5 arrays are not recommended for applications in which performance is important (with the exception of applications which never require data writes).

Recreating Lost Data

When a drive fails in a RAID-0 array, the entire array can no longer be used. In a RAID-1 array, a failed drive has little impact on performance since data can be read from the mirrored drive. In a RAID-5 array, however, the data must be reconstructed by reading and XORing the corresponding data stripes from the remaining drives in the array. This process is relatively slow. The more drives contained in the RAID-5 array, the slower the operation becomes but also the less probable case.



Rebuilding a Failed Hard Disk

When a failed drive is replaced in a RAID-1 or RAID-5 array, the SCSI controller will rebuild the lost data onto the new drive. This rebuild operation occurs during normal reads and writes in the array. In the case of RAID-1 arrays, the rebuild will occur relatively fast since it involves just copying the data contained in the mirrored drive to the newly replaced drive. In the case of RAID-5 arrays, the new data to be rewritten to the newly replaced drive must be recreated by reading and XORing the corresponding data stripes from the remaining drives in the array. This causes a considerable degrade in performance during the rebuild operation, which will take longer the more drives are contained in the array.