5 PERIPHERAL DEVICES

The CRAY EL series computer system accommodates customer requirements for peripherals by using a wide range of devices with diverse storage and communication capabilities. Refer to the OEM manuals for specific information about the equipment. Table 5-1 on the following page lists the peripheral storage devices and their characteristics.

Table 5-1. Peripheral Devices

Device Type	Manufacturer Model Number	Device Description	Storage Capacity	Transfer Rate (Peak)	Cray Research Name
Tape Drives	Archive Technology, Inc. Anaconda 2750	QIC 0.25-in. streaming cartridge †	1.3 Gbytes	600 Kbytes/s	DS-2
	EXABYTE Model 8500	8-mm helical-scan EXABYTE cartridge †	5 Gbytes	4 Mbytes/s	EX-2
	StorageTek Model 9914	0.5-in. 9-track 6 – 10.5 in. round tape	Encoding-dependent	6,250 bpi GCR mode	TD-2
	StorageTek Model 4220	18-track cartridge (IBM 3480-compatible)	260 Mbytes	6,250 bpi GCR mode	TD-3
	Hewlett-Packard C1533A DDS-2	Digital Audio Tape (DAT) drive	4.0 Gbytes on a 120-m tape	1 Mbyte/s with compression	DS-3
Disk Drives	Seagate Technology, Inc.	3.5-in. Winchester SCSI †	204 Mbytes formatted	3 Mbytes/s	HD-1
	MDB SYSTEMS, INC. DataShuttle 500	3.5-in. removable SCSI ‡	204 Mbytes formatted	3 Mbytes/s	DR-2
	Hitachi DK516-15	5.25-in. ESDI drives in 8-drive PE-3 tray	1.32 Gbytes formatted	2.75 Mbytes/s	DD-3
	MDB SYSTEMS, INC. DataShuttle 2100	5.25-in. removable ESDI with 2 drives per drawer ‡†	1.32 Gbytes formatted per drive	2.75 Mbytes/s per drive	DR-1
	Seagate Technology, Inc. ST83050K	8-in. 2-hd parallel IPI in ea. 2-drive PE-4 tray	2.7 Gbytes formatted	8.11 Mbytes/s	DD-4
	Hitachi DK516-15	5.25-in. ESDI array (10 disks per array) as controlled by the Maximum Strategy DAC	10.4 Gbytes per array	16.0 Mbytes/s per array	DAS-2
	Seagate Technology, Inc. Elite ST43401ND	5.25-in. SCSI-2 interface system disk drive	3.1 Gbytes formatted per drive	5.0 Mbytes/s	DD-5S
	Seagate Technology, Inc. ST43200K	5.25-in. IPI-2 interface system disk drive	3.4 Gbytes formatted per drive	12.4 Mbytes/s	DD-5I

† The drive connected to the Heurikon SCSI controller in IOS 0 is a dedicated system device. The tapes on this channel are standard equipment and can be used by the customer.

‡ The actual disk incorporated by MDB is the Seagate Technology, Inc. ST1239 Swift.

‡† The actual disk incorporated is the Hitachi DK516-15 ESDI drive.

All of the peripheral devices are mounted in the cabinets available. The input/output subsystem (IOS) that controls the peripheral device is normally included in the same cabinet. Figure 5-1 shows a typical layout of the various components within the system. The exact placement of any component is determined by the system configuration.

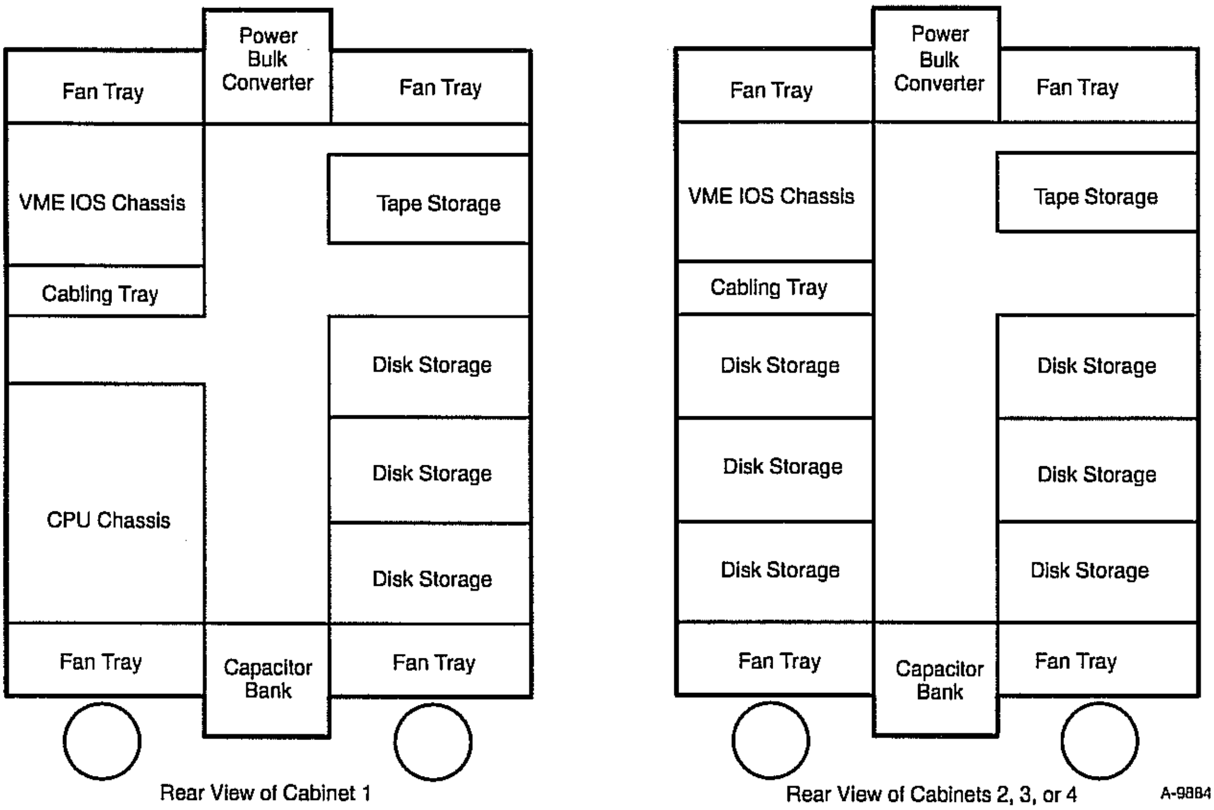

Figure 5-1. Example of Cabinet Layouts for the CRAY EL Series System

Most of the CRAY EL series system peripherals are field replaceable units (FRUs). This means that when one of the devices ceases to function, the entire device is replaced. There are no field-repairable parts inside the peripherals, except the disk array subsystem (DAS-2) controller. This component is described later in this section.

Naming Conventions

Whenever you read an error file, you must interpret the device name stored to determine which device has an error. The naming conventions used for the disk drives in the CRAY EL series system are shown in Figure 5-2. It is important to remember that each of the characteristics shown represents a 4-bit (hexadecimal) numbering system.

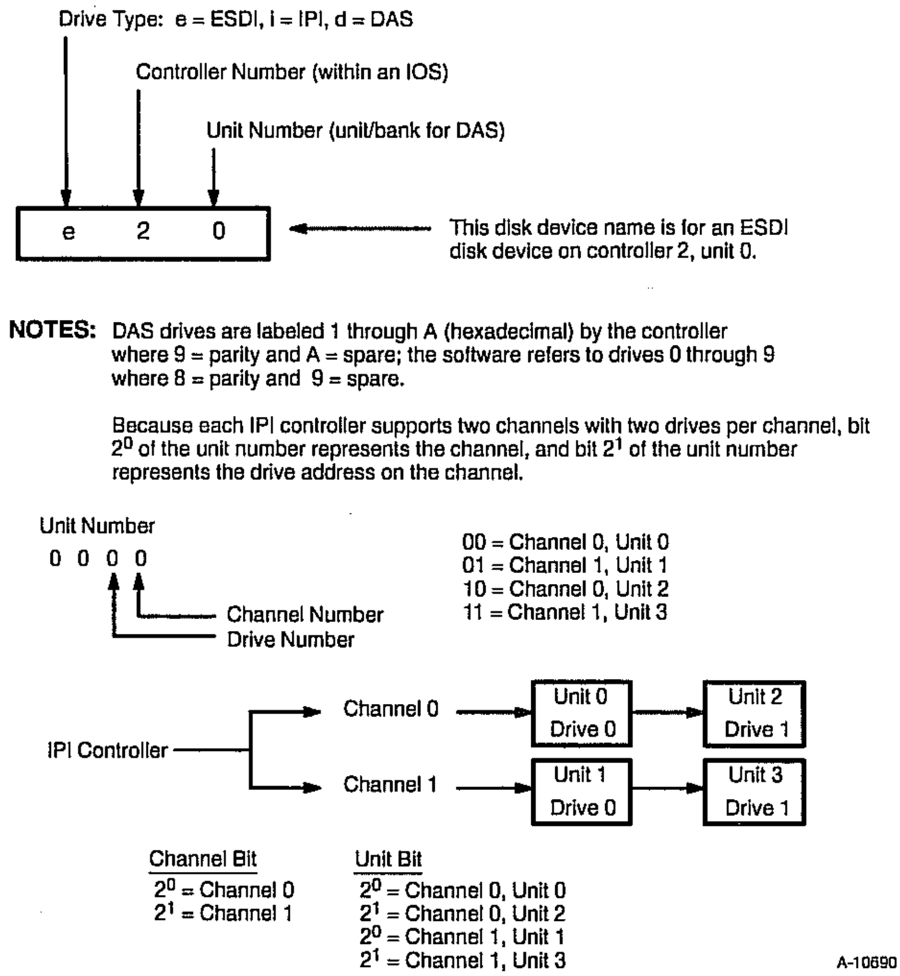

Figure 5-2. Disk Device Naming Conventions

The system tape drives also conform to a set of naming conventions. The naming conventions for the tape drives in the CRAY EL series system are shown in Figure 5-3.

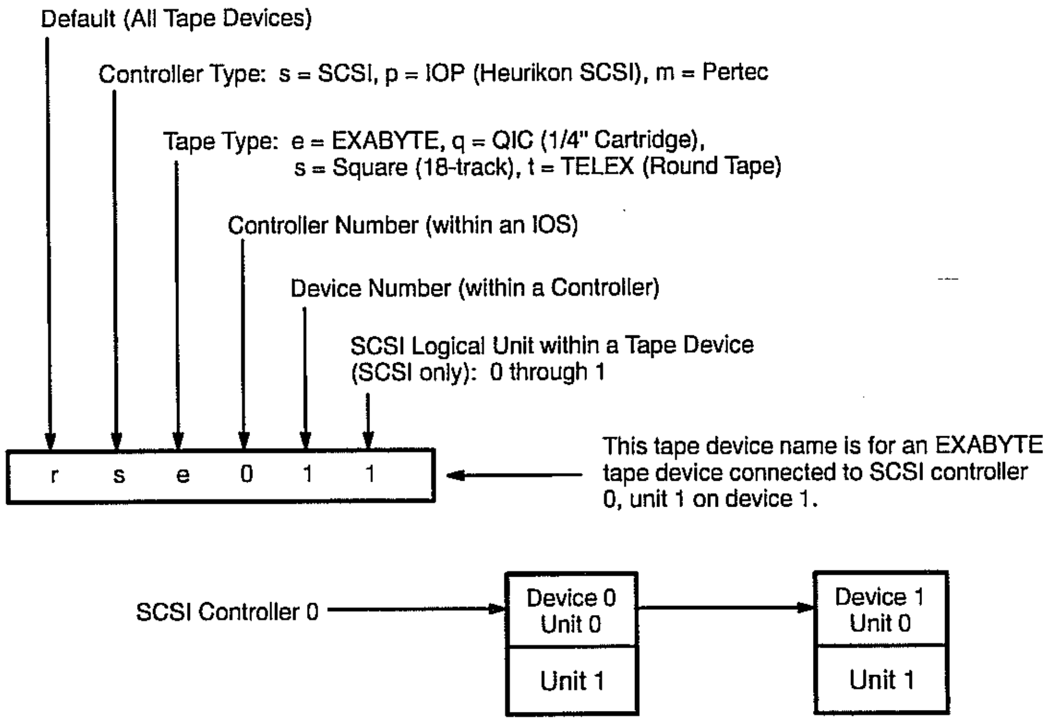

NOTES: The default label used for the TD-3 tape device is RSS040.
The right-most digit (SCSI logical unit number) is unused by the other tape devices and need not be specified; default = 0.

A-10691

Figure 5-3. Tape Device Naming Conventions

IOS Error Log

The error log is a binary data file kept on the IOS disk as `/adm/errlog`. The error log contains DAS and ESDI disk error messages. Each disk type has a different format, and the error information is largely the return status from the disk controllers for corrected and uncorrected errors. To diagnose the errors, the user must consult the disk controller manual from the vendor for explanations of error codes. The error log can be viewed by using the IOS command `errlog`. As with the `syslog` command, `errlog` can be deleted with the `rm` command. Typing **`errlog`** on an empty `/adm/errlog` file results in the following message being displayed.

```
IOS>errpt
errlog:  open failed on /adm/errlog:  file not
found
```

If the error log has a large number of errors, typing `errlog` causes them to scroll across the screen very rapidly. The output to the screen can be started and stopped using Control-s and Control-q (flow control) keyboard entries.

The ESDI controller provides a small amount of error information compared to the DAS because the increased complexity of the DAS requires additional information. Here is a sample error log:

```
IOS>errpt
Entry:  1 ESDI error status
Date:  21/03/91  11:23:27
chan:  0x4  unit:  0x0
func:  0x4  flags:  0x90
disk status :  0x0    error :  0x0
diskaddr:  0x2856f  bytcnt  :  0x20000

Entry:  2  DAS TCS (recovered error) status
Date:  18/04/91  11:55:59
cmd:  20  err code:  2
bank:  1 beg_block:  6488a  length:  400

Entry:  3 DAS ETE (unrecovered error) status
Date:  22/05/91  15:15:35
cmd:  40  err code:  bb
bank:  1   disk:  7    length:  200
block:  4545c  block count:  0

Entry:  4 DAS Disk Error Log detailed status
Date:  22/05/91  15:15:35
err code:  2  bank:  1  disk:  7  header:  1
cyl.  ea        head:  5                sector:  1
pos:  251       bitlen:1288
retries:    5  match: 0                block:
4545c
```

The ESDI entry format is the same for all error types possible on the ESDI controllers. The first fields are the date and time stamp of the error and the UNICOS logical channel and unit numbers, which are equivalent to the controller and disk number. The next field is the controller function, and the remaining fields are the controller return status codes. These are described in the *CIPRICO Rimfire 3400 Controller* manual, Ciprico Inc. publication number 21011905. The key field to decode is the flags field, which indicates whether the error type is a recovered or unrecovered error. The drive status field is decoded in the same manual, and the error codes are listed in Appendix C-1.

NOTE: With the software release IOS 9.3, disk errors are no longer logged on the IOS Winchester disk in `errlog`, but are reported to UNICOS and saved in `/usr/adm/errfile`. To aid analysis of problems when the system crashes during UNICOS startup, the IOS development group has added a new command to IOS 9.3 called `errpt`. The `errpt` command displays the last 64 disk errors detected by the IOS.

DAS entries are classified as three types: recovered error status, unrecovered error status, and disk error log status. Refer to the *Maximum Strategy Operation Manual for Strategy 2 Controller*, vendor number SOM 002. The `err code` field of the error log entry is the error correction (EC) and process identification (PID) fields of the `Transfer Command Success` return status. The EC bits are described in the OEM manual. The architecture of the disk array subsystem makes it impossible to determine the failing disk unit based on the recovered error status.

The DAS unrecovered error status is different than the recovered error status error code. These error codes are itemized in Appendix A of the DAS controller manual. The failing disk number is included in the error log entry.

The final entry type is the DAS disk error log. This is a detailed status that is returned following a DAS unrecovered error and is outlined in the DAS controller manual. The following examples are DAS error log entries.

Example 1: `IOS>reset`

```
IOS>iosdump /adm/iosdump1
IOS>load
IOS>dump -v p80 0
```

This process dumps mainframe memory to controller 8 (DAS), disk 0, starting sector 0.

Example 2:

```
IOS>reset
IOS>iosdump /adm/iosdump2
IOS>load
IOS>dump -v p20 .88288
```

This process dumps mainframe memory to controller 2 (ESDI), disk 0, starting sector 88228.

NOTE: The DAS device is not available with the CRAY EL98 system.

Further information on the `iosdump` and `dump` commands is in the *CRAY XMS Systems IOS Reference Manual,* CRI publication number SR-3085. For more information on the `idmp`, `cpdmp`, `crash`, and `/etc/coredd` commands, refer to the *UNICOS Administrator Commands Reference Manual,* CRI publication number SR-2022.

Disk Maintenance Information

CRAY EL series system flawmaps are maintained differently than flawmaps for the CRAY X-MP or CRAY Y-MP computer systems. First, it is assumed that in the CRAY EL series system, UNICOS is running on a flaw-free medium. This assumption is made to avoid maintaining flaws at the operating system level. Furthermore, the respective controllers for the ESDI, IPI, and DAS devices handle error correction at the IOS level.

UNICOS is run on the CRAY EL series system with the assumption that all disks have flaw-free media. This means that all flaw hiding, remapping, and data correction must be done at the IOS device-controller level. This also means that UNICOS does not maintain a UNICOS spares map. Because UNICOS does not recognize that flaws exist, all flaw management must be performed at the IOS kernel level. Online utilities such as `donut` and `ddms` do not apply to CRAY EL series devices.

Flaw Handling for DD-3, DD-4, and DD-5 Disks

Flaw hiding on DD-3 (ESDI), DD-4 (IPI), and DD-5 (IPI) drives is performed through a combination of sector slipping and track remapping. Each track contains a spare sector to accommodate slipping a defective sector. Spare tracks are set aside in the upper cylinders to accommodate track remapping in cases where slipping is not possible. If all spare tracks are utilized, the drive must be replaced (although this should seldom occur).

In all cases, user-added flaw information is maintained only in the ID fields of the affected sectors and is never recorded in a flaw table. The factory flaw table resides on the last cylinder of each drive and is, essentially, a read-only table placed there by the manufacturer. This table is not consulted during normal disk operations. Instead, the respective IOS controller references each sector header for any flaw information.

Each track contains one spare sector that is used to slip a defective sector and is not accessible to the user. A set of spare tracks, not available to the user, is located adjacent to the factory flaw table and is used to remap bad tracks. The IOS kernel commands used to perform DD-3, DD-4, and DD-5 flaw management include `dformat`, `dverify`, `dflawr`, and `dslip`.

Flaw Management Commands

The following list describes the characteristics of the flaw management utilities as they apply to the DD-3 (ESDI) and DD-4 (IPI) disk drives: For additional information, refer to the *Cray Research Entry Level Computer Systems IOS Reference Manual*, publication number SR-2408 2.2.

- `dformat` - This command formats the entire drive specified using the factory flaw table information. A limited range of cylinders cannot be specified. In addition to writing the ID field information, the data fields are also patterned to legitimize the ECC information for each sector.

  NOTE: Customer data is destroyed and all user-added flaws are lost when the `dformat` utility is used.

- `dverify` - This utility attempts to read all usable sectors on the specified drive and check for ID compare and/or data ECC errors. If an ECC error is detected, the utility attempts to correct the data and slip all affected sectors on the track. If slipping is not possible, the track and its data are remapped to a designated spare. If any ID error is detected, the track is always remapped to a designated spare. In all cases, the customer data is moved and/or corrected if possible. The `dverify` utility can be run on a periodic preventive maintenance schedule. Data must be backed up on the drives before running `dverify`, because data on the sector that is slipped or mapped out could be lost.

- `dslip` - The arguments for this utility must specify the drive type, controller number, and drive number (for example, E02) as well as the decimal physical block number of the sector to be tested. All of this information should be readily available on error reports. The utility attempts to read the specified sector up to 100 times while checking for errors. If errors are detected, the sector is slipped or the track is remapped. Customer data is corrected and moved whenever possible. If no errors are detected by `dslip`, the user is informed and prompted as to whether slipping/remapping should proceed.

- `dflawr` - This utility reads the factory flaw table from the specified drive and writes an ASCII file copy to /FLAW/IOS#/dev.flw on the IOS 0 maintenance SCSI drive. This command should be used whenever a new drive is installed.

- `dflaww` - This utility applies to all drives except IOS SCSI Winchester drives.

Disk Array Subsystem and DD-3 Block Assignment

The DD-3 (ESDI) and disk array subsystem (DAS) devices both use the same Hitachi DK516 spindle. Although the physical attributes of this spindle are identical in both applications, the format specification is incompatible between the DD-3 and the DAS. A DK516 spindle formatted for DD-3 applications cannot be used for DAS applications without being reformatted, and vice versa. The format specification and block allocation are significantly different for these two applications. The DK516 logical block assignments for DD-3 and DAS devices are as follows:

- DK516 connected to a CIPRICO (DC-3) controller as a DD-3
    - Locations 0 through 668399 are available to UNICOS
    - Locations 668400 through 668699 are available to diagnostics (6 cyl)
    - Locations 668700 through 669899 are reserved for spare tracks (4 cyl)
- DK516 on a Strategy 2 (DAC) controller as a DAS spindle
    - Locations 0 through 1269113 are available to UNICOS
    - Locations 1268114 through 1270051 are available to diagnostics (1 cylinder x 8 data spindles)
    - Locations 1270052 and greater are reserved by the controller

UNICOS uses these logical blocks when the disks are configured into the system. On each physical device description there is always an off_XX slice that starts at location 668400 for the DD-3 and location 1269114 for the DAS. It is required that all drives added to a machine have an off_XX slice with a correct starting sector address (for example, ESDI=170100, DAS=1269114) at the end of the physical device description in the configuration file. The offline disk diagnostics will destroy any data stored in the reserved diagnostics area on each disk drive. Note that the IOS level block descriptions for DD-3 drives are 2-Kbyte blocks, not the standard 4-Kbyte blocks.

NOTE: The DAS is not available in the CRAY EL98 system.

Disk Array Subsystem (DAS) Overview Information

NOTE: The DAS is not available in the CRAY EL98 system.

The ten drives in a DAS are controlled individually by the microprocessor (MPU) board within the disk array controller (DAC). All data recovery or correction is performed under the control of the MPU board. UNICOS treats the entire disk array as a single logical device and depends on the DAC to handle the individual spindles. For maintenance purposes, the IOS kernel can send maintenance commands to the DAC to reference individual disks within the array via maintenance commands.

Data striping is handled in the following manner. A sector write operation transfers 4,096 bytes of customer data to the array. The array buffers the data long enough to compensate for disk skewing and parity generation. The data is then transferred to the disks, bytes 0 through 511 go to the first disk, bytes 512 through 1023 go to the second disk, and so on. The byte parity generated by the DAC is written to the ninth disk. The data written to each of the nine disks is accompanied by ECC information.

DAS Flawmaps and Maintenance Commands

The DAS consists of 8 data spindles, 1 parity spindle, and 1 hot spare spindle. The DAS controller uses the first two physical cylinders as system tables. Two flaw tables are maintained by the DAS controller: the raw factory flaw table (RFT) and growth error table (GET). The RFT resides on the last cylinder of the DAS, and the GET resides on the first two cylinders of the DAS. These cylinders cannot be accessed by the system.

The flaw commands for the DAS are `dverify`, `dformat`, `ddisable`, `dreplace`, `dflawr`, `dflaww`, `dwconfig`, and `dinfo`. For additional information, refer to the *Cray Research Entry Level Computer Systems IOS Reference Manual*, publication number SR-2408 2.2.

NOTE: It may not be necessary to use some of the commands listed below if the IOS kernel support features are implemented on the DAS. Refer to "IOS Kernel Support Features for the DAS" in the following subsection.

- `dverify` performs a read-only sweep of the DAS by default. If any errors are detected, they are added to the GET. At this point the erroneous blocks are still accessible until the `dformat` (level 5) command is issued to add them as flaws.

- `dformat` is used to rewrite the ID fields on DAS devices, and employs several functional levels, as follows:
  - Level 2 commands the DAC to read the RFT, format physical cylinders 0 and 1 (the system area), and copy the RFT to cylinders 0 and 1.
  - Level 3 commands the DAC to read the current flaw table on cylinders 0 and 1 and use this information to format the entire user area of the specified spindle(s).
  - Level 5 commands the DAC to merge new entries in the GET with the RFT. The affected sectors are remapped and formatted, and data is corrected and moved whenever possible.
- `ddisable` disables the specified spindle and asks whether to substitute the hot spare. If the answer is yes, the hot spare is enabled and data from the bad spindle is reconstructed to the hot spare based upon information from the parity spindle and the remaining seven data spindles. At the conclusion of `ddisable` execution, the defective spindle may be physically replaced without loss of customer data.
- `dreplace` is used to reconstruct data from a currently active hot spare to a newly replaced spindle and to reassign the hot spare to a standby state.
- `dflawr` reads the RFT and writes it to an ASCII file on the IOS 0 SCSI drive.
- `dflaww` reads the ASCII backup file on the IOS 0 SCSI drive and writes it to the RFT on the disk array.
- `dwconfig` reads configuration parameters from a specified file and writes them to the selected controller. Refer to "IOS Kernel Support Features for the DAS" in the following subsection.
- `dinfo` displays the current state of each spindle in the DAS.

IOS Kernel Support Features for the DAS

The IOS release versions 8.9 and above have additional fault-resilient features for the DAS controller. One feature is automatic bad sector remapping (for example, autoflawing), which is done in the background during system operation. A UNICOS error log entry is made each time a media flaw is encountered.

Another feature is the automatic switch to the DAS hot spare drive. This feature is initiated during system operation when a drive exceeds a selectable threshold for unrecovered read or write errors. The thresholds can be changed to meet site preferences, but the release defaults to 10 unrecovered read errors, or 1 unrecovered write error per disk.

The threshold values and the individual running error counts for each drive in the array are stored in an unused area in the DAC controller memory, which is retained during system reboots and system power-down. This provides an accurate long-term record of disk errors.

The threshold values and error counts must be initialized when installing the IOS release. This initialization is done only once; it does not need to be performed again unless the user wishes to completely reinitialize the DAS controller.

The initialization commands are as follows:

- For a DAS connected to IOS 0:
    - `dwconfig d0 /install/das_init.cfg`
- For a DAS connected to an IOS other than IOS 0:
    - `rcmd ios# dwconfig d0 /install/das_init.cfg`

The commands to view or review the DAS threshold and error counts are as follows:

- For a DAS connected to IOS 0
    - `dstat d00`
- For a DAS connected to a different IOS:
    - `rcmd ios# dstat d00`

If any spindle exceeds the read- or write-error threshold and if the hot spare is available, then data from the failing spindle is written to the hot spare using parity reconstruction. This takes approximately 45 minutes. A UNICOS error log entry is made at the start and at the completion of reconstruction. An IOS syslog entry is also made of both events. This entire operation is done during UNICOS operation. If the reconstruct operation encounters uncorrectable errors on individual sectors, these disk blocks are reported to the IOS console and the UNICOS error log. In these situations the system administrator must restore the lost files from backups.

After a failing spindle is swapped and the hot spare is activated, the system administrator must contact maintenance personnel to schedule replacement of the defective spindle. At an appropriate time, the repaired/replaced drive can be reconstructed for use. Procedures for initializing a replacement drive for use are covered in a following subsection.

Changing Threshold Values

If a site does not wish to use hot-spare switching, then the threshold counts should be set to 0. A site may also choose to use a different value, which is limited to 255 maximum by the byte size.

An engineer may copy the file `/install/das_init.cfg` to a new file name, edit the new file using `ed` or `vi` directly on the IOS, and execute the appropriate `dwconfig` command to load the values from the file into the DAS. The file is a text file with comments that clarify which parameters to change.

Replacing a Failed Spindle within a DAS

When the DAS logs a disabled spindle, it is necessary to identify and replace the defective spindle. You must bring the operating system down before beginning the replacement procedure. The system needs to be down only long enough to physically install the replacement spindle.

The array spindles are located in the tray directly under the DAC. To determine the physical location of the failing unit, extend the disk tray. Each drive's data cable is labeled with its physical number. The parity spindle is physical number 9 and the spare spindle is 10.

CAUTION
Always install a replacement spindle in the same position as the failed spindle being removed. Repositioning of spindles within the DAS will result in loss and corruption of customer data.

When the replacement drive has been physically installed, the system can be brought up and returned to production. The spare spindle will still act as an active replacement for the replaced spindle. Reconstruction of the data to the newly replaced spindle must be done with the `dreplace` command, which can be performed in the background. Data reconstruction procedures are covered in the following subsection.

Reconstructing Data to a Replacement Spindle Using IOS Kernel Commands

This procedure applies only to DAS firmware level 1.2 or above. Perform this procedure after the faulty drive is physically replaced. Follow this procedure exactly! For additional information, refer to the *Cray Research Entry Level Computer Systems IOS Reference Manual*, publication number SR-2408 2.2.

CAUTION
Be certain that the disabled spindle, respective bank, and controller number are specified or the entire file system could be lost. The `dinfo` command can be used to confirm which spindle is currently disabled or substituted.

1. Establish communication with the IOS kernel from either the MWS-EL or the console terminal.

   Execute the `dformat` command with the appropriate parameters and specify format level 2.

   Example: **dformat d00 B010 2**

   This example initializes the system area for the 5th data spindle (logical spindle 4) in IOS 0, controller 0, bank 0. The individual spindle is specified by a hexadecimal bit map (B010). This procedure takes a few seconds to complete.

2. Repeat the `dformat` command with the appropriate parameters, but this time, specify format level 3 at the end of the command string.

   Example: **dformat d00 B010 3**

   Except for the last argument in the command string, all other values should be the same. This formats the entire user area of the spindle, which takes 20 to 30 minutes.

3. If no problems were encountered in formatting, execute the `dreplace` command using the appropriate controller, bank, and drive parameters.

   Example: **dreplace d00 4**

CAUTION

The spindle specifier for the `dreplace` command is NOT bit mapped. The spindle number must be entered as a decimal value with the first spindle number being 0 to properly define the replacement spindle.

The terminal screen displays a dot for each minute elapsed during the data reconstruction process. The `IOS>` prompt reappears when the reconstruction is complete. This process lasts approximately 20 to 30 minutes but can vary, depending on system activity.

Reconstructing Data to a Replacement Spindle Using the DAC COM Port

The following procedure can be used to initialize a replacement spindle into the DAS and reconstruct data to the new spindle. Use the previous IOS kernel procedure whenever possible. This procedure requires a terminal and appropriate connecting cable.

1. Connect the console terminal to the COM port of the DAS.
2. Set the terminal's baud rate to 9600.
3. Press `^c` to view the list of maintenance functions.
4. Enter **4** to select the `FORMAT` function.
5. Enter **1** to select the DAS bank to be formatted.
6. Enter **2** for the `FUNC` value (i.e. Format Mode 2).
7. Enter physical drive numbers for the `DSEL` (Drive SELect) value. This value is a bit-mapped representation of the replacement drive in hexadecimal format.

Examples:

3FF – all 8 data drives + 1 hot spare + 1 parity
200 – hot spare
100 – parity drive
80 – physical data drive #8
40 – physical data drive #7

8. The format takes only a few seconds because only the system area of the drive is formatted and initialized (cylinders 0 and 1).

9. Press `^c` to return to the list of maintenance functions.

10. Enter `C`; you will see a page of drive data.

11. Press ↵ ; all drives should be mounted and each disk unit should have a status value of 30.

12. Press `^c` to return to the list of maintenance functions.

13. Enter **4** to select `FORMAT` option.

14. Enter `3` to discard GET and format the user area (i.e. Format Mode 3).

15. Enter physical drive numbers for the `DSEL` value. (Enter the same value you used in Step 1). This takes about 20 minutes.

16. Press `^c` to return to the list of maintenance functions.

17. Enter `F` to select the `Scrub` option. (Do not select `Scrub Utility`).

18. Select data option **F** to run one pass of random data on the drive. This takes 20 to 30 minutes.

19. Press `^c` to return to the list of maintenance functions.

20. Enter **4** to select `FORMAT` option.

21. Enter `5` to merge new flaws into the RFT and to format the affected areas.

22. Enter physical drive numbers for the `DSEL` value. (Enter the same value used in Steps 1 and 3).

23. Disconnect the console terminal from the DAS COM port and reconnect to the console port.

24. Set the baud rate back to 19,200 on the console terminal.

25. At the `BOOT>` prompt, enter **`load`** to load the IOS kernel.

26. Enter **`cd/flaw`**; there should be a set of XXXX.DAF files, one for each drive.

27. Enter **`dflawr d00 <logicaldrive#> <filename>`** to rewrite a factory flaw table on the IOS SCSI drive.

NOTE: `<logicaldrive#>` is 0 through 7 for the data drive, 8 for the parity drive, or 9 for the hot spare. `<filename>` is D1 through D8 and corresponds to the flaw table for the data drive; use DH for the hot spare, or DP for the parity drive.

28. Enter **`dreplace d00 <logicaldrive#>`**. (Use the same drive value used previously in Step 7). This command rebuilds the data on the new disk and releases the hot spare. This takes about 30 minutes.

29. Enter **`dinfo`**; all status readings should be 30. Nothing should be disabled or substituted in the status.

30. The system is now ready for the operating system to be booted.

TD-2 Tape Drive

The TD-2 tape drive (TD-2) is a 9-track tape drive that handles open-reel 0.5-in. tape. Any reel between 6 in. and 10.5 in. fits on the unit. The TD-2 is manufactured by Storage Technology Corporation as model 9914. It is designed to handle four types of data encoding:

- NRZI - nonreturn to zero indicated (800 bpi)
- PE - phase encoded (1,600 bpi)
- DPE - double phase encoded (3,200 bpi)
- GCR - group coded recording (6,250 bpi)

The TD-2 is controlled by a pertec cache interface controller card, which is supplied by Ciprico Inc. as model TM3000. This pertec interface connects to another pertec cache contained in the TD-2, and it is a field-replaceable unit (FRU). Other FRUs in the TD-2 are:

- Data control board (DCB)
- Digital data processor board (DDP)
- Analog data processor board (ADP)
- Servo control board (SCB)
- Power supply board (PSB) that contains two fuses
- Hub sensor board (HSB)
- In-chute sensor board

Any other parts that can malfunction are too complex for field repair, because they require an oscilloscope or other test equipment. In this case, the entire TD-2 must be removed and replaced.

TD-2 Operator Functions

The customer operator can perform the following tasks by using the switches located at the front of the TD-2:

- Switch power on or off.

- Select the recording density, which determines the data encoding.

- Place the TD-2 in online mode. The UNICOS operating system controls the TD-2 when it is online.

- Place the TD-2 in offline mode.

- Rewind the tape to beginning of tape (BOT).

- Unload the tape.
- Clean the tape read/write heads and the tape path.

Refer to Figure 5-4 for a diagram of the front (operator's) panel. The shaded areas of the panel indicate membrane-type switches; the other items represent lighted displays. The area showing the legend OFFLINE is a liquid-crystal display capable of displaying a large variety of messages. The other displays shown are back-lighted, single-message displays that are illuminated when the particular condition exists in the TD-2. For a more detailed explanation of these controls and displays, refer to the manufacturer's manual for the TD-2.

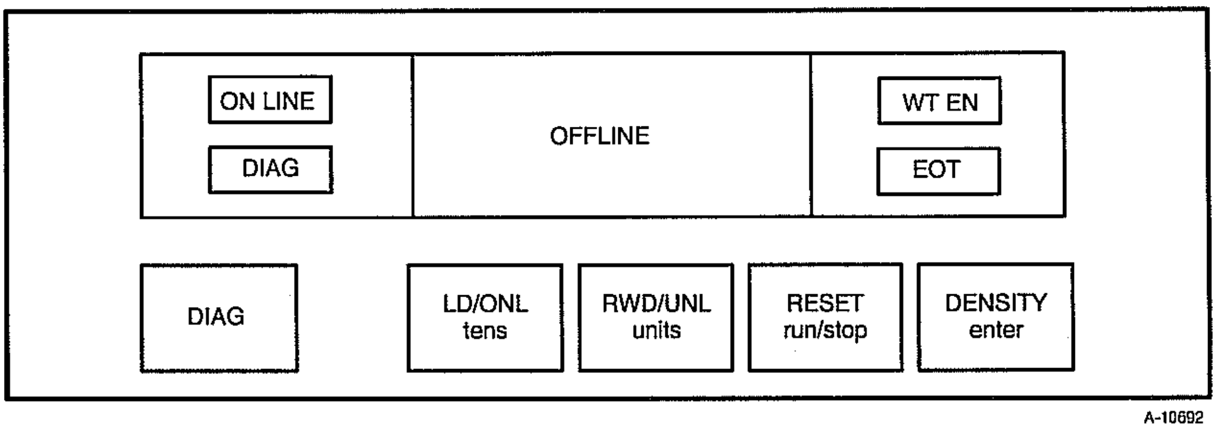

Figure 5-4. Operator Panel

When power is applied to the TD-2, a set of internal diagnostics is run. These diagnostics are referred to as power-on/reset (POR) diagnostics and reside in non-volatile random access memory (NVRAM). While these diagnostics are running, the liquid-crystal display (operator's display) displays the message TESTING. When the POR diagnostics successfully complete, the operator is prompted to load a tape into the tape loading chamber and begin operation.

If there is a malfunction of the POR diagnostics, the operator display shows a fault message. The messages and possible causes for failure that may be displayed are shown in Table 5-2.

Table 5-2. POR Messages

Message	Possible Cause
Blank display	Main power is not available, the +5-volt supply is not operating, or the servo control bus is faulty.
CON VAL x	This usually indicates a missing or disconnected subassembly (x is a number from 1 through 7).
OK	The POR checks have completed successfully.
POWER	A fault has been found with the power supply board while all the internal power supply lines were being checked for normal limits.
TESTING	This power-up message indicates that the POR checks are in progress.

The following process occurs when power is applied to the TD-2:

1. The operator panel indicates TESTING. This means the POR diagnostics are running.

2. The operator panel momentarily indicates OK. This means the POR diagnostics have successfully completed.

3. The operator panel displays LOCATING. While this display is illuminated, the supply hub is rotating back and forth for several seconds.

4. The operator display indicates NO TAPE. This indicates completion of the POR sequence and prompts the operator to begin normal operation.

If you want to load a tape at this time, you are prompted to press the RESET on the operator panel. It is important to note that the tape loading door cannot be opened manually; there is no manual latch release mechanism. The only way that the loading door can be opened is by releasing the solenoid that locks the door during operation. If it becomes necessary to manually open the tape loading door as a fault/maintenance procedure, you must extend the TD-2 from the mainframe chassis. Next, open the tape path cover by releasing the three latches holding the tape path cover closed. Then, using a pointed instrument, manually push the closing latch away from its locked position, thus opening the tape loading door.

Once the tape loading door is opened, it is possible to load or unload a tape. The TD-2 is capable of containing tapes ranging from a 6-in. diameter to a 10.5-in. diameter. To unload a tape, simply grasp the tape and remove it. To load a tape, place the tape into the tape chamber and close the tape loading door. At this time, the TD-2 takes over the load operation.

The TD-2 first causes the supply hub to rotate back and forth, moving the tape reel into position over the hub. Once in position, the hub locks the reel into place, turns on the vacuum motor, and begins rotating in a counter-clockwise (rewind) direction. Thus, the tape is loosened on the reel, pulled into the tape chute by the vacuum, then rewound by the supply reel. This action can occur two or three times, and ensures that the tape reel is inserted right side up, which is determined by an optical sensor that "sees" the tape (being pulled by the vacuum), then does not "see" the tape (being rewound). The result is that the supply motor reverses and starts unwinding tape. The tape is pulled through the tape path by the vacuum and wraps itself around the take-up reel. Once the take-up reel has secured the tape, the tape transport mechanism takes over and performs several functions.

If the load operation begins to unwind tape during the initial sensing cycle (that is, if the optical sensor does not see, then sees, the tape), the TD-2 determines that the tape reel is inverted. In this case, the operator display will display the message REEL INV, the supply hub motor rewinds the tape by doing a fast-forward operation, the supply hub unlocks the reel, and the operator panel begins indicating that RESET must be pushed to remove and invert the reel.

When the tape load operation is completed correctly, the TD-2 will move a small portion of the tape forward and reverse across the read head rapidly several times. This is referred to as a "shoe-shine" operation. During this operation, the hardware is doing a read operation on actual data written on the tape. Then, firmware compares the written data with imbedded samples and determines the density that was used to write the tape. During the comparison period, the operator display contains the message ANALYZE. When the analysis is finished, the display shows the density, and the firmware is set to decode that tape. The TD-2 now begins a period of mechanical motion, searching forward and backward, until it locates the beginning of the tape (BOT). When this operation completes, the message BOT is displayed, alternating with density. This completes the load sequence. During this sequence, several messages are displayed. If the sequence fails, several other messages can be displayed. These messages are detailed in Table 5-3.

Table 5-3. TD-2 Operator Display Messages

Message	Indication
ANALYZE	The tape reached BOT and the data circuits are reading the ident burst to determine its recorded density.
BOT	Tape is at BOT and loading is complete. When set online, the display changes to LD POINT and the TD-2 is able to respond to host commands.
DOOR	The loading door (or the tape path cover) is open after a loading sequence started.
HUB LOCK	Shortly after loading was initiated, the reel was clamped onto the supply hub.
HUB ERR	Incorrect seating of the supply reel on the supply hub.
IN LIMIT	The tension arm reached the limit if its travel and tape tension has been lost.
LID OPEN	Loading has been initiated, but the tape path cover latches are not secured.
LRG REEL	Near the end of the loading sequence, the firmware detected that a large (10.5-in.) reel has been used.
LOADING	The reel was clamped and the tape is being threaded along the tape path. **NOTE:** If LD/ONL is pressed during loading, the ONLINE legend illuminates immediately; at BOT, the TD-2 will be under the control of the host.
LOCATING	Immediately after loading has been initiated, the reel is being maneuvered so that it lies flat on the supply hub, prior to clamping.
LOCKING	After loading has been initiated, the reel is clamped onto the supply hub.
MED REEL	Near the end of the loading sequence, the firmware has detected that a medium (8.5-in.) reel is being used.
NIC	Not in chute. During the early stages of loading, tape was not detected in the tape path chute.
NO TAPE	After loading has been initiated, the reel locating process concludes that no tape is present on the supply hub.
NTU	No take-up. During the later stages of loading, tape was not gripped onto the take-up reel.
OK	The POR checks were successfully completed.
REEL INV	During the early stages of loading, the tape reel was found to be inserted with the write-protect ring up.
RESET	During the loading sequence, the RESET button was sensed as permanently depressed. Engineering assistance is required.
**TAB	The BOT tab was not detected near the end of the loading sequence. This tab must be present and within ANSI/ECMA specifications from the physical end of the tape.
REWIND	The tape is rewinding following either a host or operator rewind command. This sequence completes at BOT. If RESET is pressed during rewinding, tape motion stops and the REWIND indication is replaced by an OFFLINE indication.
REW/UNLD	An operator rewind and unload command has been given by simultaneously depressing the RESET and RWD/UNL buttons.

Table 5-3. TD-2 Operator Display Messages (continued)

Message	Indication
SML REEL	Near the end of the loading sequence, the firmware has detected that a small (6- or 7-in.) reel is being used.
TAPE NOT IN CHUTE	Scrolled message: the tape end has not entered the tape path.
UNLOAD	An unload sequence is in progress following the pressing of the RWD/UNL button with BOT indicated. When the unload sequence is complete, the display changes to READY.
UNLOCK	The reel is being unclamped. **NOTE:** UNLOAD followed by UNLOCK could be the result of a failure to detect the BOT marker.

At the product specialist level, there are no preventive maintenance measures required. However, periodic maintenance is required by the operator. This maintenance consists primarily of cleaning the various components in the tape path. It is possible for operator maintenance to cause or reveal a fault condition. Therefore, it is good practice to be aware of the required maintenance procedures, which are in the OEM manuals available with the CRAY EL series system.

TD-3 Tape Drive

The TD-3 tape drive (TD-3) is also supplied by Storage Technology Corporation as the model 4220 cartridge tape subsystem. The TD-3 is an 18-track tape drive that is IBM 3480-compatible.

The TD-3 is used in several subsystems. These subsystems include:

- Tape transport
- Servo electronics
- Disk drive electronics
- Read/write electronics
- Pneumatics
- Operator control panel
- DC power supply

The disk drive subsystem provides the TD-3 with both functional and diagnostic microcode. The floppy disk drive uses a 3.5-in., 720-Kbyte format, and it is controlled by a single integrated circuit mounted on the system board (SB). Generated signals are sent to or from the drive through the bottom card (BT).

As is true in any electromechanical device, three types of failure can occur in the TD-3: power-level electrical, mechanical, and signal-level electronic.

Power circuitry for the TD-3 is designed to accommodate either 100 to 120 Vac or 200 to 240 Vac; these ranges are switch selectable from 47 to 63 Hz, single phase. From this input, the TD-3 supplies + 5 Vdc, +/- 12 Vdc, and + 24 Vdc to the various subassemblies. The only location in the TD-3 where AC power voltages exist is at the DC power supply board (DCPS). All voltages internal to the TD-3 are DC voltages.

The mechanical processes of the TD-3 are contained primarily on the tape deck. These mechanical processes deal with loading, unloading, and moving the tape after the cartridge is inserted into the deck.

The TD-3 tape drive is considered a field-replaceable device; no internal repairs are possible.

Figure 5-5 shows the basic TD-3 menus available with a SCSI interface.

```
Offline Menus

[OF:DIAG:      *]
[+SUBSYS STATUS]
      [*Online Request]
      [*Exit      ]
[+SET OR DISPLAY]
      [**CONFIGURATION]
            [> > Set Diags:None]
            [> > Diag EC LVL:0]
            [> > Set Host: SCSI]
            [> > Lang: English]
            [> > Set Part: Auto]
            [> > Set Sync Tm: 2]
            [> > Passkey: Enab]
            [> > Sync Nego:Enab]
            [> > Set BRN:    00]
            [> Exit       ]
      [*Retension Tape]
      [*Write Tape Mark]
      [**Set Scsi ID: 0]
      [**Set Lun:    0]
      [**Set Clean:  16K]
      [*Exit      ]
[+DIAGNOSTICS    ]
      [*Diag: Test 1  ]
      [*Diag: Test 2  ]
      [*Diag: Test 3  ]
      [*Diag: Test 4  ]
      [*Diag:Load/Unld]
      [**Diag: Loop 1]
      [Diag: Test Disp]
      [*Loader Sensors]
      [*Disply Pressure]
      [*Disply Tension]
      [*Exit      ]
[+SAVE TRACE  ]
      [*Save Trc: All]
      [*Exit      ]
```

```
Online Menus

[             *]
[ON:IDLE:      *]
[+SUBSYS STATUS]
      [*Offline Request]
      [*Exit      ]
```

Figure 5-5. Basic TD-3 Menus with a SCSI Interface

Disk Array Controller

The only non-FRU peripheral device in the CRAY EL series system is the disk array controller (DAC) unit supplied by Maximum Strategy, Inc. The DAC is an intelligent disk controller capable of storing and retrieving data from one to four banks of eight standard ESDI disk drives in parallel. The DAC also controls an additional drive for each bank of eight that is used to store a parity bit for fault protection as well as an operational standby (hot spare) drive used to replace a faulty drive from the bank.

The drives supported by the DAC are the DD-3, 5.25-in. ESDI serial data drives, which are capable of transferring data at a rate of 16 Mbytes/s in this configuration.

Each of the ten disk drives in the bank is connected to a disk interface and data buffer PC board in the DAC, which are FRUs. The DAC also contains a parity PC board that is separate from the parity drive disk interface PC board, and a CPU PC board. The CPU contains a high-speed interface (HSI) unit. The DAC is designed to interface with up to four banks of ESDI drives. If more than one bank of drives is to be controlled, an additional HSI PC board must be installed for each bank. A multiplexer subassembly must also be installed to link the four banks to the one controller.

Figure 5-6 indicates which components must be removed to access the PC boards, and Figure 5-7 shows the location of each of the PC boards within the DAC.

The DAC has a small cut-out on the front panel (refer to Figure 5-6) in which the READY LED, a RESET button, and RS-422 serial communication (COM) port are visible. The READY LED, located on the CPU card in slot 12, is a visual indicator of the state of the DAC. When the READY LED is glowing solid red, it indicates that the DAC is offline. During a normal power-on sequence or a reset sequence, the LED glows red as the disk drives spin up. The normal time for this sequence should not exceed 120 seconds. If the LED has not turned green after 120 seconds, it indicates a problem with the DAC or with one of the disk drives. Note in Figure 5-6 that each of the disk interface PC boards has an LED visible through the front shield. As the disk drive associated with the interface board spins up, this LED glows red. When the disk drive is up to speed and when the internal diagnostics for the disk interface PC board have been run, the LED turns green.

If the CPU READY LED is red, no communication is possible with the DAC. If any one of the data buffer/disk interface LEDs is red, either the disk drive failed to reach speed or the internal diagnostics for that interface board failed. In either case, the next step is to initiate disk drive diagnostics.

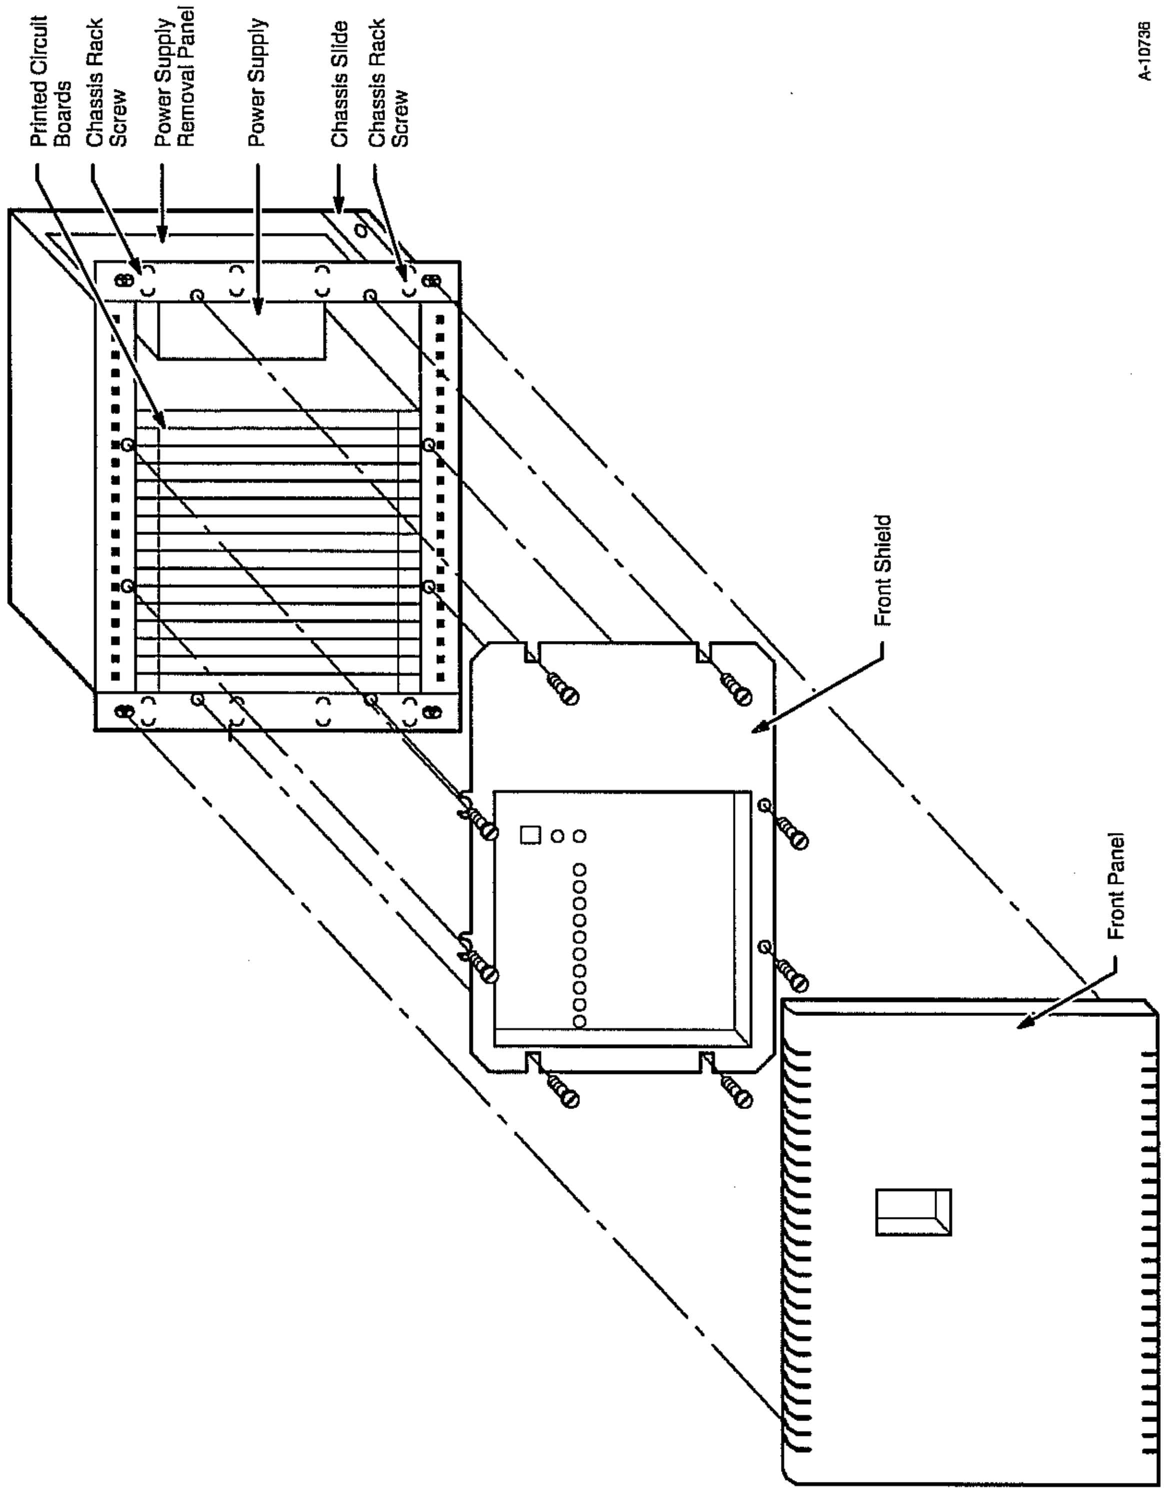

Figure 5-6. DAC (Exploded View)

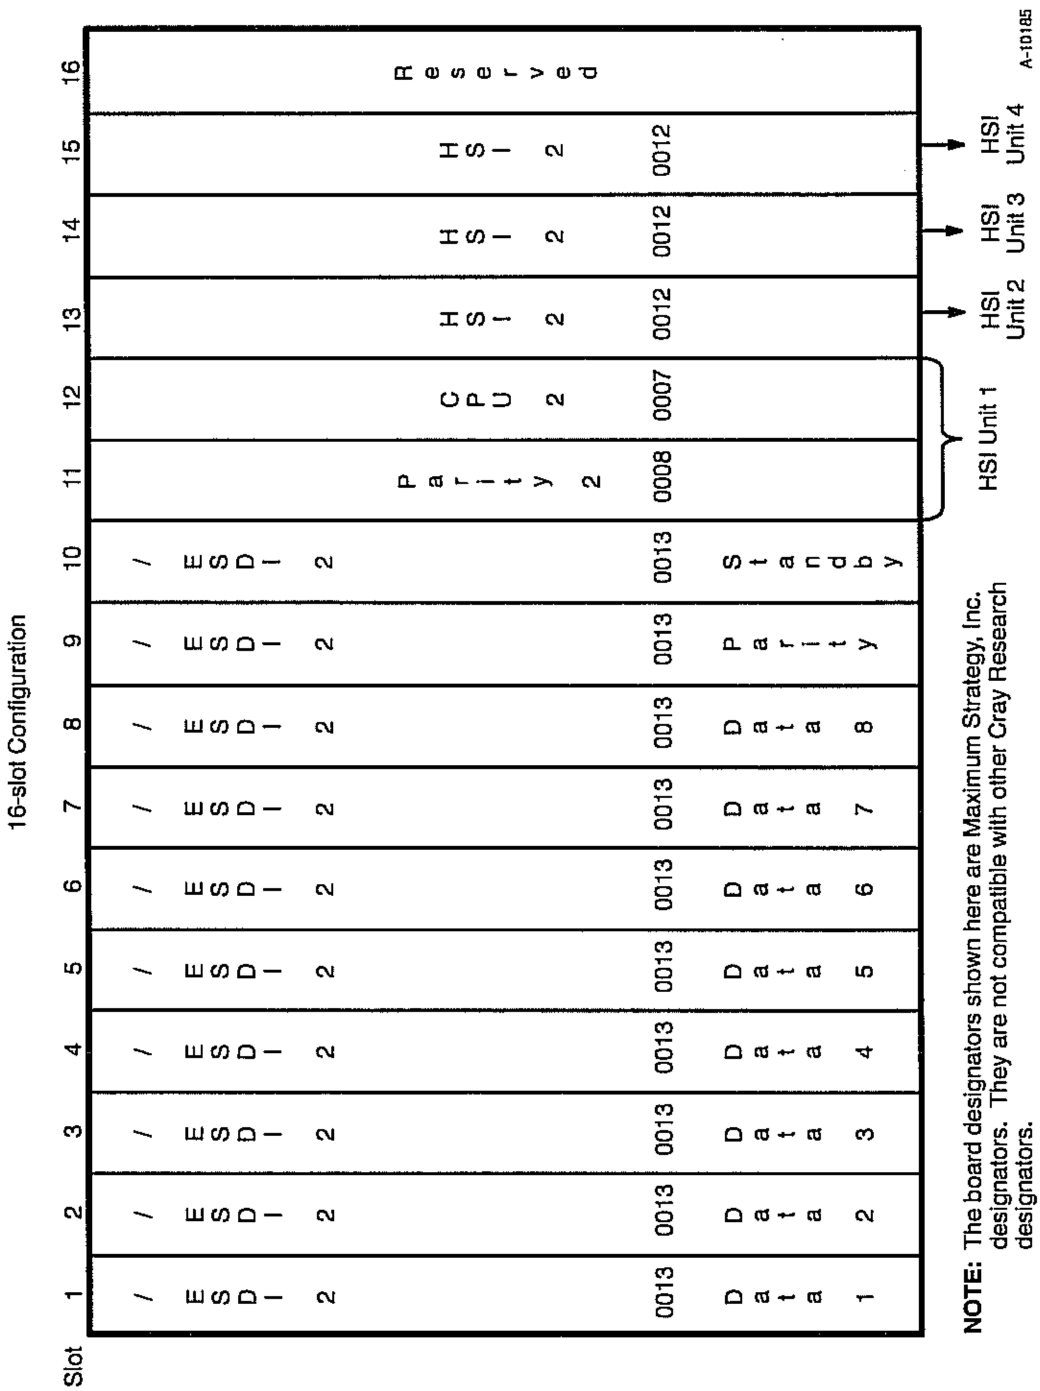

NOTE: The board designators shown here are Maximum Strategy, Inc. designators. They are not compatible with other Cray Research designators.

Figure 5-7. DAC Printed Circuit Board Locations

During normal operation, do not use the RESET button. Initiate the reset function using IOS software. During a power-up sequence, however, you may press the RESET button once if the READY LED fails to turn green within 120 seconds.

The COM port located on the front of the DAC is a serial RS-232 connection that can be used for offline diagnostics or for online system configuration. These functions require connecting a terminal to the COM port found on the front of the DAS in order to invoke any of the offline commands to operate the DAS or any of its drives. These commands are described in the *CRAY EL Series System IOS Reference Manual.* If you need more information regarding the use of the COM port, refer to the Maximum Strategy, Inc. documentation supplied with the system, or contact Hardware Product Support at Cray Research, Inc.

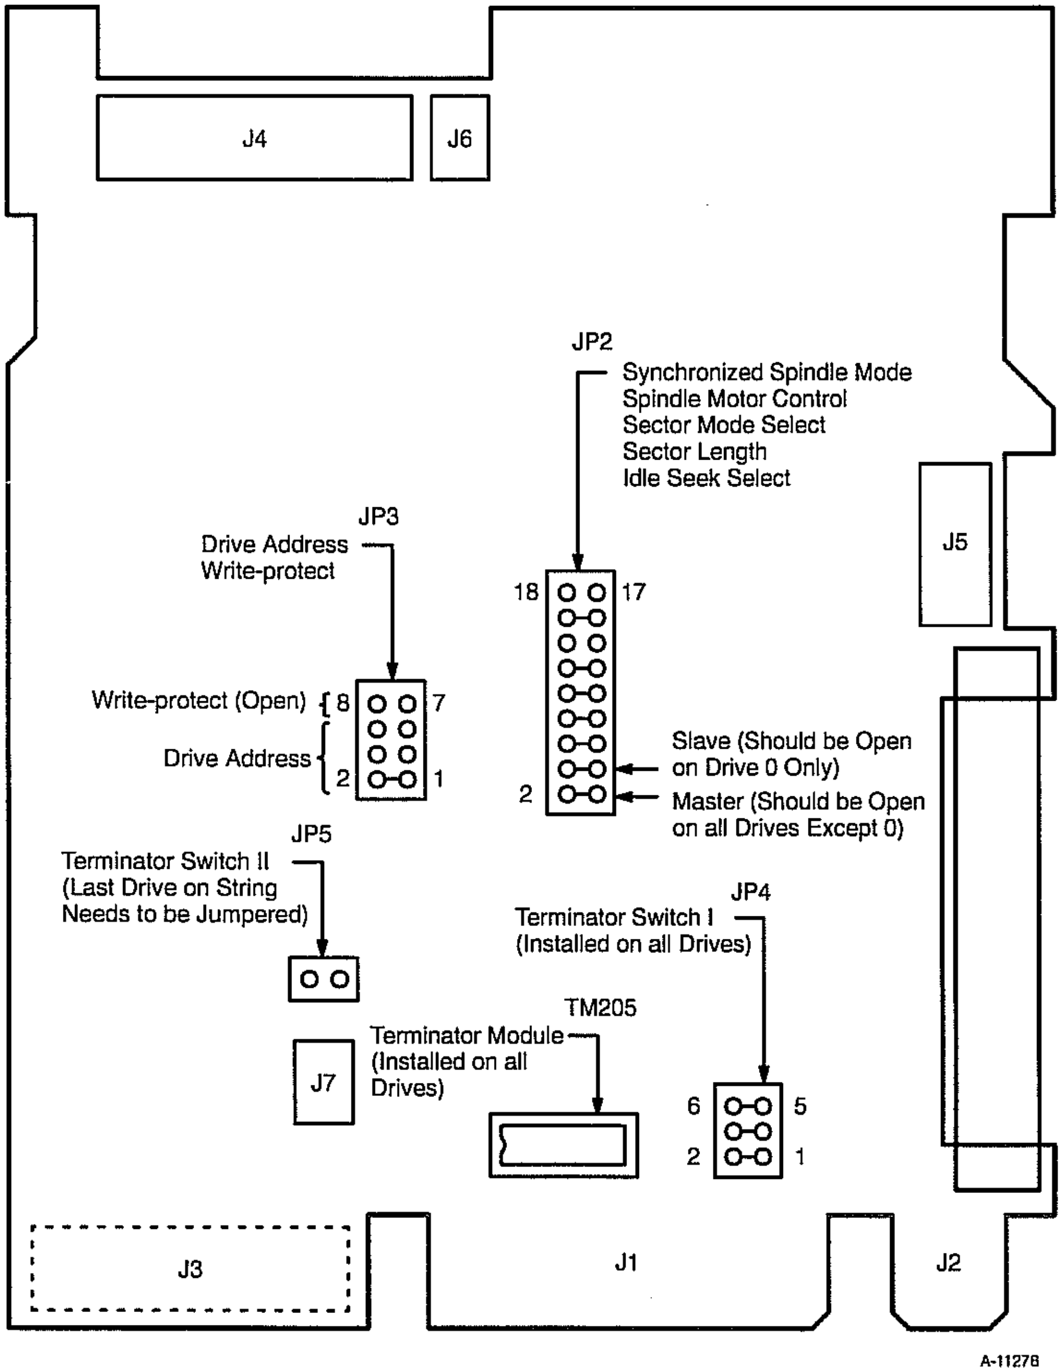

Figure 5-8. ESDI DK516 DAS SZ963 PCB Layout

DataShuttle 2100 DR-1 Removable Disk Drive

The following characteristics describe the DataShuttle 2100 disk enclosure (DR-1).

- Manufactured by MDB Systems, Inc.
- Contains up to two 5.25-in. disk drives
- Currently uses two DD-3 drives
- Contains internal power supply and cooling fans
- Can be free standing or rack mounted
- Green LED, when illuminated, indicates no DC power
  - Safe to remove drive canister
- Yellow LED, when illuminated, indicates DC power is present
  - Do not remove drive canister when yellow LED is blinking. The drive is spinning down.

Figure 5-9 is an overview of the DataShuttle disk enclosure; Figure 5-10 shows the DataShuttle front panel.

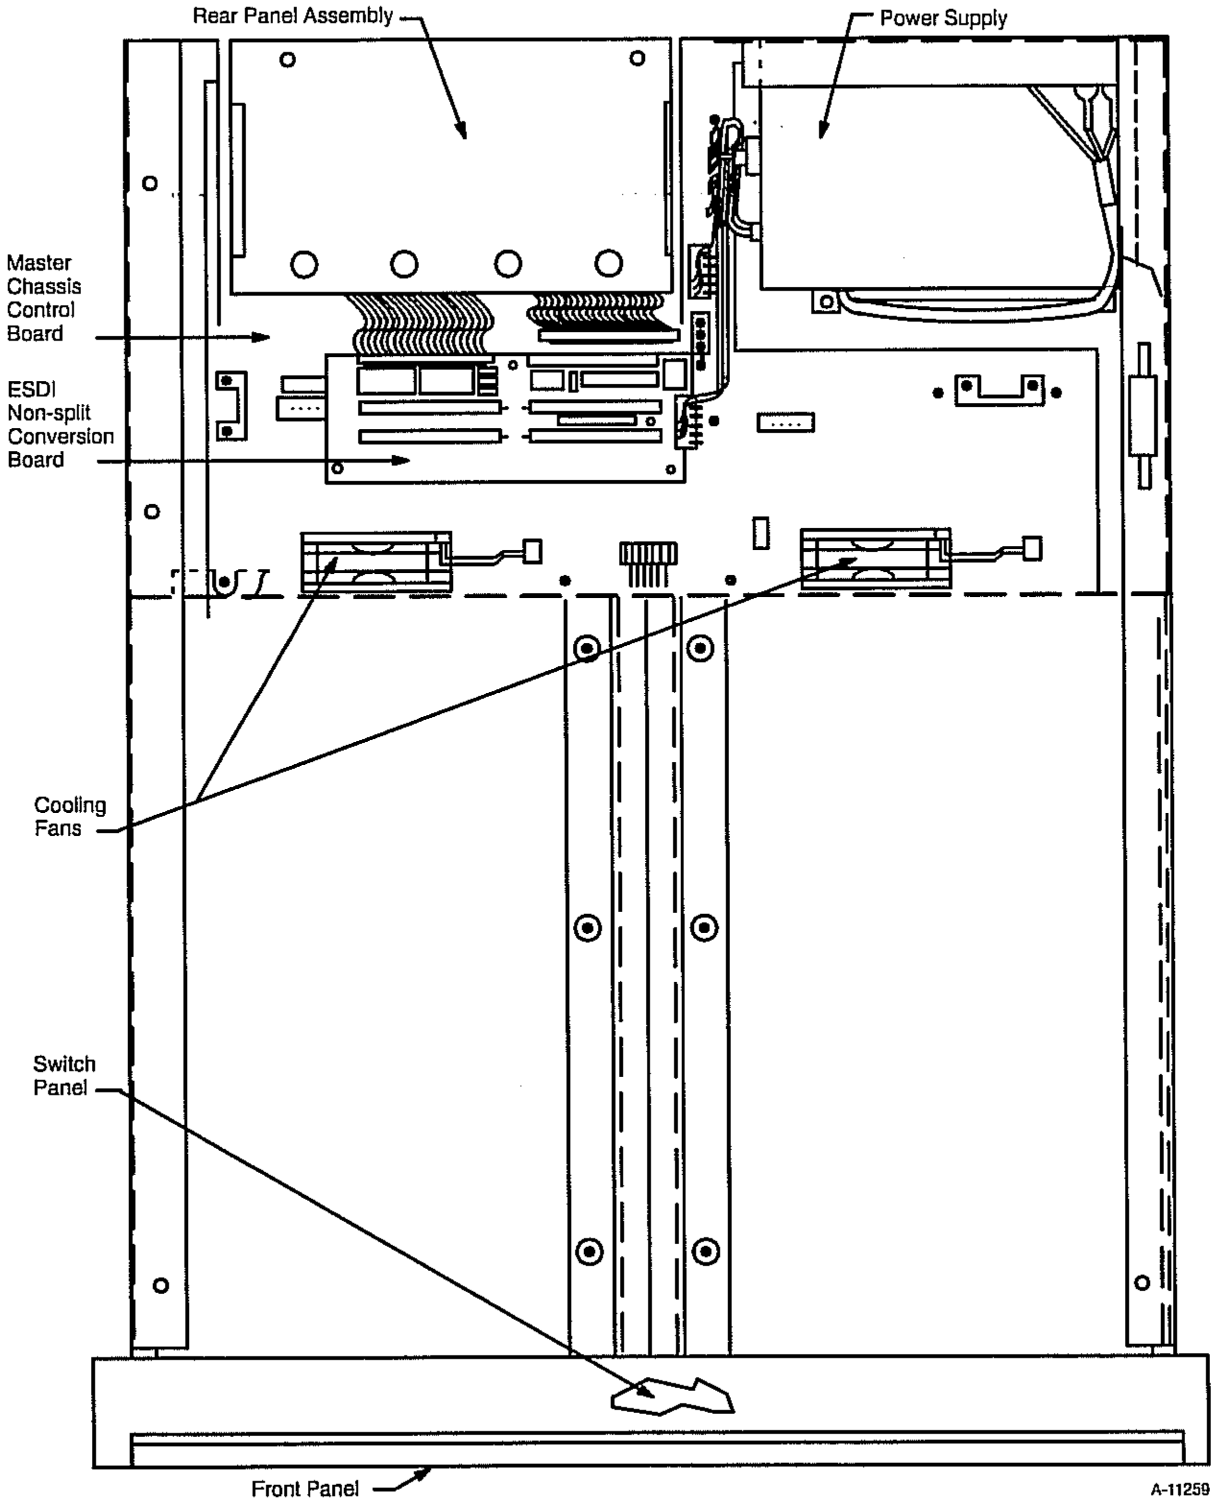

Figure 5-9. DataShuttle Chassis (Top View)

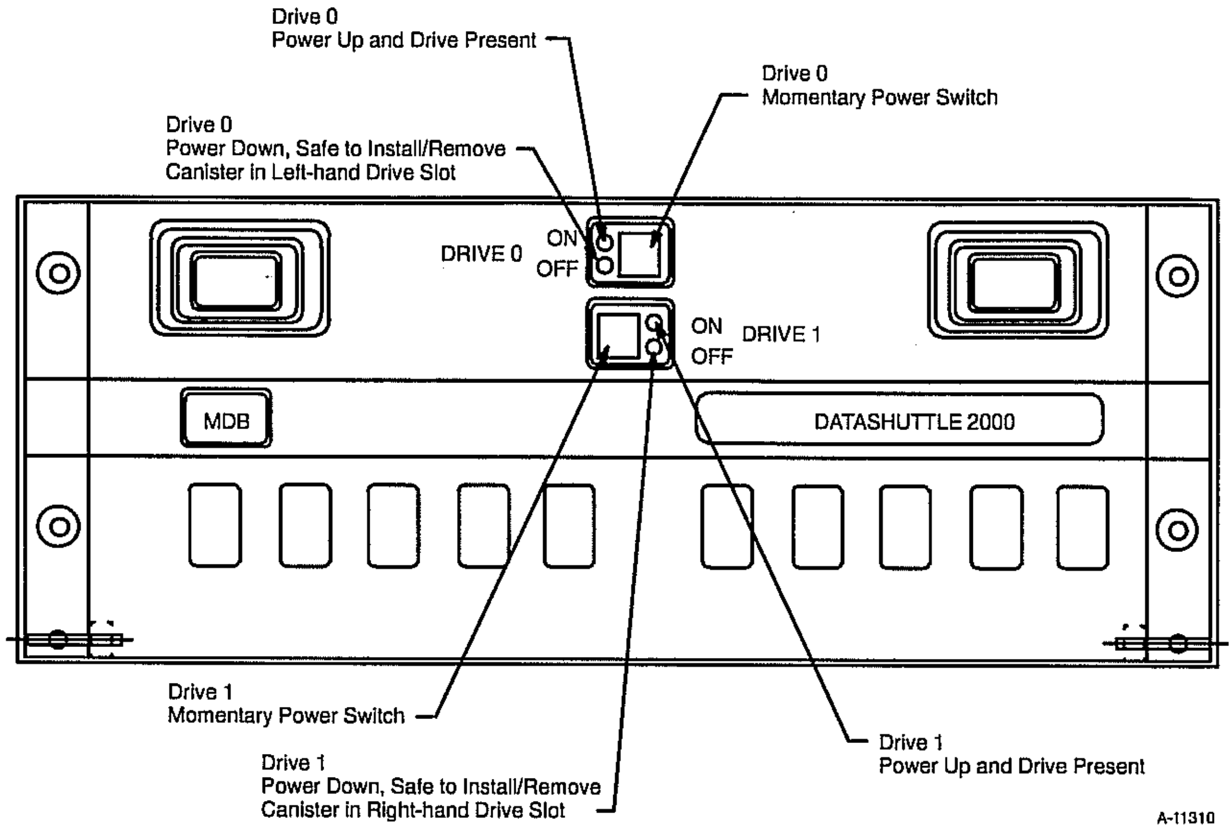

Figure 5-10. DataShuttle Front Panel Assembly

Rear Panel Assembly

The rear panel of the DataShuttle chassis contains the cooling fan exhaust openings, the main power switch, the AC line connector that accommodates the AC line cord, and the primary voltage select switch, which may be set to either 115 Vac or 230 Vac nominal AC line operation. In the CRAY EL series system, this selector switch should be set to 230 Vac.

The power ON/OFF switch that provides input line voltage is located on the left side of the chassis as viewed from the rear.

The input/output connectors located on the recessed rear panel provide the host interface connections as well as provisions to expand the subsystem (daisy chain).

Connector locations J1 through J4 on the back of the recessed rear panel are for future use and are covered with blank panels. J5 and J6 are female 50-pin subminiature D-connectors located on the bottom of the recessed rear panel assembly. Refer to Figure 5-11.

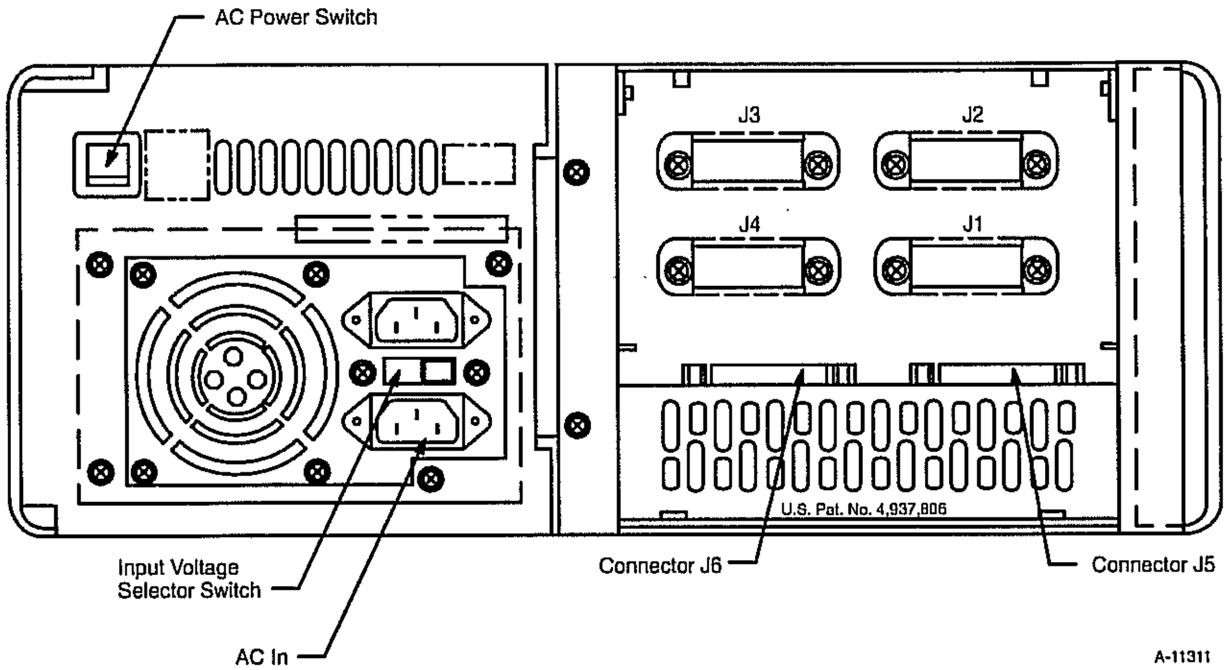

Figure 5-11. DataShuttle Rear Panel Assembly

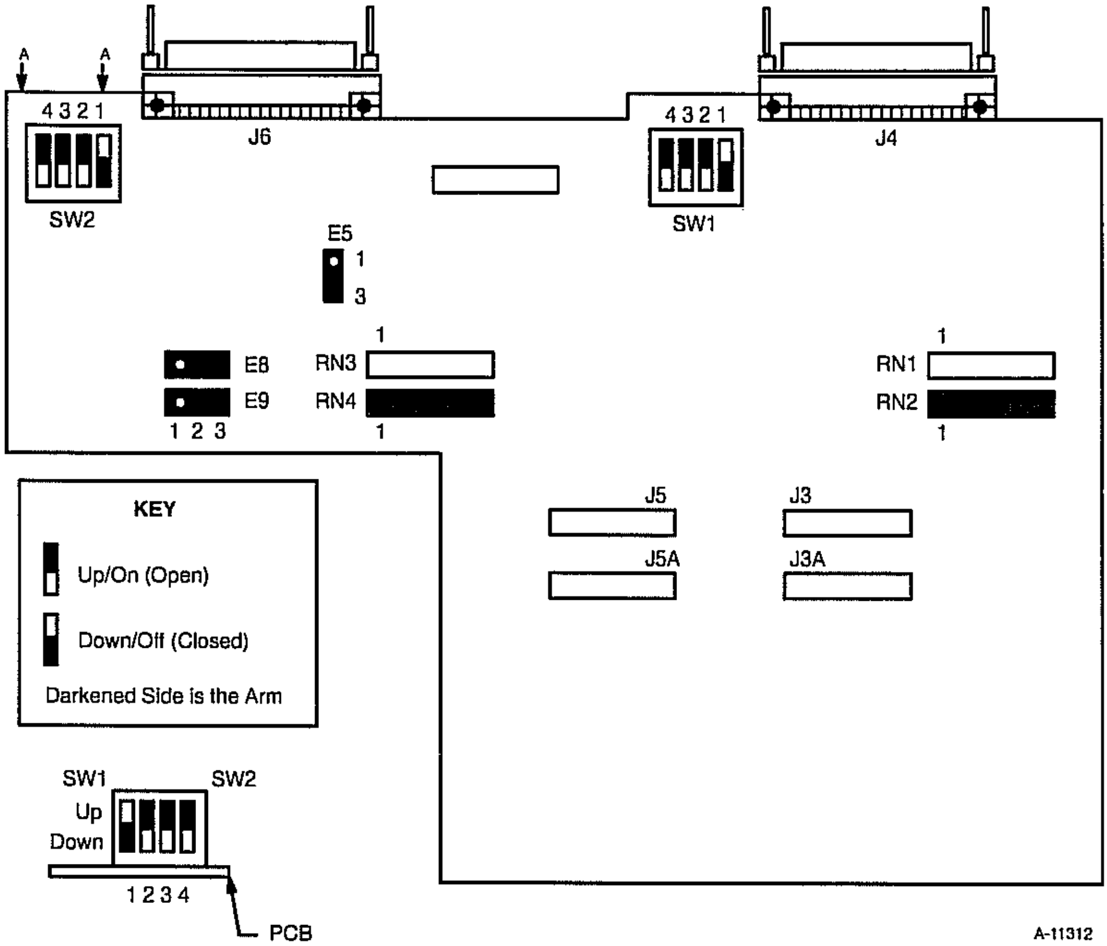

Figure 5-12. Master Control Board Jumper and Switch Locations for the DataShuttle

Table 5-4 lists the configurations for the DataShuttle switches and jumpers shown in Figure 5-12.

NOTE: When the split option is selected, two DC-3 disk drive controllers must be used for each DR-1. If the non-split condition is selected, both drives in the DR-1 are controlled by the same DC-3.

Table 5-4. Jumper and Switch Configurations for ESDI Interface

Location	Placement	Affected Logic †
E1	2-3	Right drive
E2	2-3	Left drive
E3	1-2	Left drive
E4	1-2	Right drive
E5 ‡	**1-2 (Split)**	**Right drive**
	2-3 (Non-split)	
E6	1-2	Left drive
E7	1-2	Right drive
E8 ‡	**1-2 (Split)**	**Right drive**
	2-3 (Non-split)	Right drive
E9 ‡	**1-2 (Split)**	**Right drive**
	2-3 (Non-split)	Right drive
E10	Removed	Master control
E11	1-2	Master control
E12	1-2	Master control
E15	Removed	Master control
E16	Removed	Master control
E17	2-3	Master control
SW1 1	**Closed**	**Left drive**
SW1 2-4	**Open**	**Left drive**
SW2 1	**Closed**	**Right drive**
SW2 2-4	**Open**	**Right drive**
SW4	2-3	Master control
RN1 (220/330)	**Removed**	**Left drive**
RN2 (150) ‡	**Installed (Split)**	**Left drive**
RN2 (150) ‡†	**Removed (Non-split)**	**Left drive**
RN3 (220/330)	**Removed**	**Right drive**
RN4 (150) ‡†	**Installed**	**Right drive**

† Locations shown in bold print are user configurable. Please select split option.
‡ User should select split option.
‡† ESDI termination 150-Ω resistor network MDB P/N20070111.

The logical device addresses are shown with the physical locations for the DR-1 disk drives in Figure 5-13, Figure 5-14, Figure 5-15, and Figure 5-16.

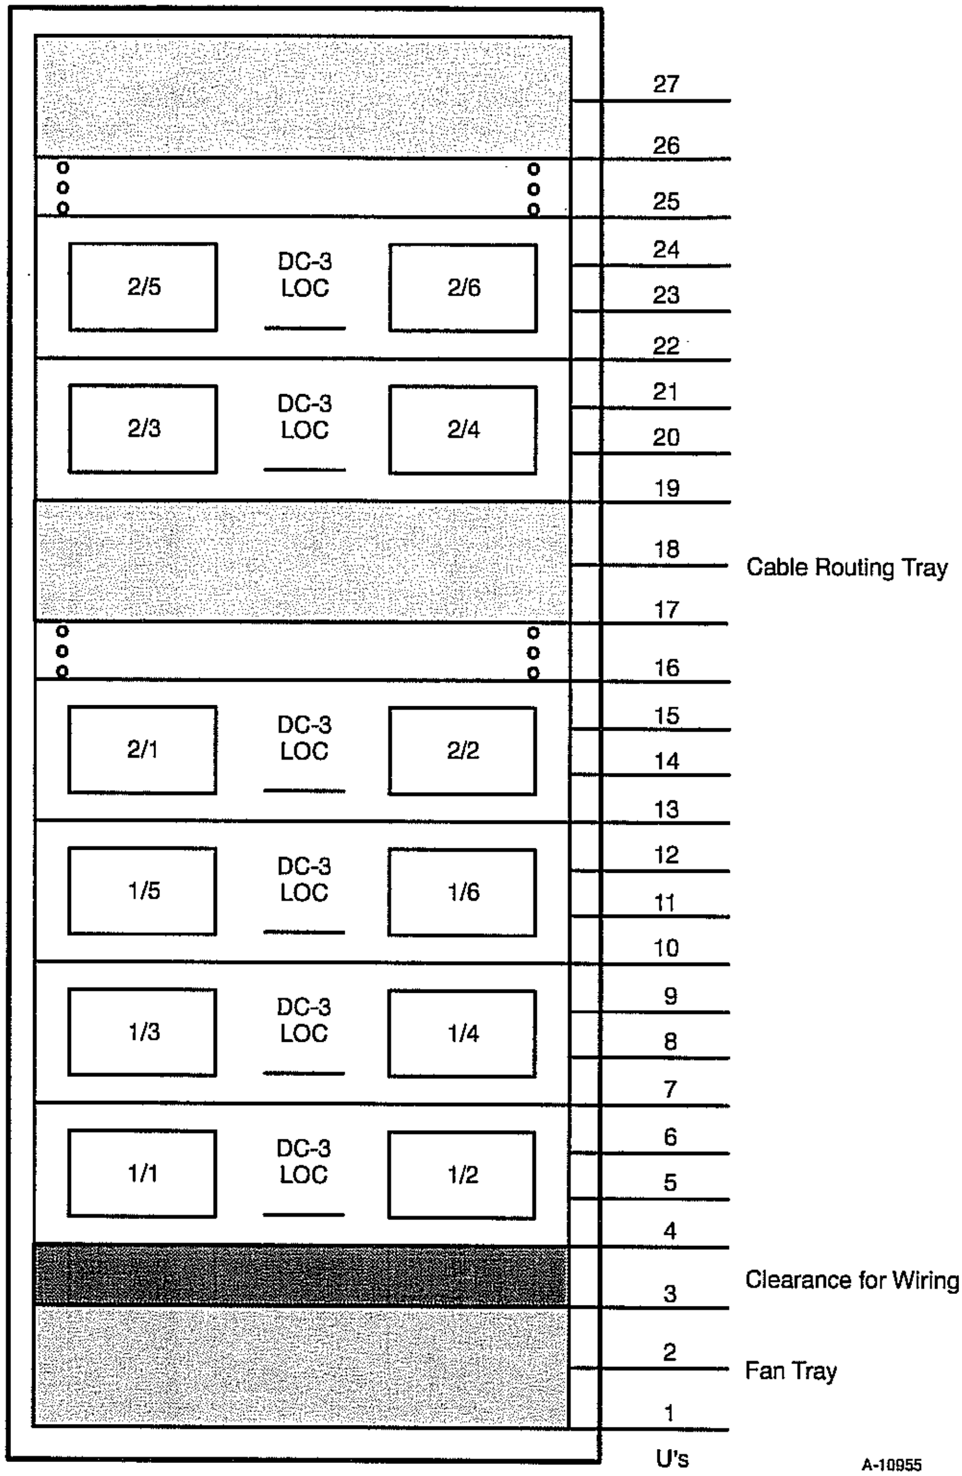

Figure 5-13. DR-1 Keying Process (Cabinet 1)

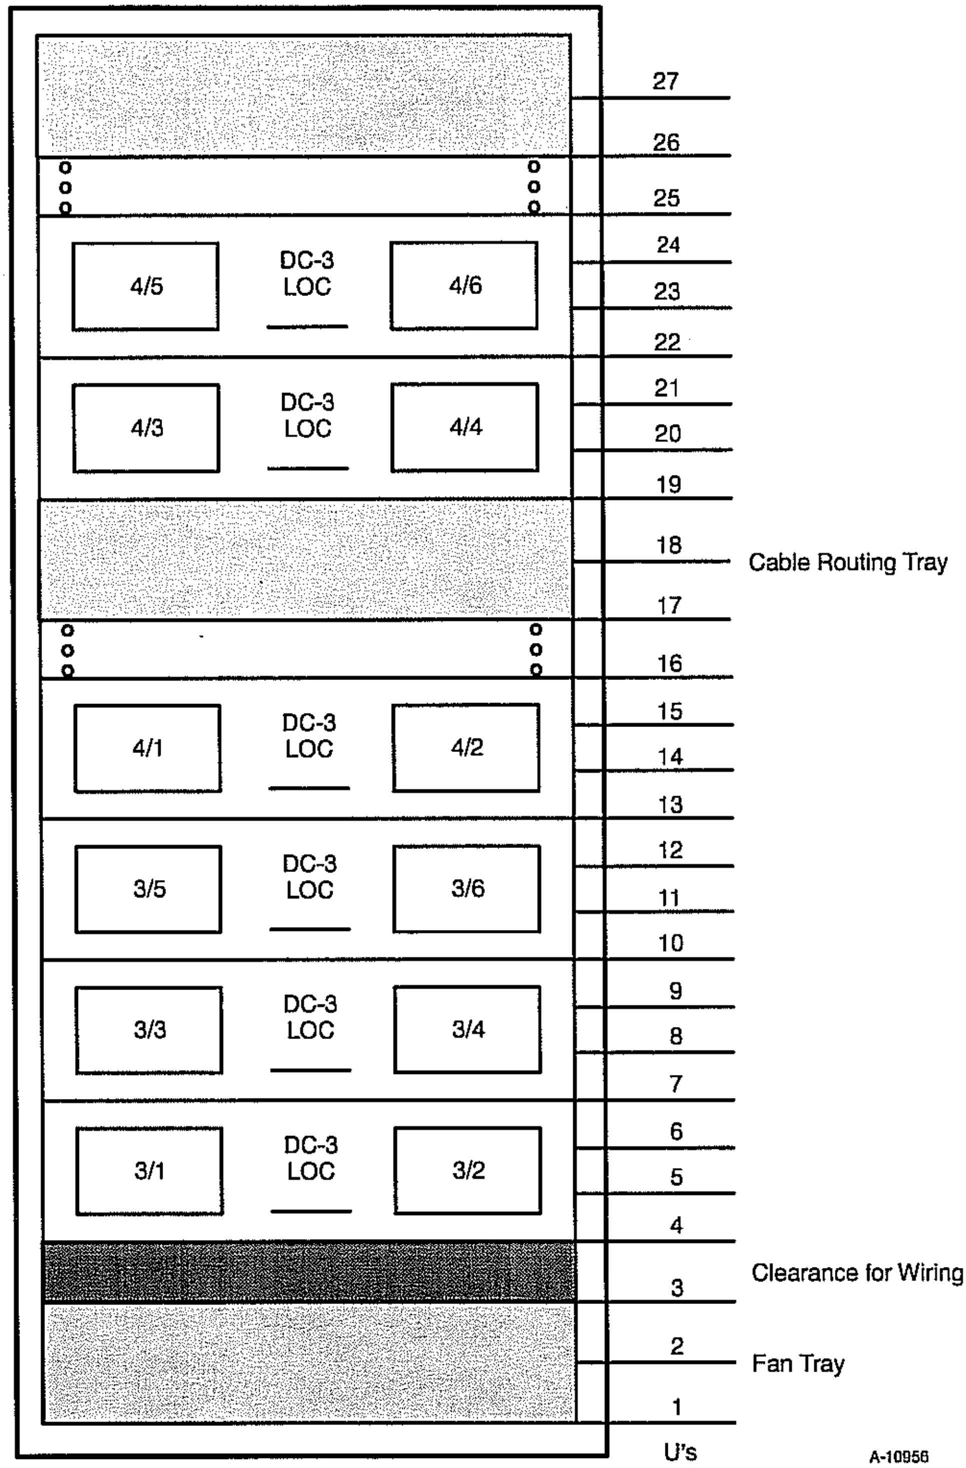

Figure 5-14. DR-1 Keying Process (Cabinet 2)

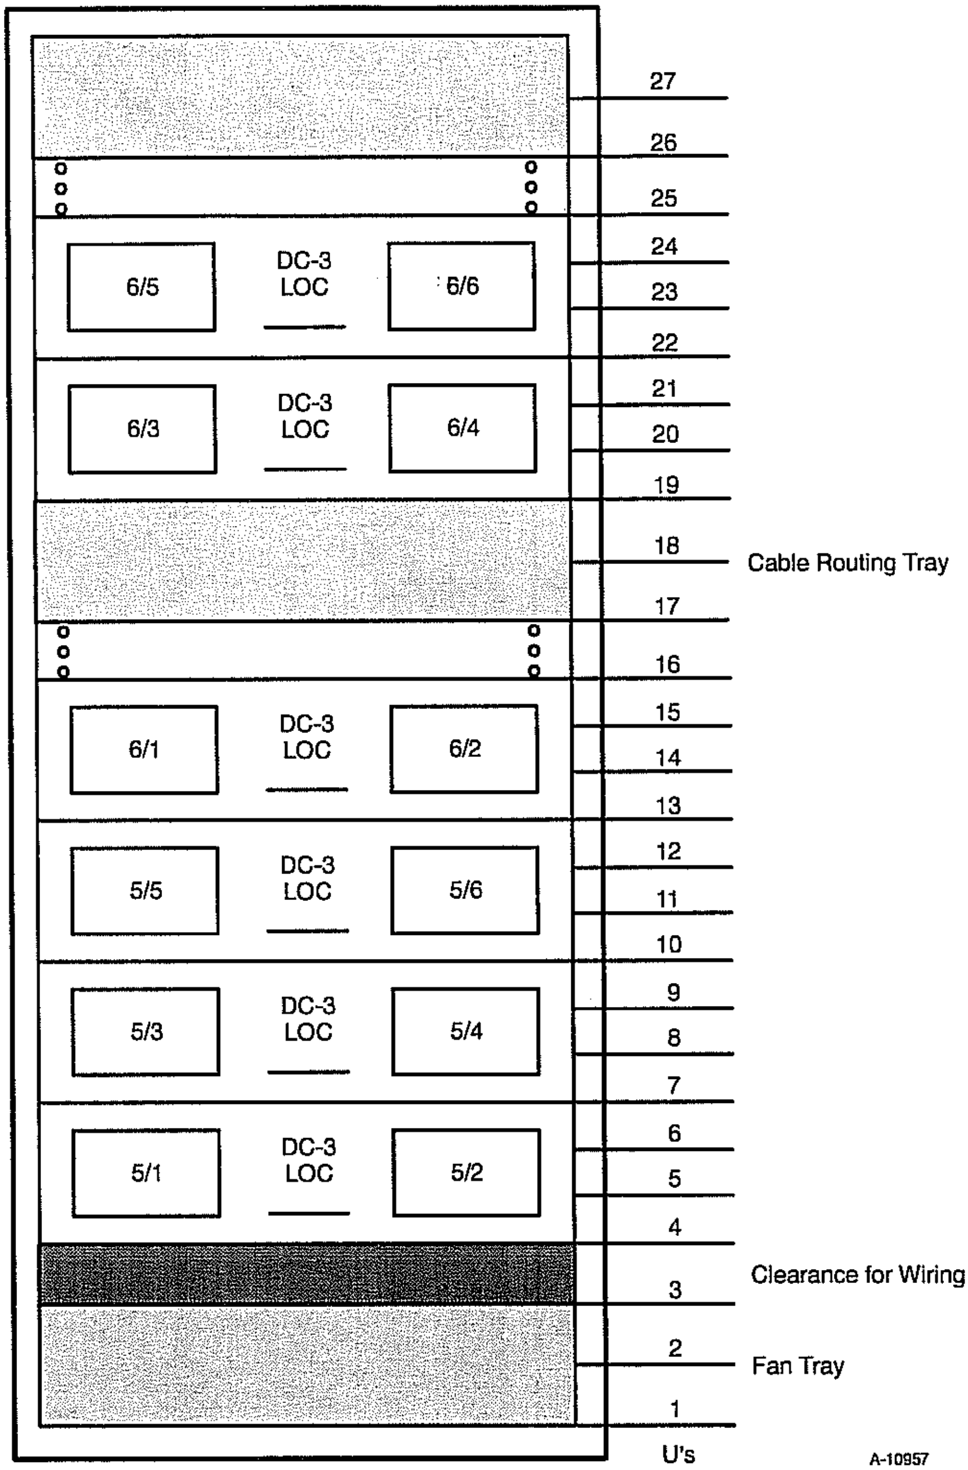

Figure 5-15. DR-1 Keying Process (Cabinet 3)

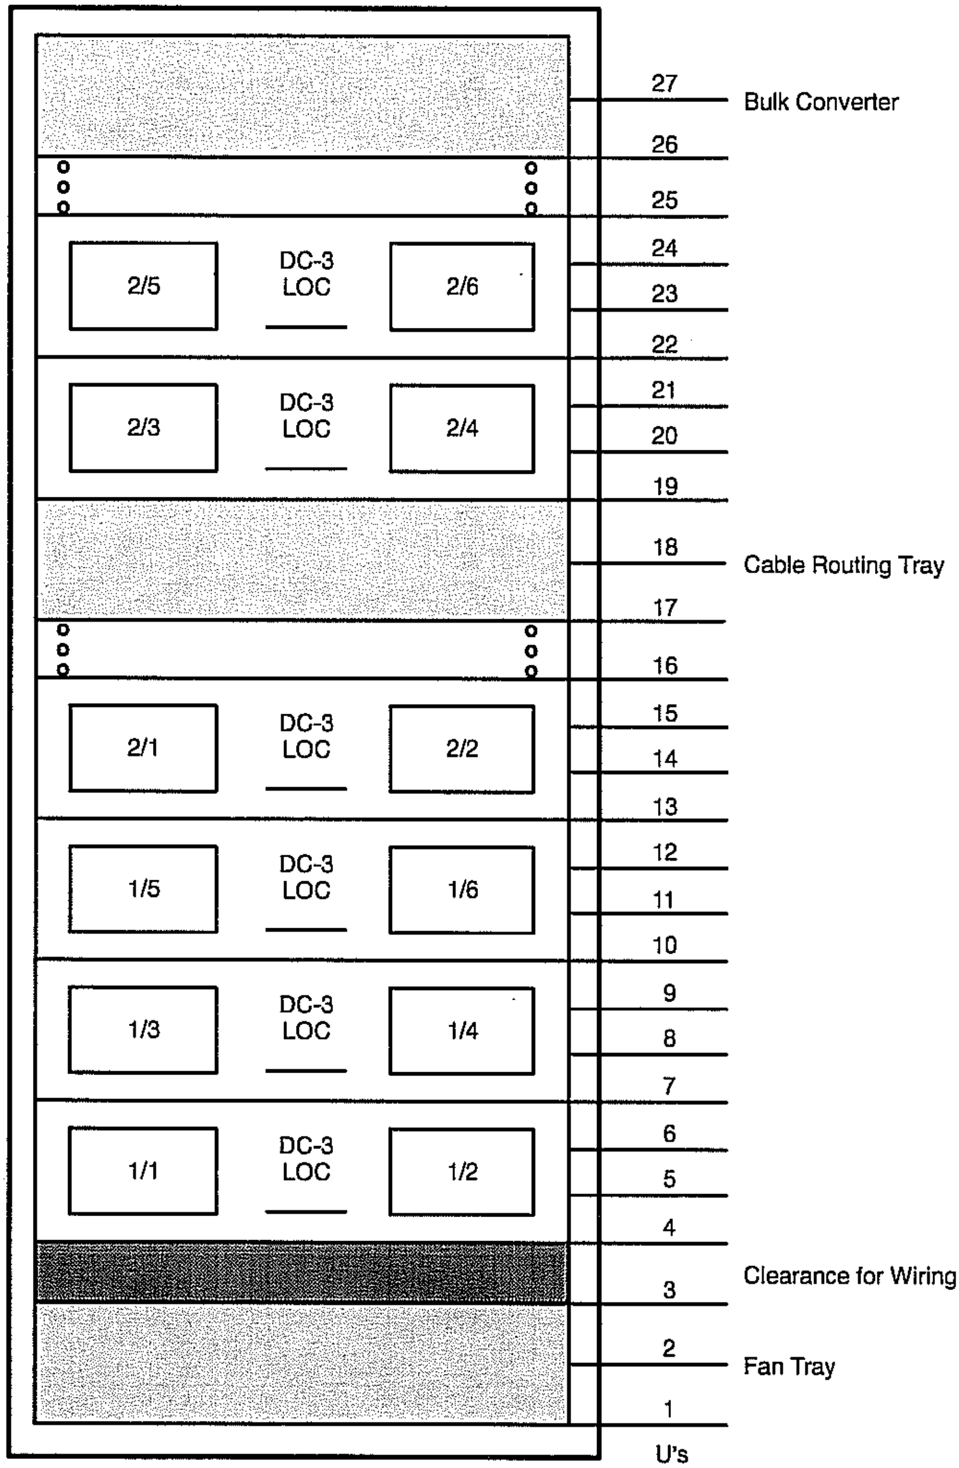

NOTE: Cabinet 4 has the same keying locations as Cabinet 1 because there are only 36 keying options. If more than one cabinet is to be used, the customer must be informed of the possible addressing conflict.

A-10958

Figure 5-16. DR-1 Keying Process (Cabinet 4)

DD-3 (ESDI) Disk Drive

The following characteristics apply to the CRAY EL series DD-3 disk drive, which is a Hitachi DK516 Winchester Disk Drive.

- Is a 5.25-in. Winchester disk drive
- Has a 1.2-Gbyte capacity
- Interfaces to an ESDI controller
- Supports a synchronized spindle option that can establish higher capacity and faster transfer rates to allow it to be used in the DAS-2 subsystem
- Has a data transfer rate of 2.753 Mbytes/s
- Specification sheet is found in the OEM product specification manual
- The format is established by the OEM
- The power connector is configured as follows:

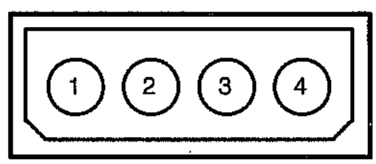

1 : + 12 Vdc
2 : + 12 V Return (Ground)
3 : + 5 V Return (Ground)
4 : + 5 Vdc

- Error codes are listed in the OEM product specification manual
- Media defect information is recorded on the first sector of each head on the upper cylinders
- No offline maintenance is possible, except flaw management
- No preventive maintenance is required

Figure 5-17 shows the PC board layout and jumper locations for the DD-3. Figure 5-18 through Figure 5-25 are enlarged portions of Figure 5-17.

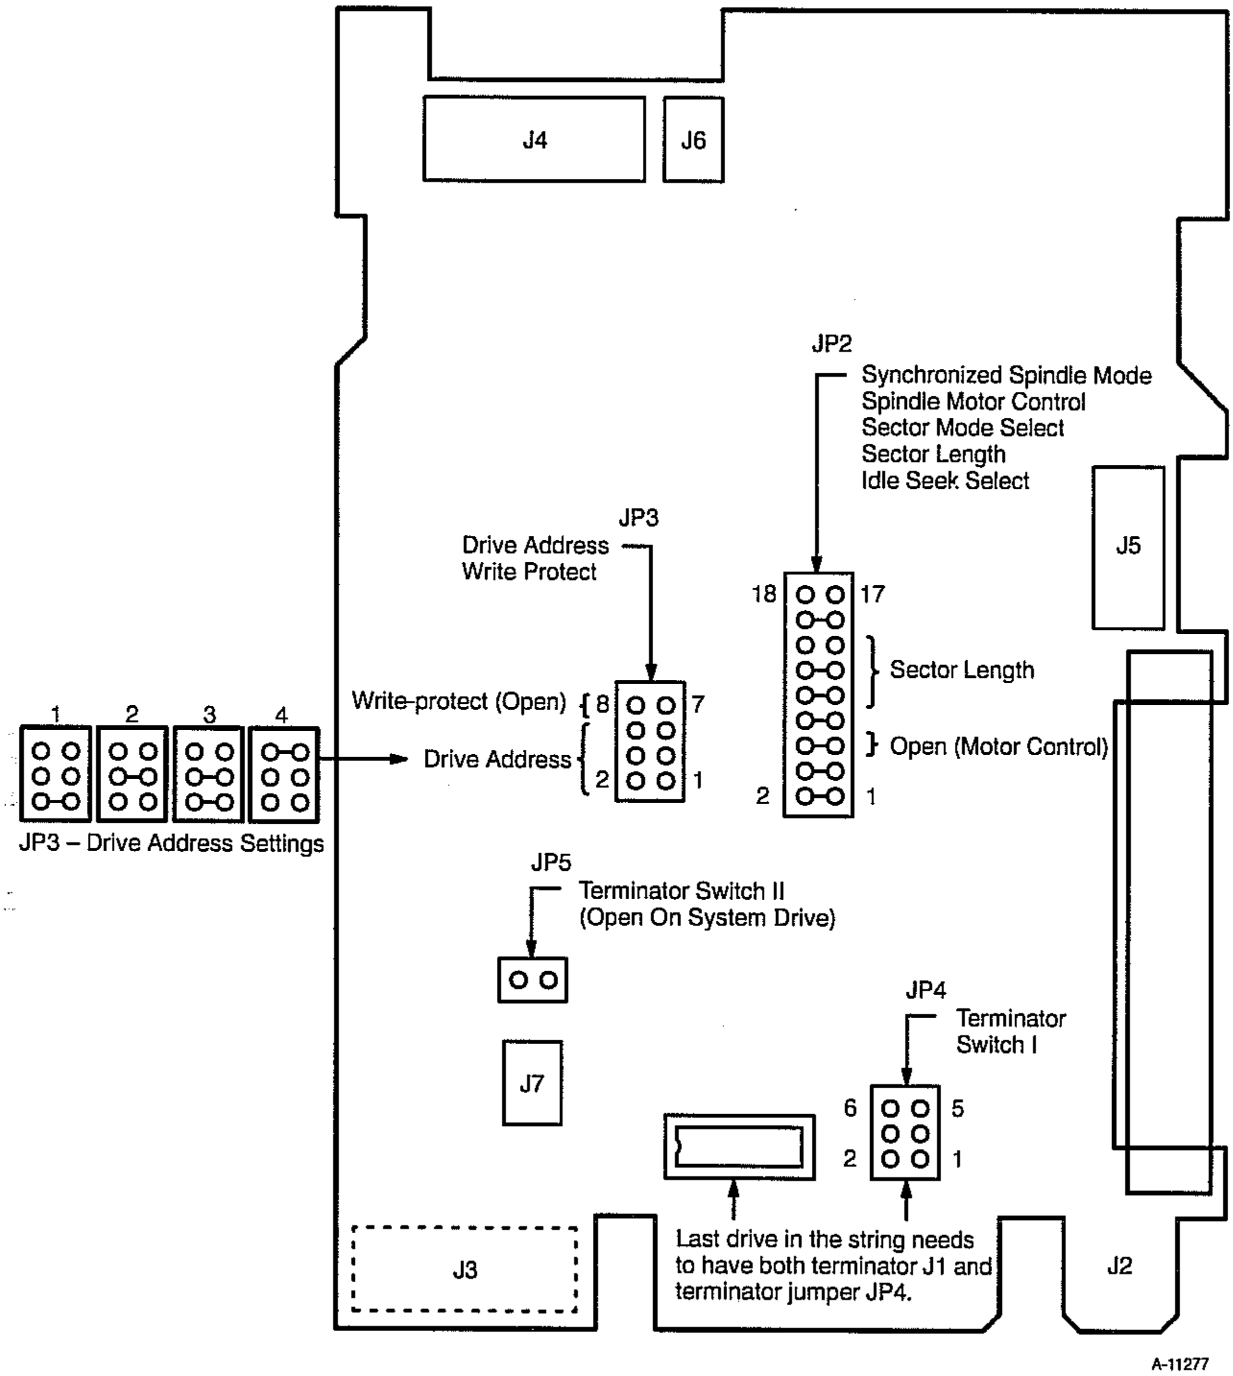

Figure 5-17. ESDI DK516 System SZ963 PCB Layout

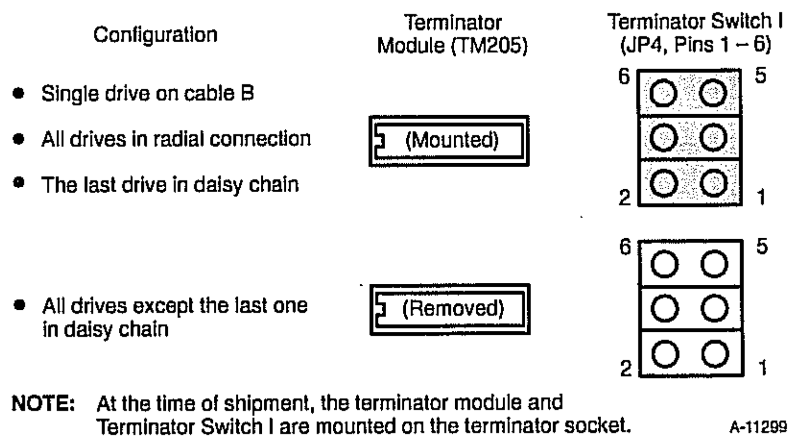

Figure 5-18. Terminator Module and Terminator Switch I

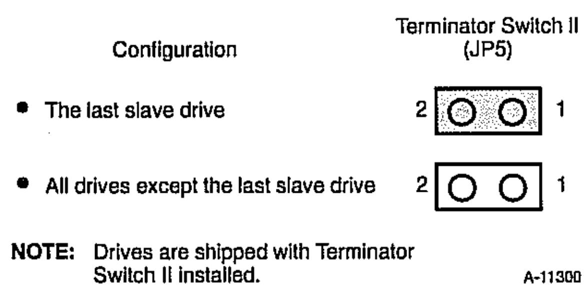

Figure 5-19. Terminator Switch II

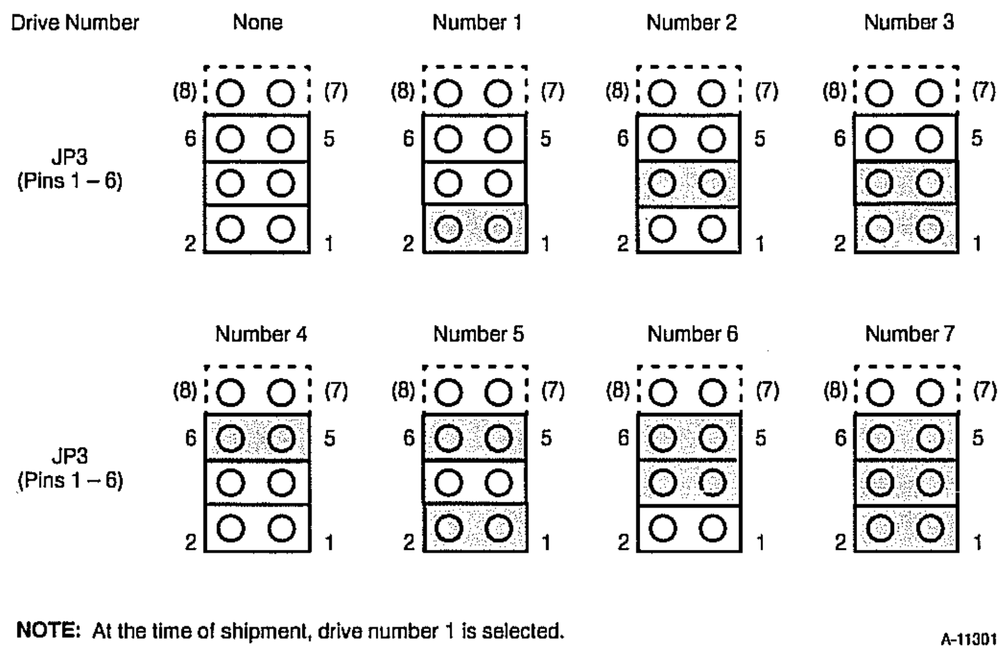

Figure 5-20. Jumper Setting for Drive Address

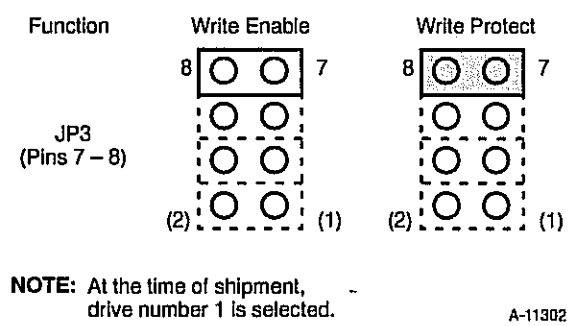

Figure 5-21. Jumper Setting for Write Protect

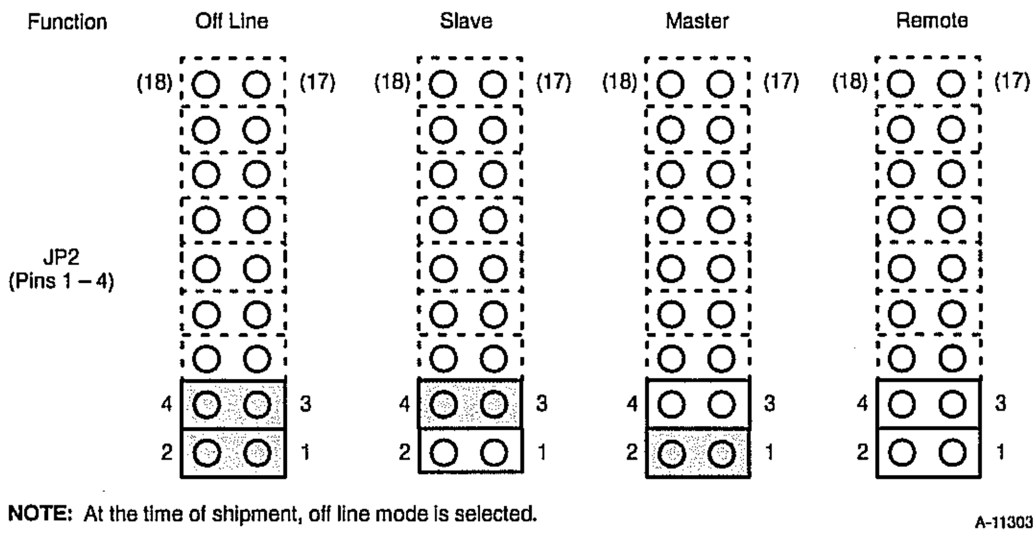

Figure 5-22. Jumper Setting for Synchronized Spindle Mode Select

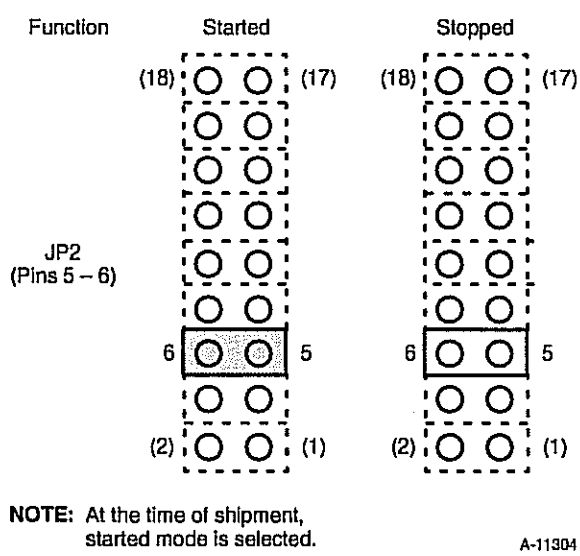

Figure 5-23. Jumper Setting for Spindle Motor Control

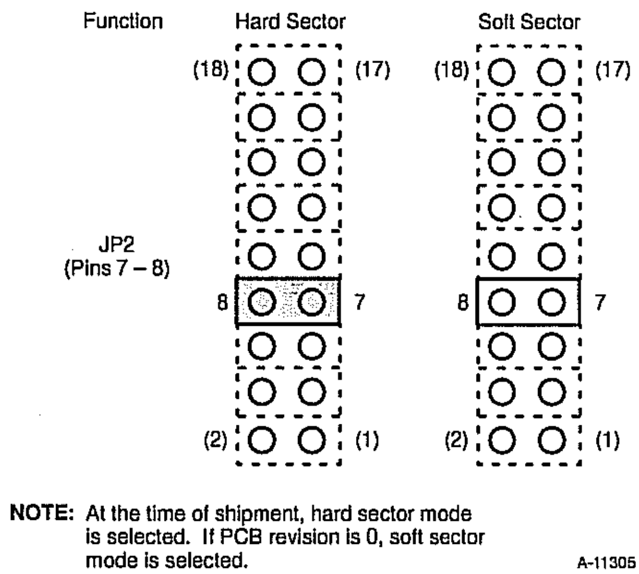

Figure 5-24. Jumper Setting for Sector Mode Select

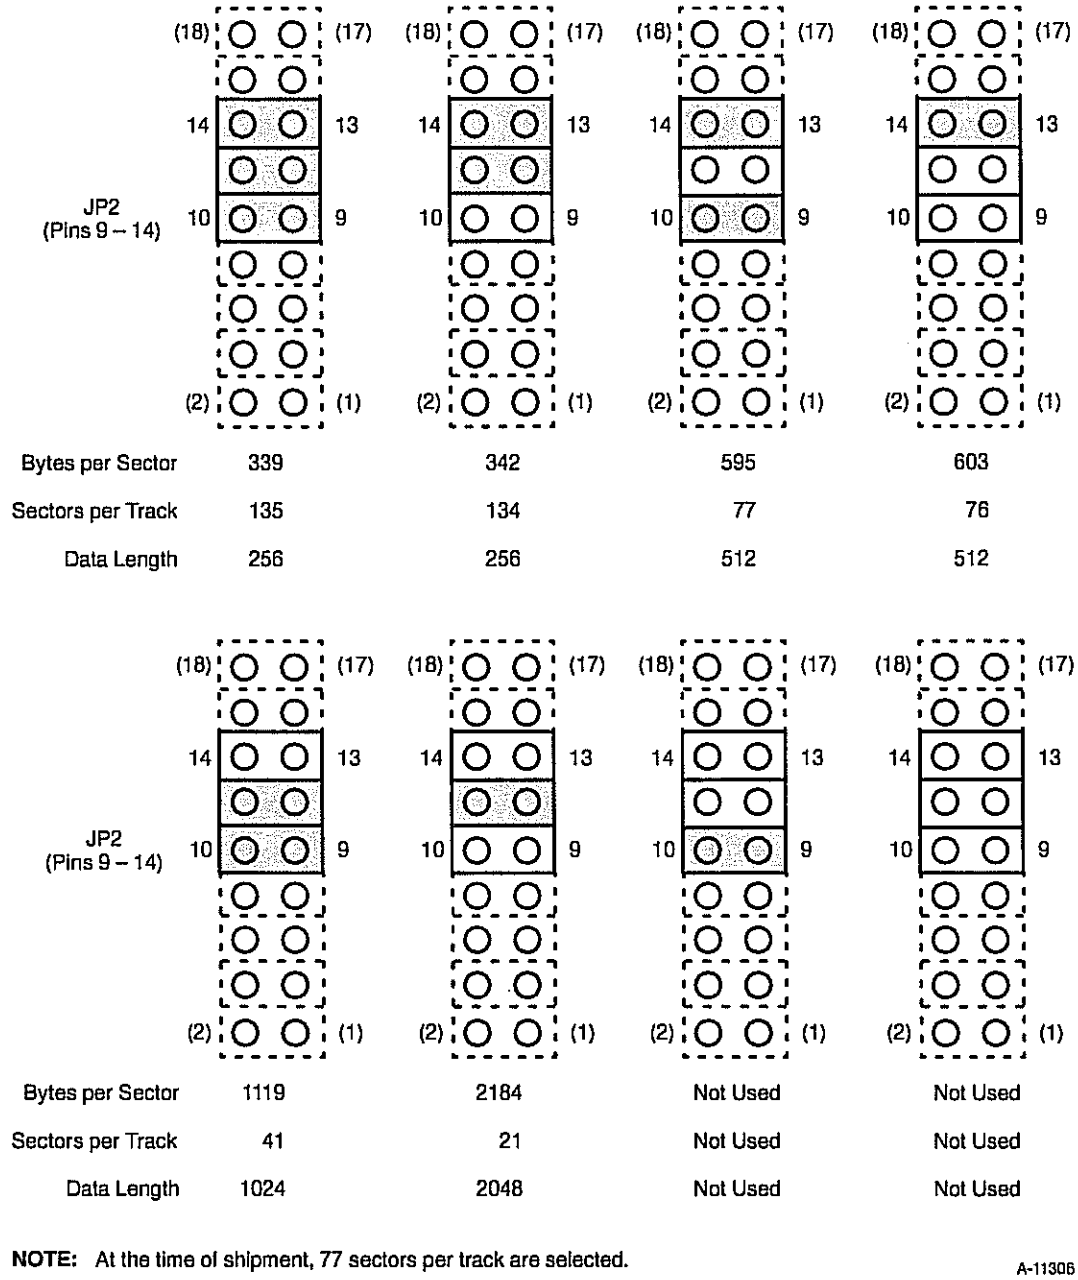

Figure 5-25. Jumper Setting for Sector Length

DD-4 (IPI) Disk Drive

The following characteristics describe the DD-4 disk drive, which is a Seagate Technology, Inc. Sabre disk drive (ST825000K).

- The DD-4 uses an IPI-2 protocol.
- Drive specifications are listed in the OEM manual.
- The optional status/control panel is not supplied.
- Normal ESD precautions are required when removing or replacing the DD-4.
- The DD-4 uses a single-phase switching-type power supply that conforms to the following specifications:

  177 to 264 Vac
  50/60 HZ (48 to 62 Hz)
  1.4 – 1.3 A

- I/O cable pin assignments are listed in the OEM manual
- When you receive a replacement drive, ensure correct input voltage selection by following the procedure in the OEM manual.
- In the absence of the control/status board, the only power-on check available is to ensure that the cooling fan on the power supply is operating. If the power supply fan is not operating, ensure that the power supply ON/STAND BY switch is in the ON position.
- After system power up, the DD-4 immediately starts its spin-up cycle, which completes within 90 seconds.

Figure 5-26 shows the cable configuration for the DD-4 disk drive.

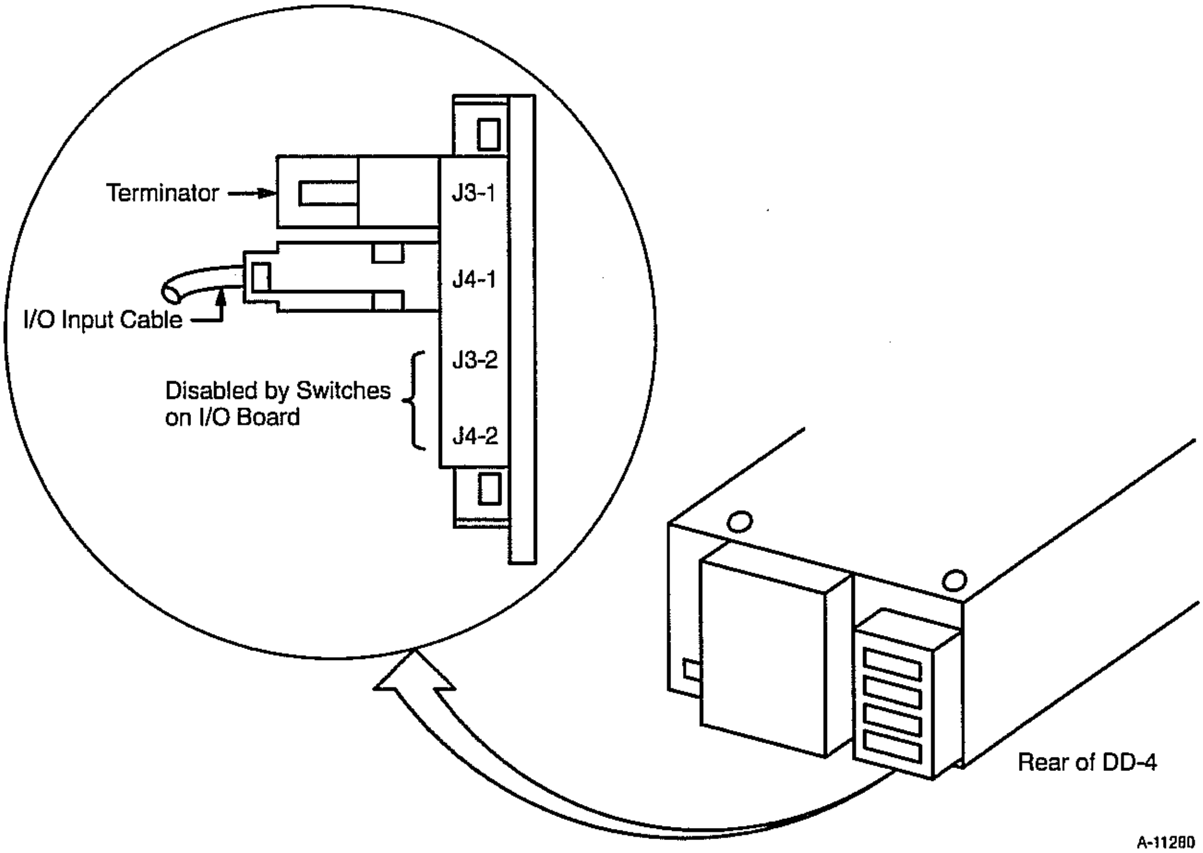

Figure 5-26. I/O Cable Location

DD-4 Installation and Addressing

Each DC-4 (IPI) controller supports two channels with up to two drives per channel. Bit 2^0 of the logical unit address represents the channel, and bit 2^1 of the logical unit address represents the drive's address on the channel. When a DC-4 is connected to four DD-4 drives (two per channel), the drives are addressed as units 0 through 3. Refer to Figure 5-27.

If only two DC-4 drives are connected to a DC-4 controller, one drive is connected to channel 0 and the other to channel 1. Both drives are set to logical address 0, and both reside in the same PE-4 drawer.

When four DC-4 drives are connected to a DC-4 controller, the drives are divided between two peripheral equipment (PE) drawers (two per drawer). Channel 0 of the controller daisy chains across the two drives in one drawer, and channel 1 daisy chains across the two drives in the second drawer. Within each drawer, one drive is set to logical address 0 and the other is set to logical address 1. Refer again to Figure 5-27.

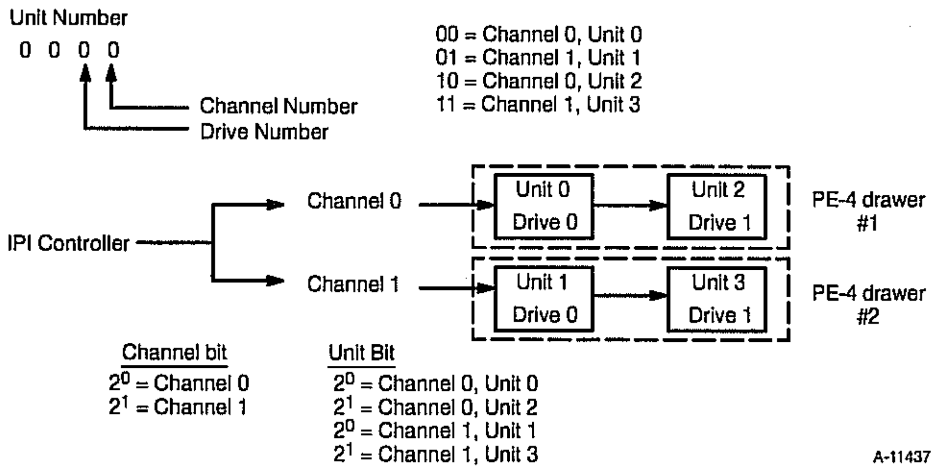

Figure 5-27. DD-4 (IPI) Disk Device Naming Conventions

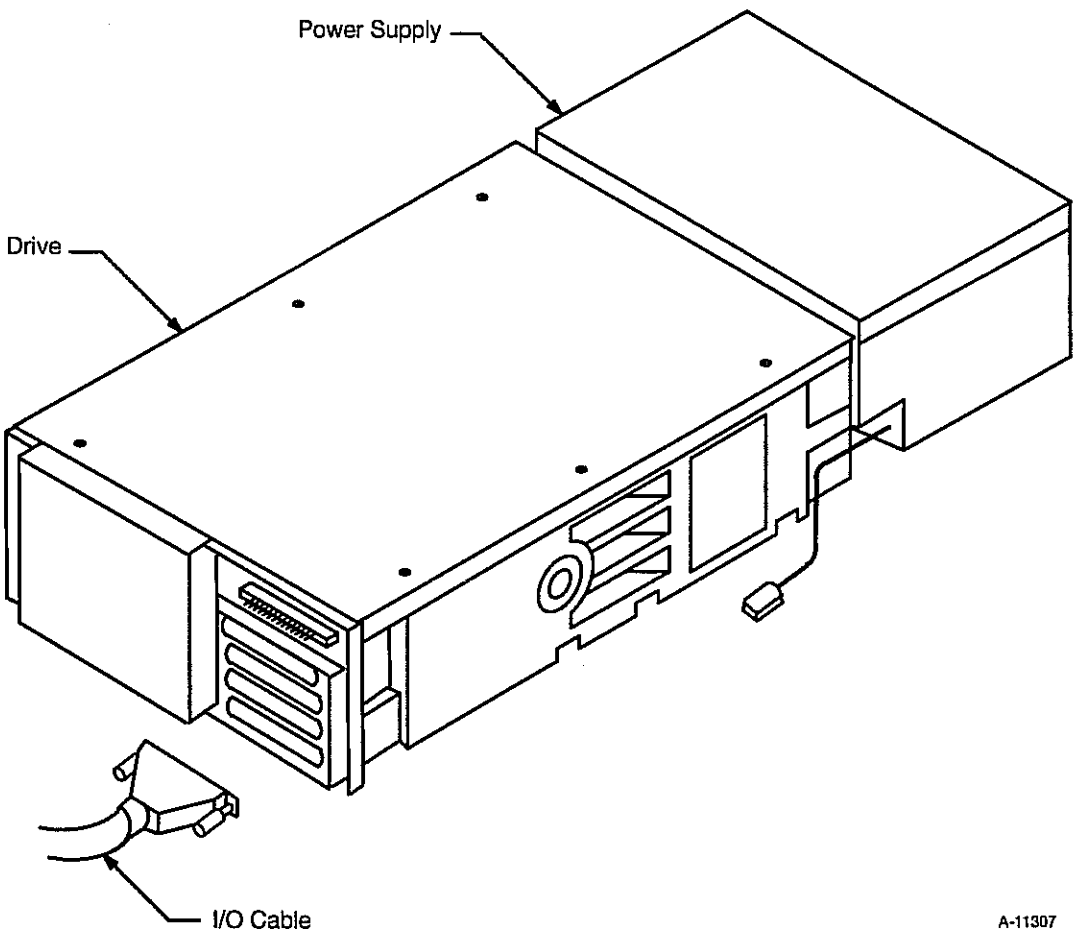

Figure 5-28. DD-4 (IPI) Disk Drive

DD-4 Power Supply

Site power enters the optional power supply through the AC power cable. The power supply converts the AC voltage to the DC voltages needed by the drive. These voltages reach the drive through the DC power cable. Refer to Figure 5-29.

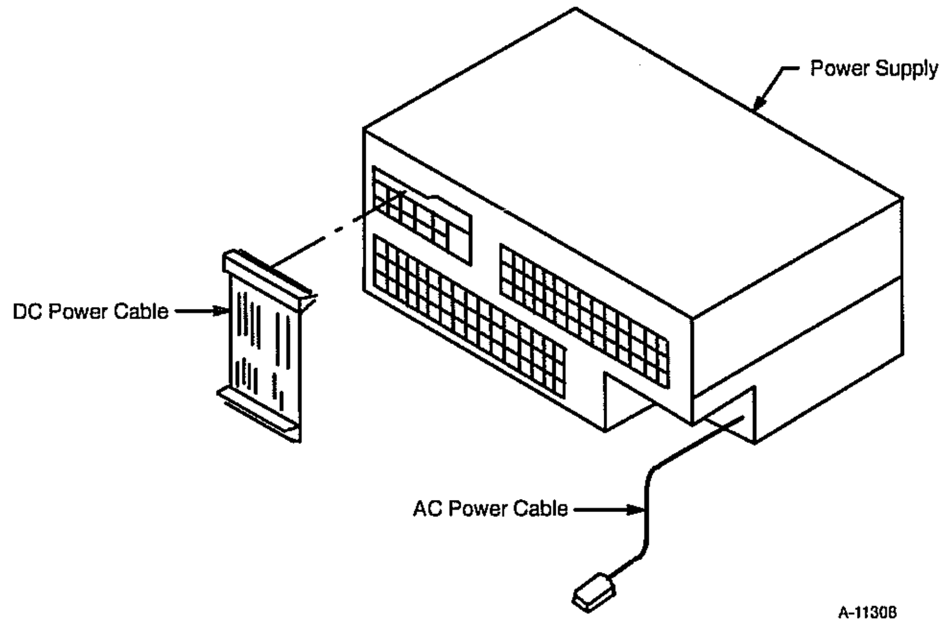

Figure 5-29. DD-4 Power Supply

Setting the Circuit Board Switches

Two groups of dual in-line package (DIP) switches enable you to tailor the drive to your exact needs. One set of DIP switches is on the I/O board; you can reach them through a hole at the left side of the rear panel. The other set of DIP switches is on the control board; you can reach them through holes in the drive top cover.

The DIP switches may be one of two types: rocker or slide. You can move a rocker switch by pressing one end of the actuator or the other (rocking it) to turn the switch ON (closed) or OFF (open). You can turn a slide switch on or off by sliding the actuator one way or the other. Use a slender ball-point pen, a straightened paper clip, or any similar object to change the switch settings. Refer to Figure 5-30.

CAUTION
Do not use a lead pencil point to move a DIP switch. The lead may break off and lodge in the switch or cause the switch to malfunction.

The switches are mounted in a plastic case and are usually numbered. Other labels may appear next to the switches on the circuit board, or a label may appear at the sides of the access hole in the drive. The position of the labels may not always coincide with the switch setting that enables the function. Always use the switch settings shown in Figure 5-31 and Figure 5-32 to properly set the switches for your needs. A switch is considered closed in the ON position and open in the OFF position.

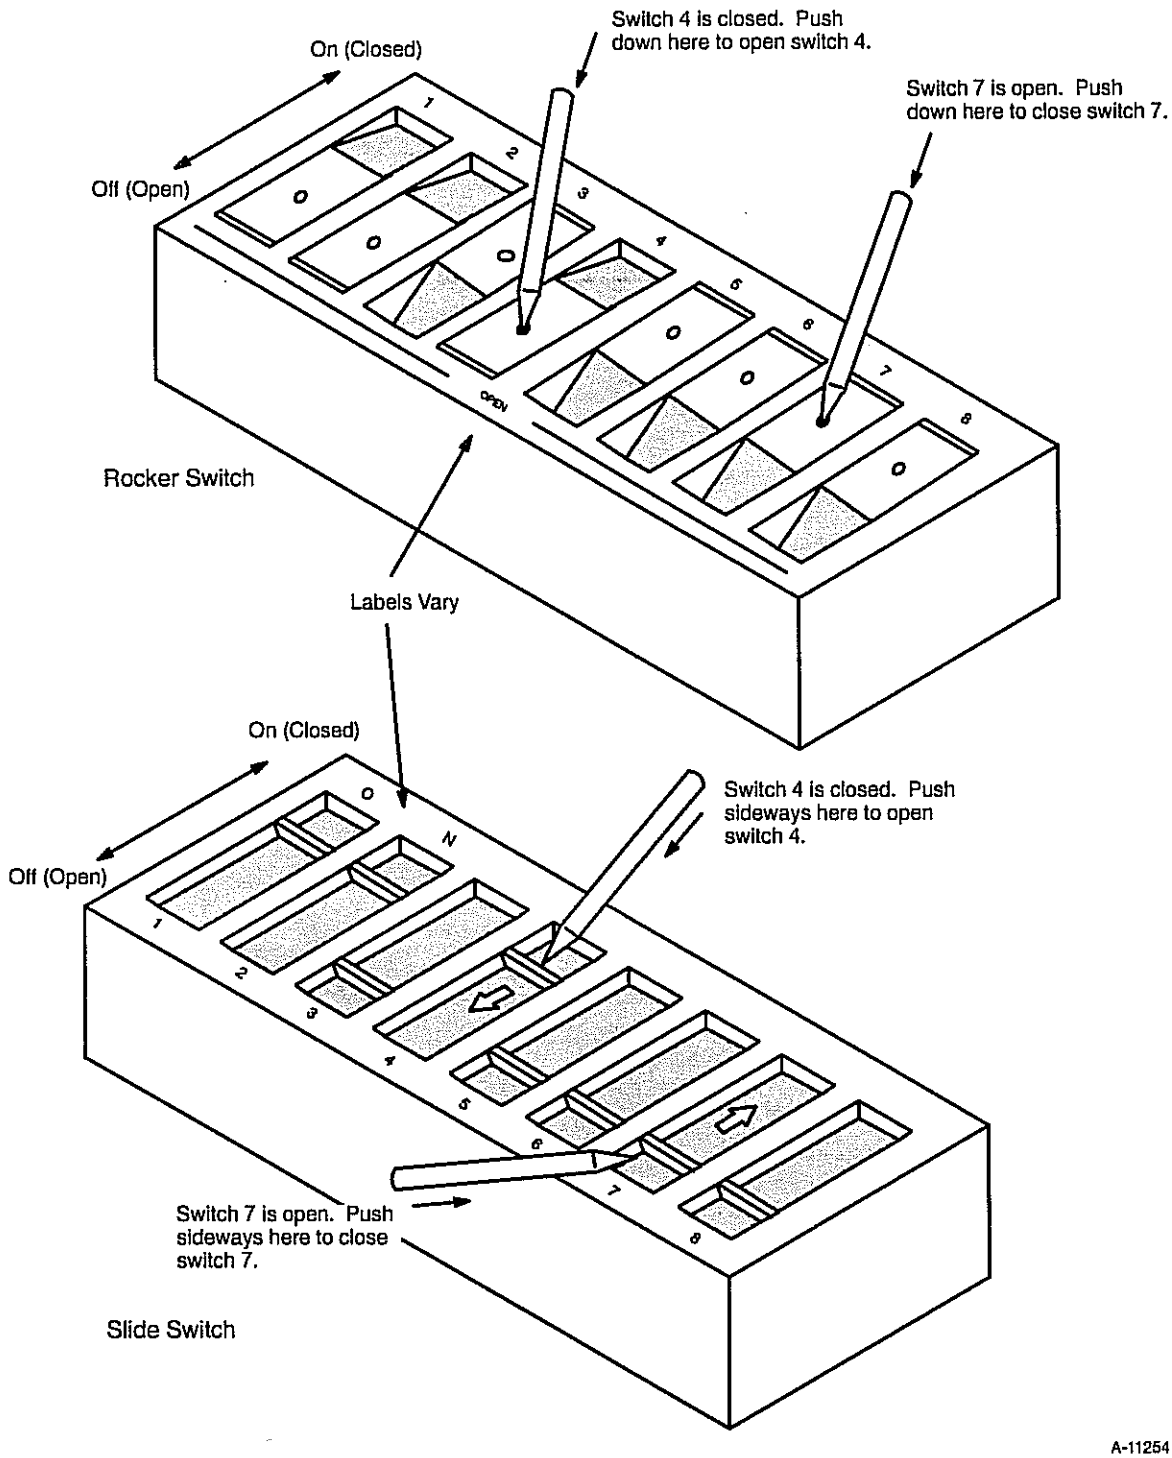

Figure 5-30. Setting the DD-4 Circuit Board Switches

Identifying Switches on the I/O Board

The switches on the I/O board are accessible through a hole in the lower left side of the rear panel. Refer to Figure 5-31.

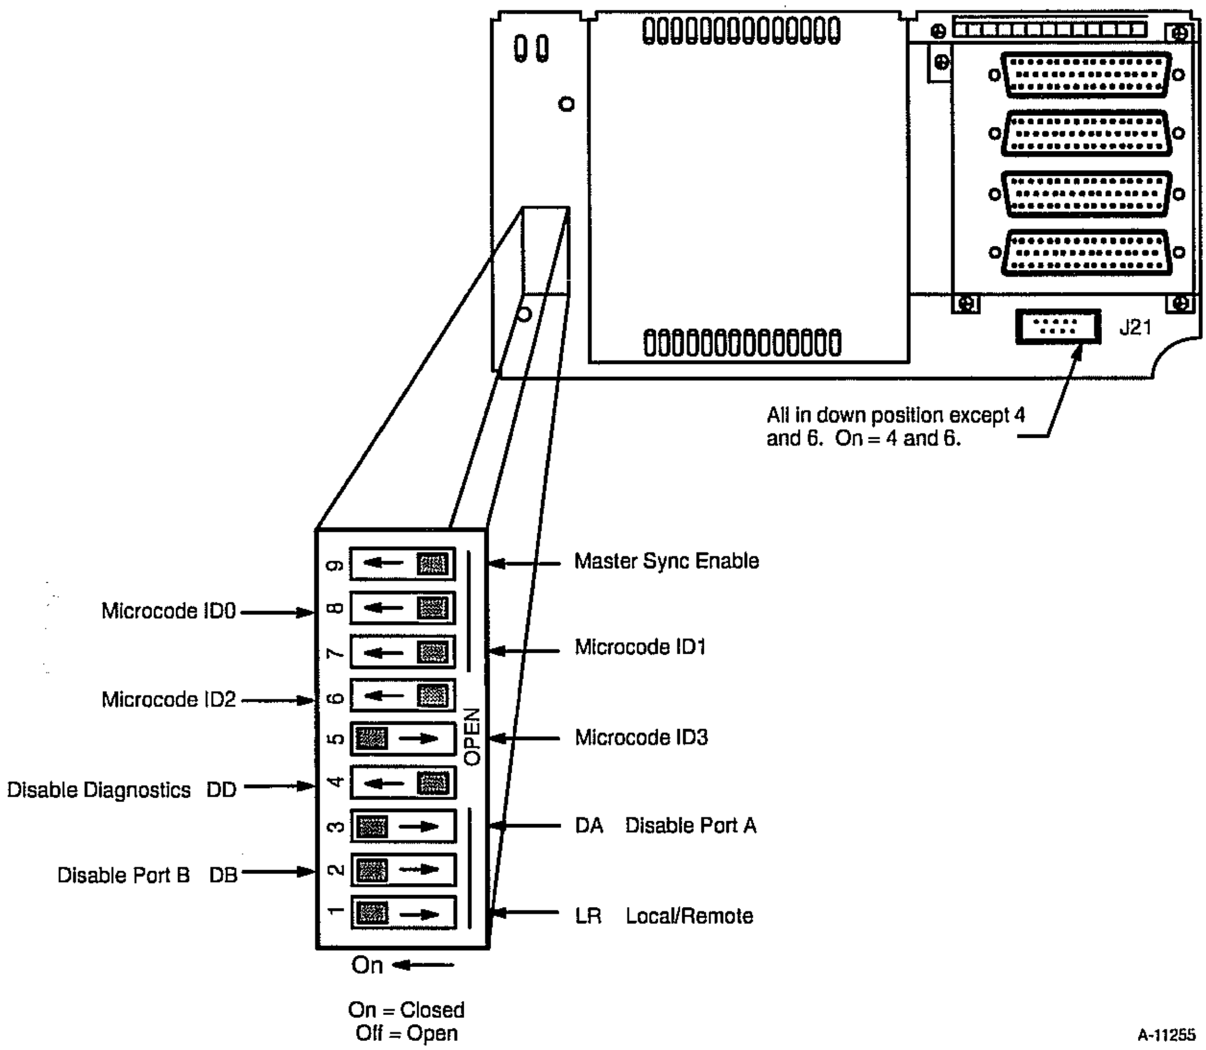

Figure 5-31. I/O Board Switches

Control Board Switch Identification

Refer to Figure 5-32 or Figure 5-33 to locate and identify the switches on the control board. When you have located and identified the control board switches, set the switches according to the VME slot location. The original control board includes labeled LED indicators.

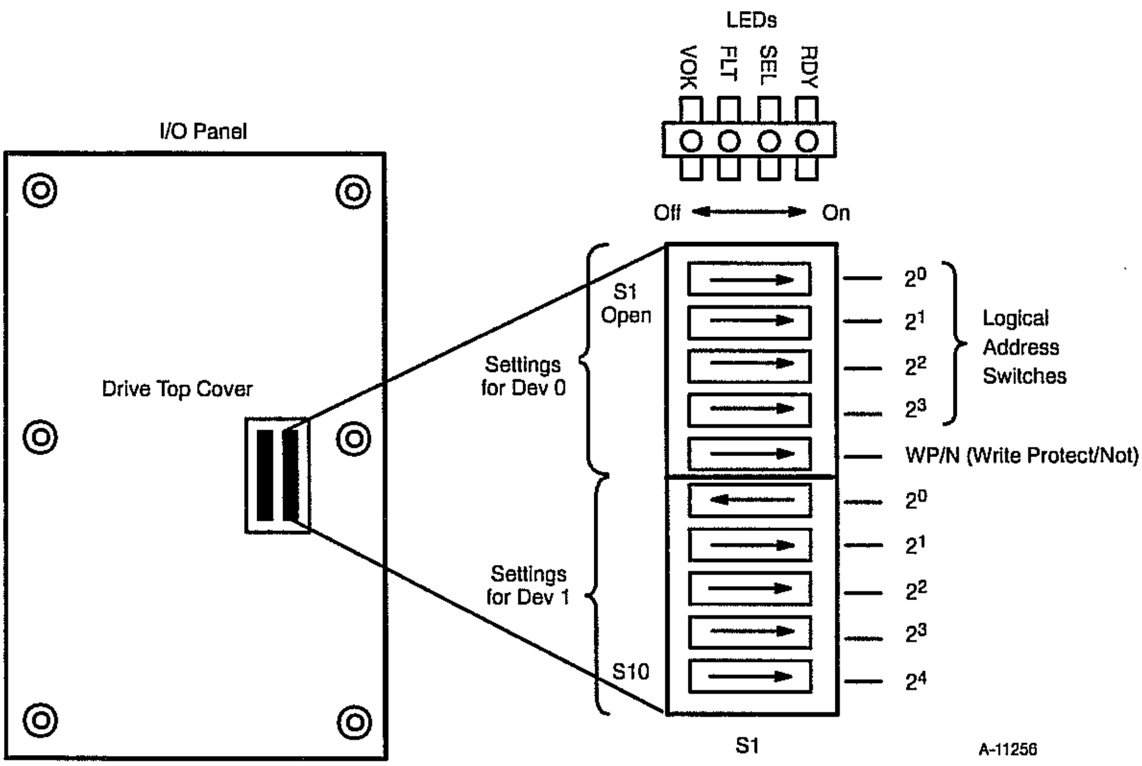

Figure 5-32. Original Control Board Switch Settings

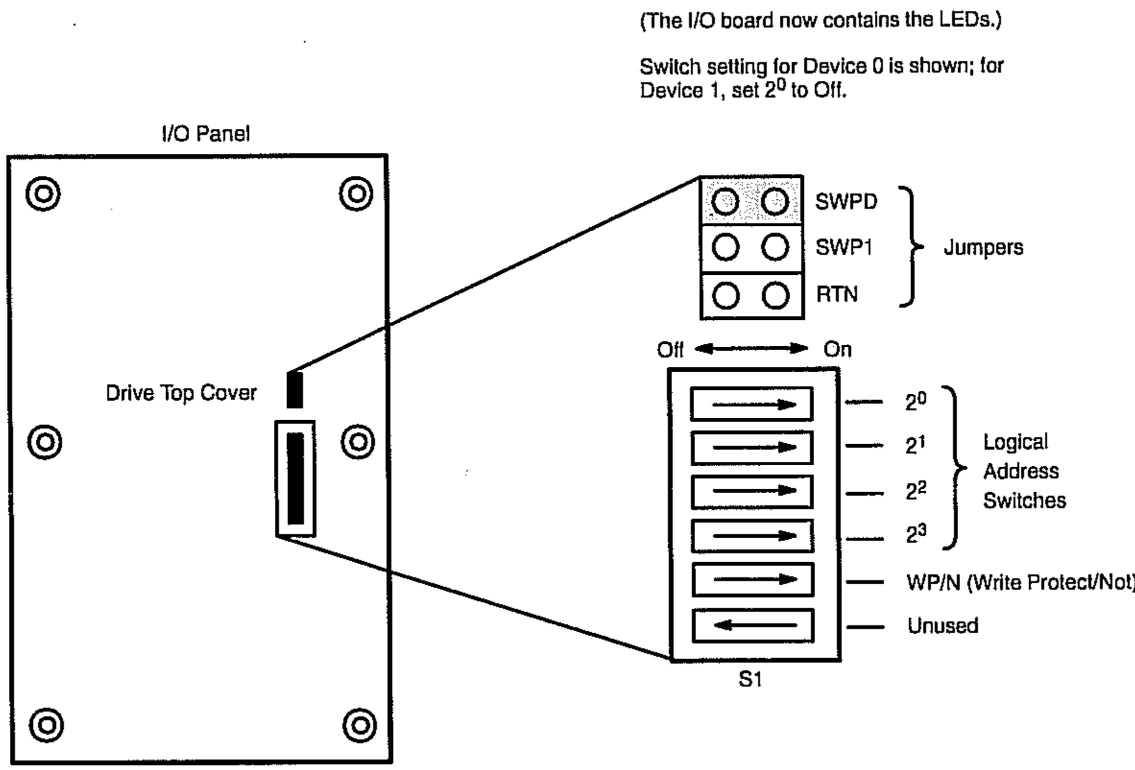

Figure 5-33. Revised Control Board Switch Settings

EXABYTE EXB-8500 8-mm Cartridge Tape Subsystem (EX-2)

The following characteristics describe the EXABYTE EXB-8500 8-mm cartridge tape subsystem.

- 5 Gbytes data storage capability
- Dual read/write head pairs with helical-scan recording technology
- Maximum sustained transfer rate of 500 Kbytes/s
- Peak transfer rate of 4 Mbytes/s
- Incorporates ECC and full read-after-write verification for data integrity
- If the EX-2 is at the end of the SCSI bus (the normal location), it must have external bus terminators installed
- Can write and read in either EXB-8200 or EXB-8500 format, but can use only one format type per tape

- Explanation of the front LEDs indicating drive condition is in the *EXABYTE User's Manual,* number 510201
- A complete explanation of all fault codes is in Appendix E of the *EXABYTE User's Manual,* number 510201

Setting the SCSI ID with the DIP Switches

Refer to Figure 5-34 to locate and set the DIP switches on the EXB-8500 cartridge tape subsystem.

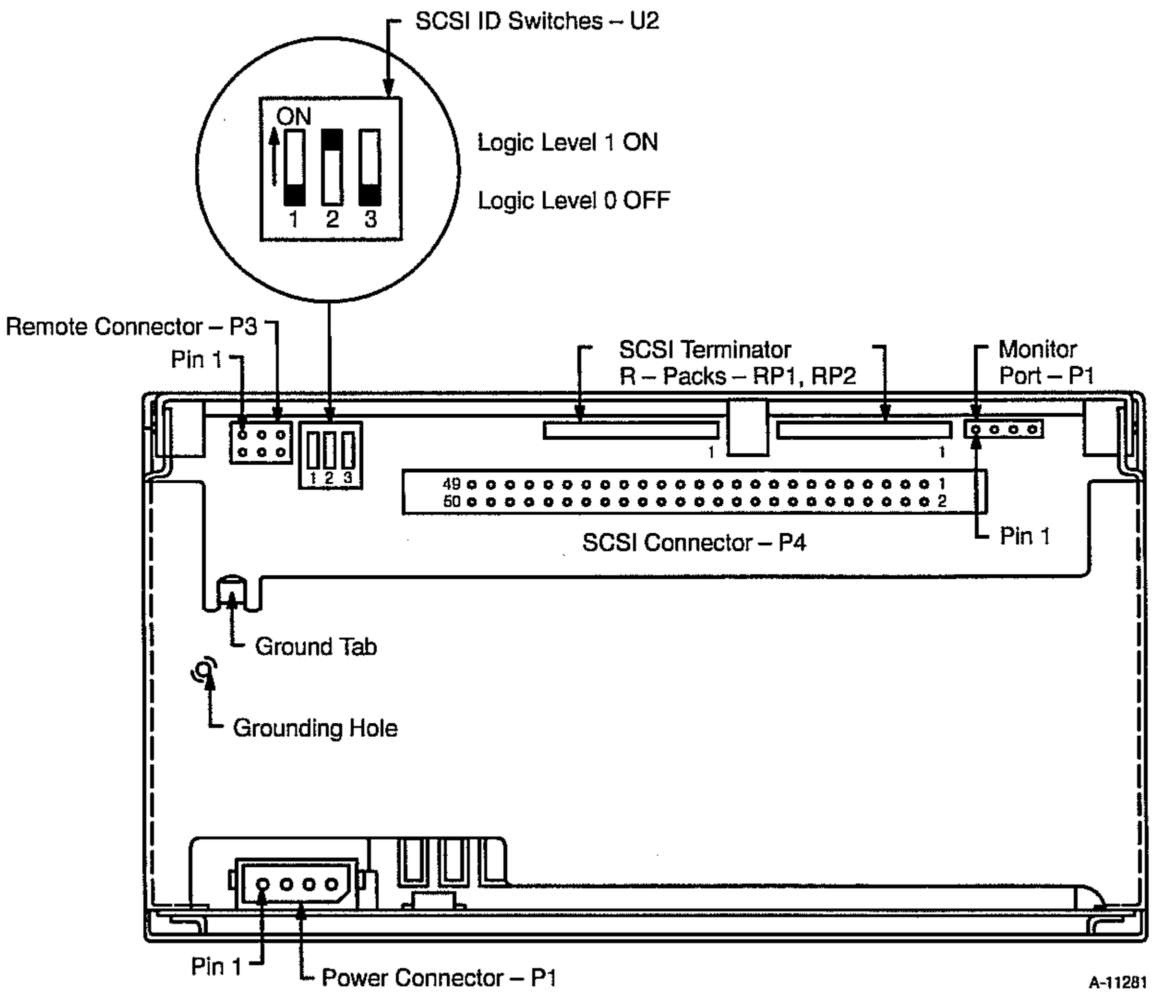

Figure 5-34. Connectors and Controls on the Back Panel of the EXB-8500

Anaconda 2800 Tape Drive

The following characteristics describe the Anaconda 2800 0.25-in. cartridge (QIC) streaming tape drive.

- Data storage capacity is determined by the type of data cartridge used
- Uses 5.25-in. chassis
- Contains 512-Kbyte data buffer
- Can support only logical unit number (LUN) 0
- Error descriptions and reports are explained in the Archive Technology, Inc. *Technical Reference* manual

Figure 5-35 shows logic for the Anaconda 2800 tape drive.

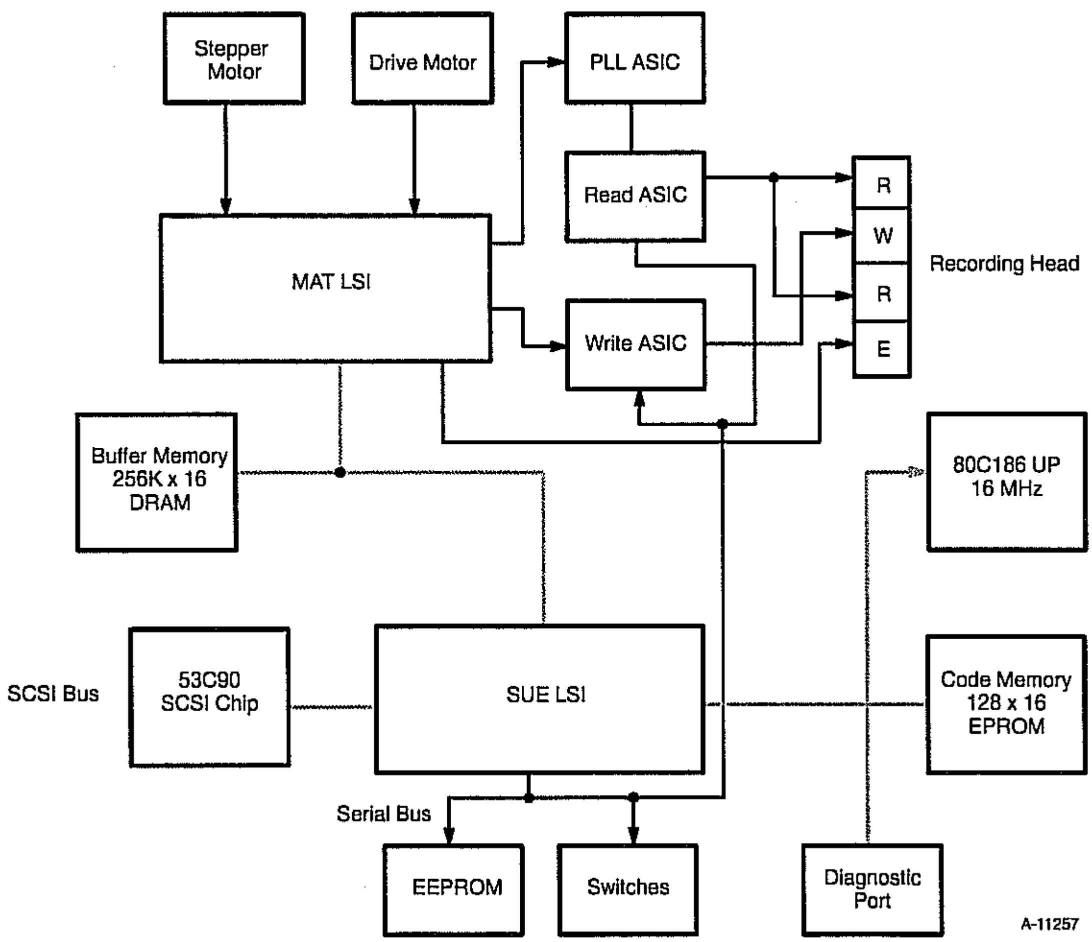

Figure 5-35. Anaconda Logic Block Diagram

Figure 5-36 is a block diagram of the firmware modules within the Anaconda 2800 tape drive.

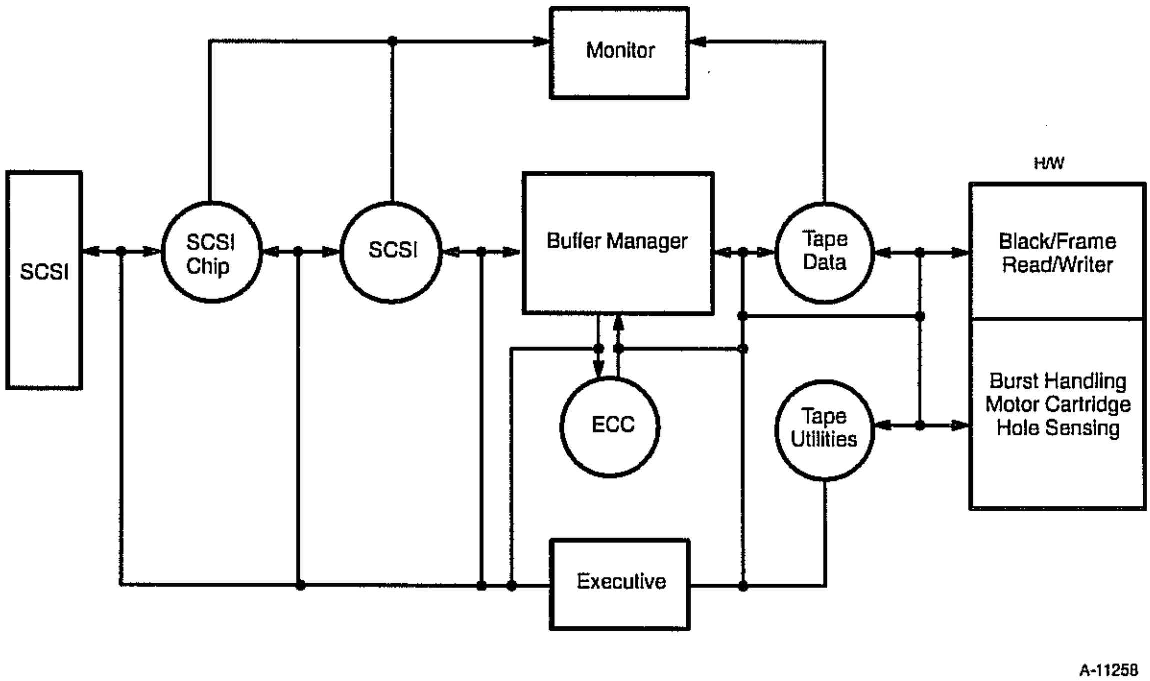

Figure 5-36. Anaconda Firmware Modules

Unloading a Cartridge (Manual Operation)

If a power outage or malfunction of the autoload feature occurs while a cartridge is loaded, you can remove the cartridge by following this procedure. Use a standard 0.125-in. (3.2-mm) allen wrench to remove the cartridge manually.

CAUTION
Ensure that power to the Anaconda 2800 tape drive is turned off before attempting to remove the cartridge manually or tape drive damage may occur.

1. Determine the position of the cartridge. Refer to the Archive Technology, Inc. *Technical Reference* manual.

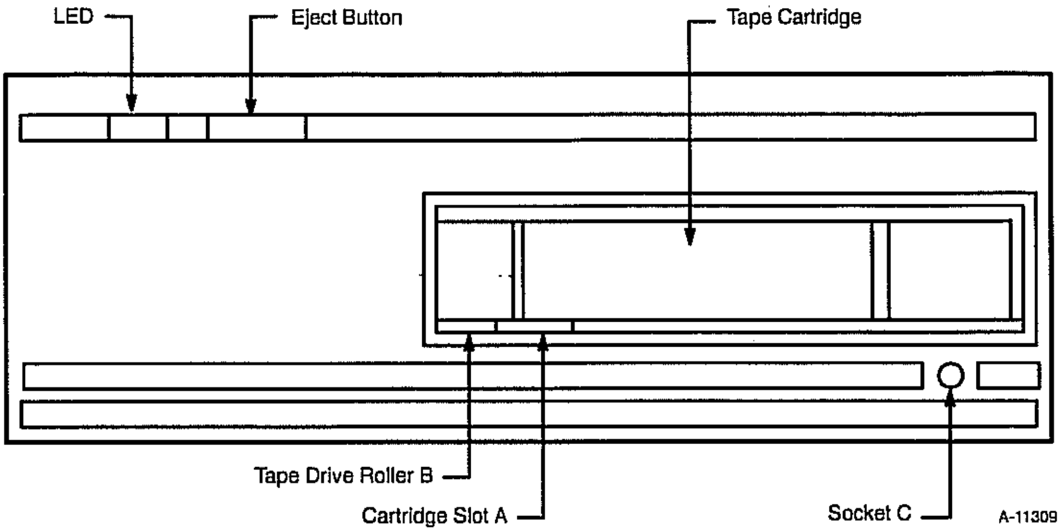

Figure 5-37. Anaconda Tape Drive

In position 1, the failure occurred with the cartridge fully loaded. The cartridge slot A is over the roller B. Proceed to Step 3.

In position 2, the failure occurred with the cartridge not fully loaded. The cartridge slot A is not over the roller B or it is recessed. Follow the instructions in Step 2, then the instructions in Step 3.

2. Insert the allen wrench into socket C in the lower right corner of the Anaconda bezel. Push the wrench in and turn it clockwise until the cartridge reaches the fully-loaded position with slot A over roller B.

CAUTION
Do not force the wrench if you feel heavy resistance, or you may damage the tape drive.

3. Fully insert the allen wrench into socket C in the lower right corner of the Anaconda bezel. Push the wrench and turn it counter-clockwise until the cartridge reaches the normal unloaded position. It takes approximately 5 full turns to unload the cartridge.

CAUTION

Do not force the wrench if you feel heavy resistance, or you may damage the tape drive.

If this procedure fails to allow removal of the tape cartridge, it will be necessary to replace the entire tape drive.

Seagate Swift ST1239N Hard Disk Drive (HD-1)

The following characteristics describe the Seagate Technology, Inc. Swift ST 1239N hard disk drive. Refer to Figure 5-38 for jumper settings.

- Used as IOS boot device
- 204.2 Mbytes formatted capacity
- High density 3.5-in. Winchester disk drive
- Uses SCSI bus driven by HK68/V30
- Can be rack mounted or installed in the DR-1 disk carrier
- Error codes are provided in the Seagate ST1239N product manual, revision F

MDB-500 DataCartridge DR-2

The following characteristics describe the MDB-500 data cartridge.

- Manufactured by MDB Systems, Inc.
- Single, removable, shock-isolated 3.5-in. disk subsystem
- Is a 1/2-height 5.25-in. chassis
- Uses SCSI differential bus
- Driven by Heurikon HK68/V30 controller board

- Drive used is the Seagate ST1239N hard disk drive
- Drive inside carrier must have jumpers correctly installed

Refer to Figure 5-38 for DR-2 jumper locations.

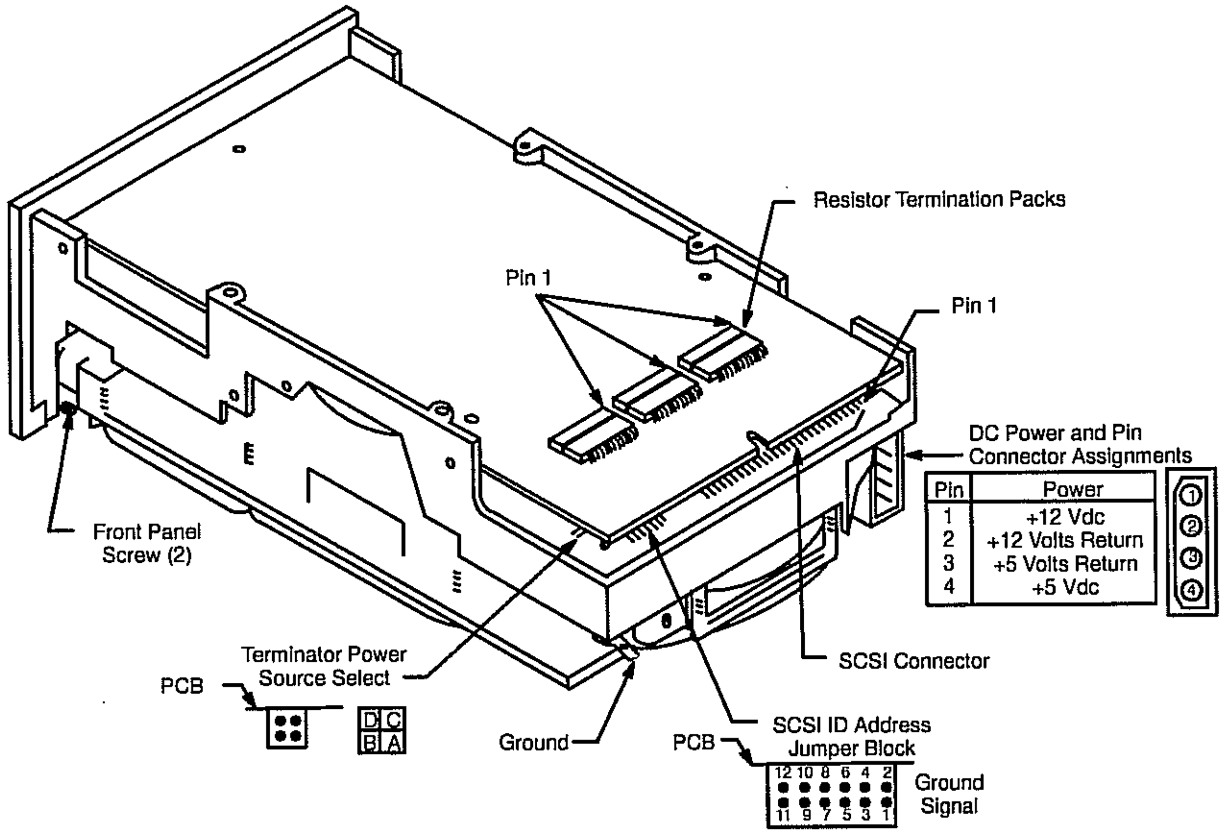

Jumper	Termination Power
A - B	Drive from Power Connector (Factory Default)
A - C	Drive from SCSI Bus
A - C and B - D	Drive from Power Connector and Provide to SCSI Bus
B - D	Provide Only to SCSI Bus

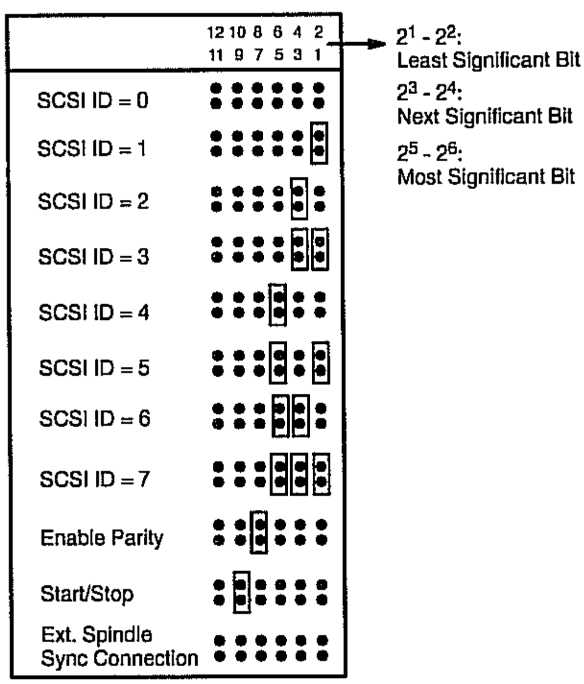

Figure 5-38. DR-2 Jumper Settings

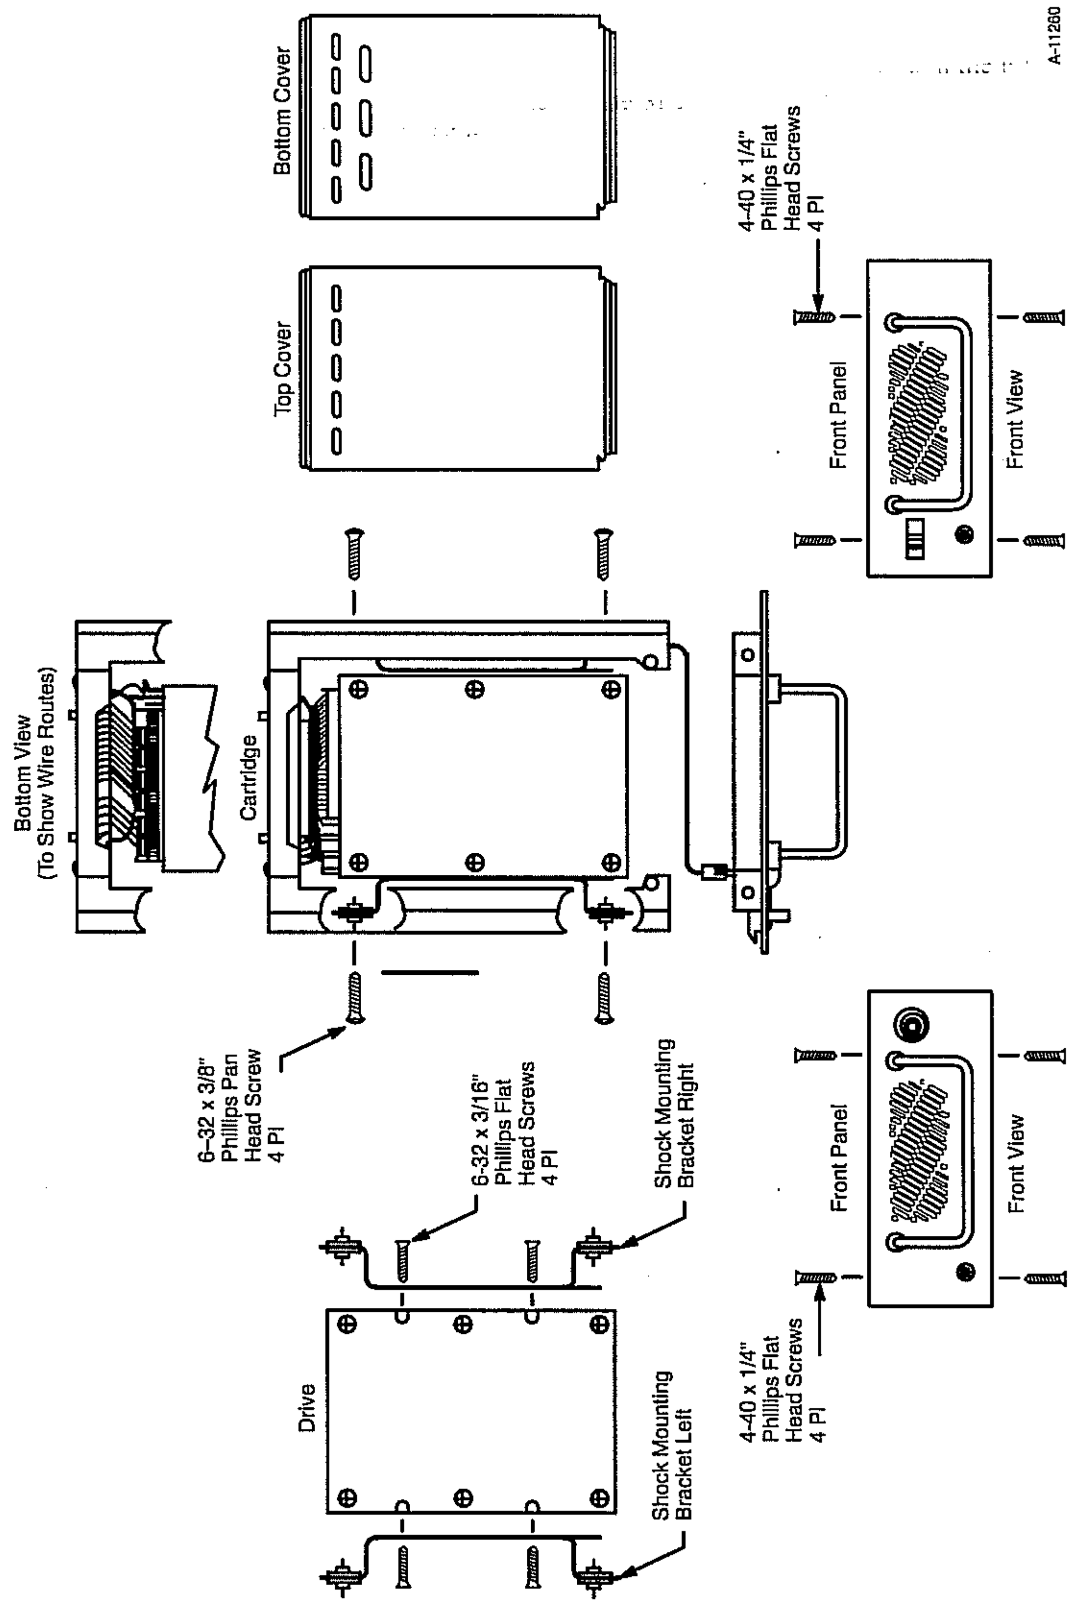

Figure 5-39. CRAY EL Series DR-2 Disk Drive

DD-5S Disk Drive

The DD-5S disk drive is a Seagate 5.25-in. hard disk drive consisting of a sealed disk assembly and supporting electronics. The DD-5S is connected to the SI-2 VME controller board.

The following characteristics apply to the DD-5S disk drive.

- 5.25-in. disk drive with 3 Gbytes of storage
- Supports SCSI-2 command set
- 20-Mbyte/s transfer rate (limited to 10 Kbytes/s with controller)
- No preventive maintenance is required

Each disk drive in the computer system must have a unique system address. Figure 5-40 shows the jumper locations and settings for the DD-5S controller board, and Figure 5-41 shows the maximum configuration of DD-5S disk drives for a CRAY EL92 or CRAY EL94 computer system and the SCSI ID jumper settings for each individual disk drive. The configuration of DD-5S disk drives for a CRAY EL98 system varies, depending on customer requirements.

Diagnostics

The following diagnostics and utilities can be used to maintain the DD-5S disk drive and SI-2 controller.

- `eddtest` – This IOS based diagnostic can be run offline or concurrently with UNICOS. If the test is run online, the onboard SCSI controller diagnostics are not invoked.

- `iob2test` – This IOS based diagnostic can be used to test any configured disk device.

- `/etc/errpt` – This UNICOS command processes disk error reports for a specific device.

- `olhpa` - This UNICOS utility (online hardware performance analyzer) can be used to format disk error reports.

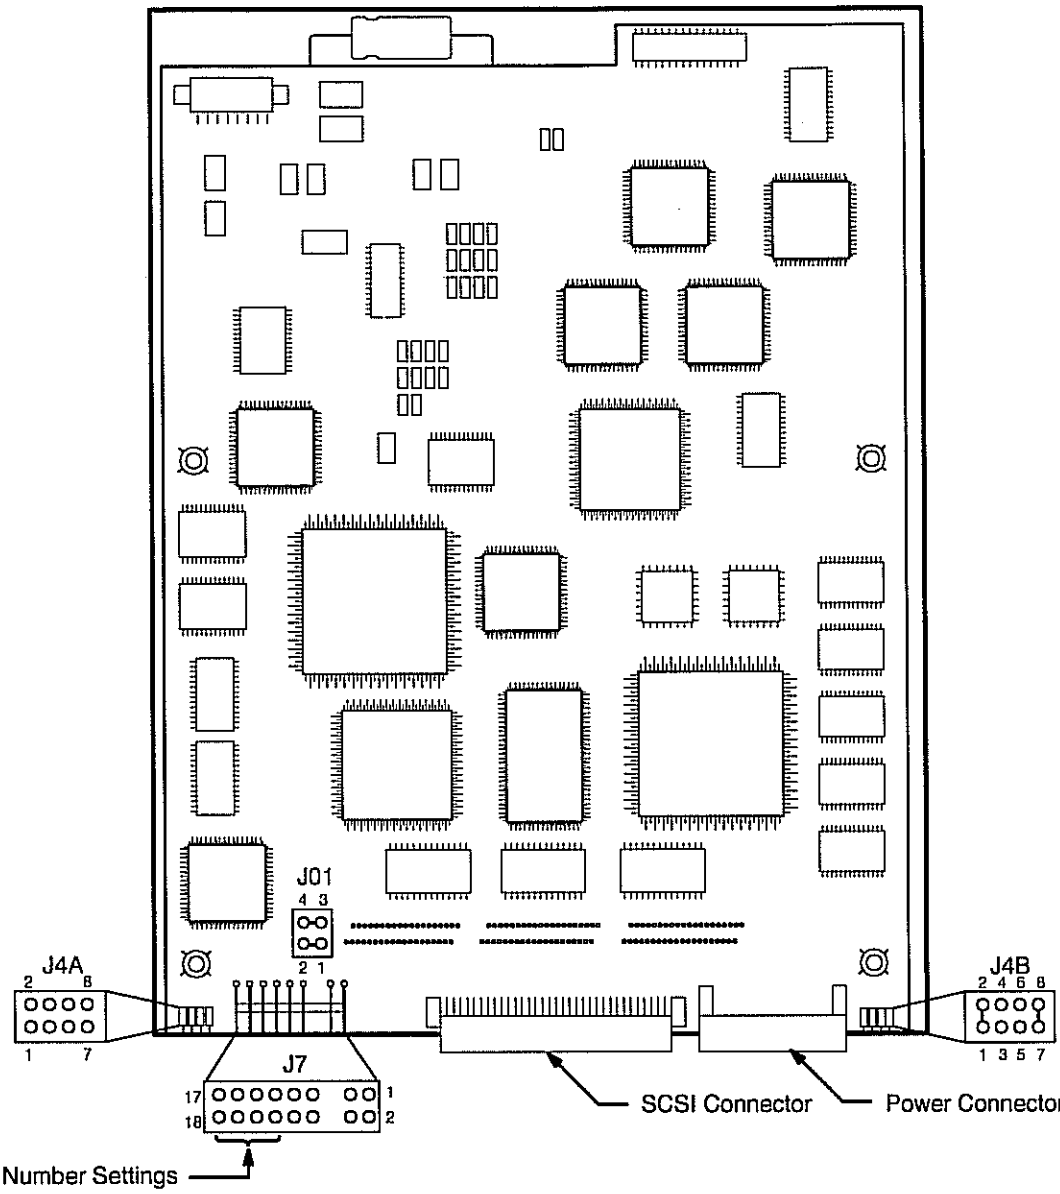

Jumper	Notes	Setting
J4A	All Jumpers	Off
J4B	Spin up Delay Option (Pins 1 and 2) Start Command Option (Pins 3 and 4) SCSI Parity (Pins 5 and 6) Sweep Cycle Option (Pins 7 and 8)	On Off Off On
J7	Only appropriate ID number jumpers should be on. All other jumpers should be off. Refer to Figure 5-41.	
J01	Termination power. Jumper pins 1–2 and 3–4 as shown when the drive is the last drive in the string. Do not jumper other drives in the string.	

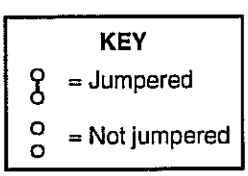

Figure 5-40. DD-5S Jumper Settings

The DD-5S disk drives (refer to Figure 5-41) are connected to the SI-2 controller by two channels. Even though all the disk drives are connected to the same physical board, under the disk device naming conventions used in the IOS operating system, each channel on the SI-2 is considered a distinct controller. For example, a DD-5S with a unit number of 0, connected to channel 0, would be designated S00. A disk with a unit number of 0, connected to channel 1 would be designated S10.

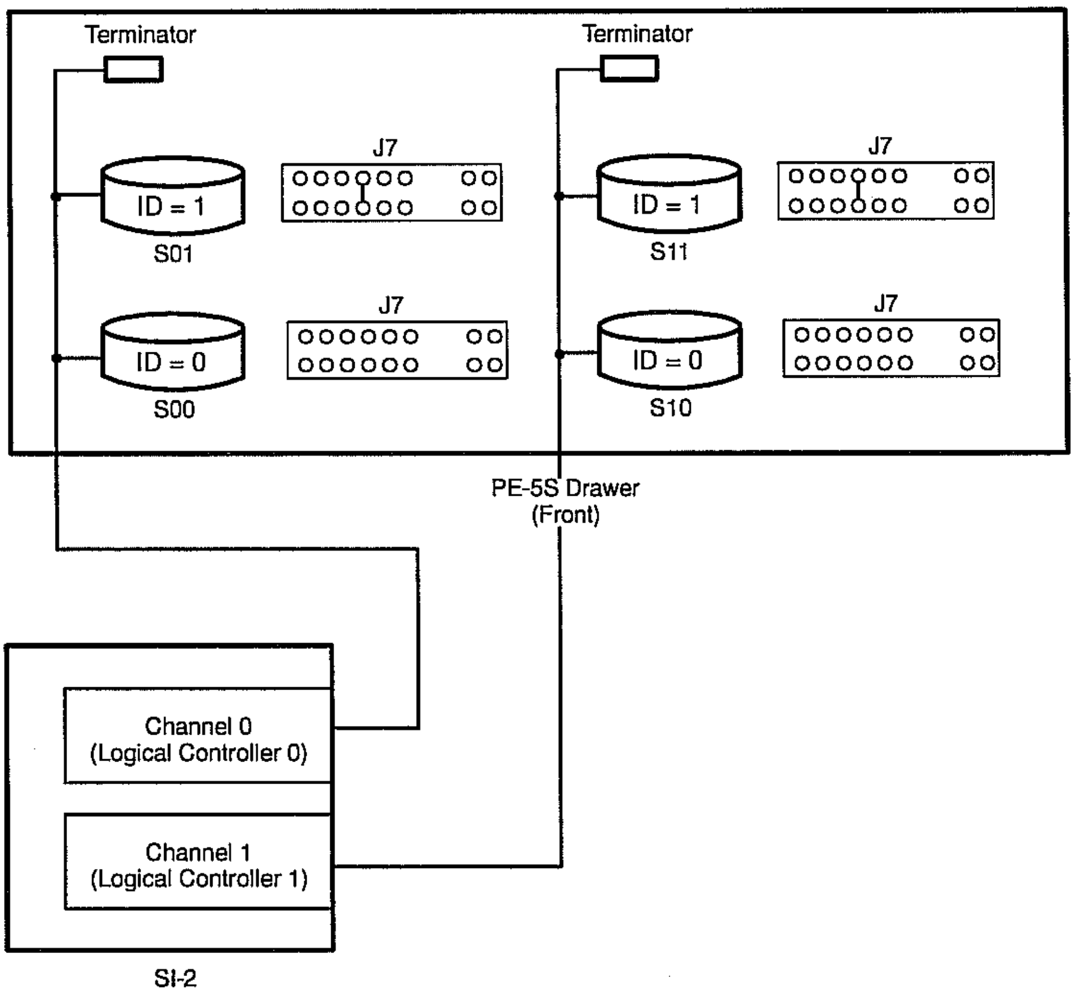

Figure 5-41. Two DD-5S Disk Drive Connections on Each Channel

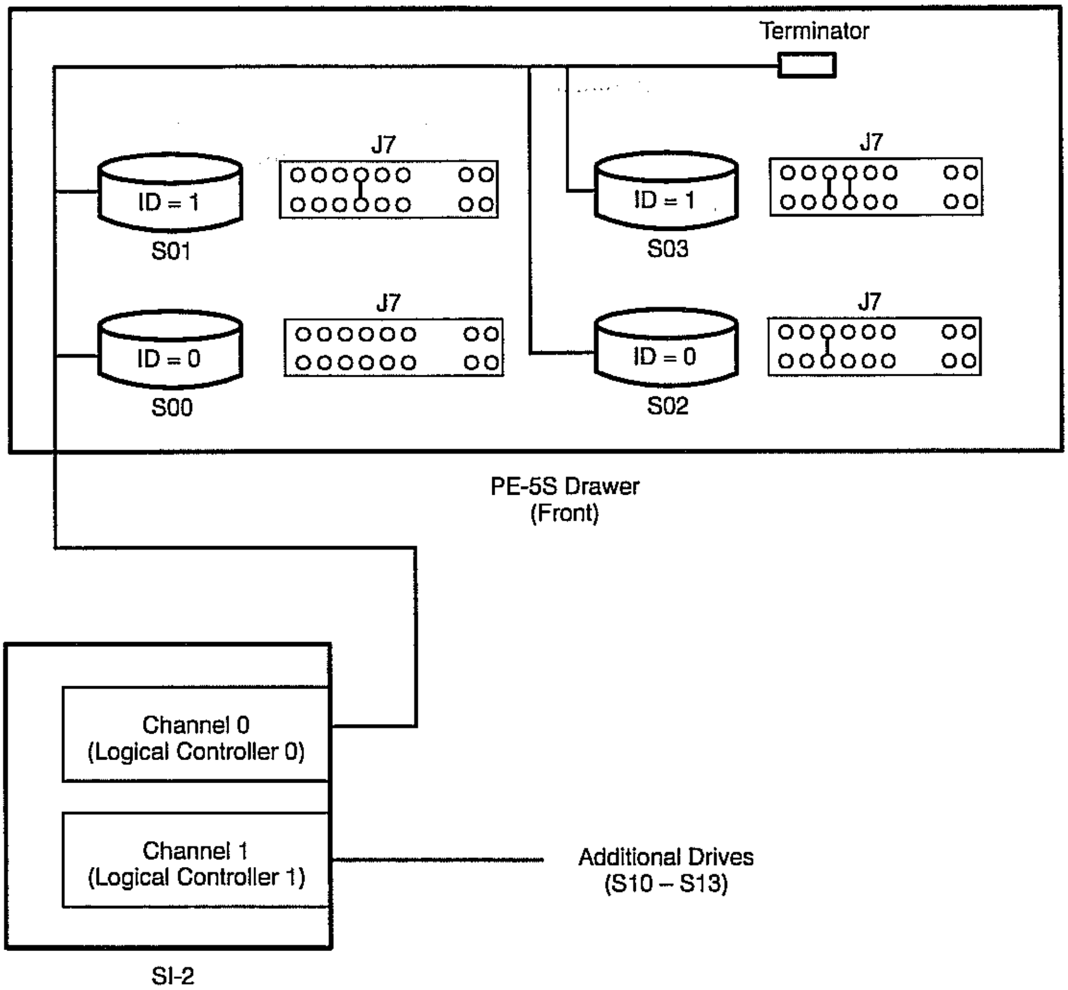

NOTE: A second SI-2 will control drives S20 – S23 on Channel 0 and S30 – S33 on Channel 1, and so on, for additional SI-2s.

Figure 5-42. Four DD-5S Disk Drive Connections on Each Channel

DS-3 Tape Drive

Newer models of CRAY EL series systems use the DS-3 digital tape drive instead of the 1/4-in. streaming cartridge tape drive (DS-2). The DS-3 uses digital audio tape (DAT) technology that has been adapted for computer data storage. The DS-3 is connected to the VME CPU by the SCSI daisy chain cable and is normally assigned SCSI ID number 1 on the VME CPU SCSI bus. The IOS hard disk drive (HD-1) is SCSI ID number 0.

The following characteristics describe the DS-3 tape drive.

- Data storage capacity is determined by the type of data cartridge used (up to 4 Gbytes)
- 510-Kbyte/s sustained transfer rate

Figure 5-43 shows the SCSI ID number jumper location and switch settings for the DS-3.

Diagnostics

The `tapetest` diagnostic is an IOS based confidence test for the DS-3 tape drive.

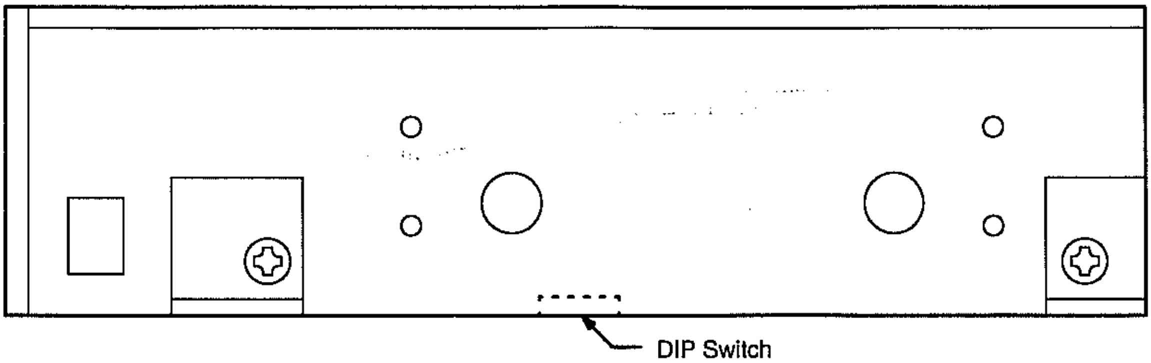

Front View with Front Faceplate Removed

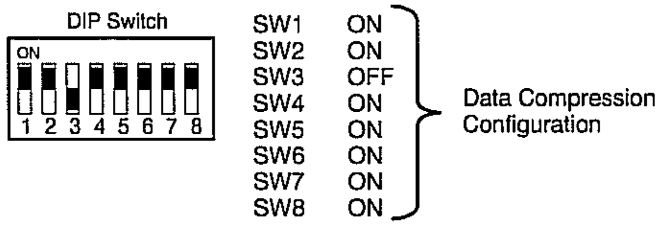

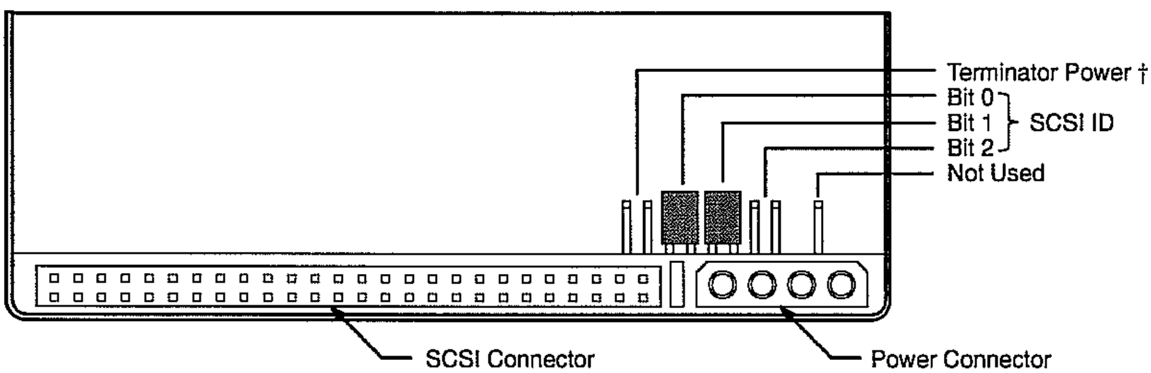

† The terminator power jumper should be installed if the DS-3 is the last device on the SCSI daisy chain cable. Usually, the DS-3 is not the last device; the HD-2 is usually the last device.

	SCSI ID Number (X = Jumpered)							
Jumper	0	1	2	3	4	5	6	7
Bit 0		X		X		X		X
Bit 1			X	X			X	X
Bit 2					X	X	X	X

NOTE: Jumper SCSI ID address 3 by installing a jumper on bit 0 and bit 1. Leave bit 2 jumper off.

Figure 5-43. DS-3 Jumper and Switch Settings

Front Panel Indicators

The front panel indicators display the drive status. The front panel has two LEDs: the Tape LED (green) shows tape activity such as tape loading or unloading and read or write activity; the Clean LED (amber) indicates either fault conditions or a need for maintenance. When the amber LED is flashing, the heads need cleaning or the cartridge is worn out.

When the amber LED illuminates steadily, either a hard fault has occurred or the self-test diagnostic has detected an error. During power-up, the drive runs a diagnostic self-test sequence; the green LED on the front panel flashes at a rate of 1 cycle per second when the diagnostic sequence is active. If the self-test fails, the amber LED illuminates continuously.

Inserting and Removing Cartridges

CAUTION

Use only cartridges labeled "DDS" in the DS-3 tape drive. Never use DAT cartridges because the media is not certified. DAT cartridges also have a different mechanical specification, which can cause them to jam in the mechanism.

Ensure that only one label is applied to the label area of the cartridge. Never use nonstandard labels, and never stick anything onto the cartridge other than in the label area; doing so may damage the tape drive.

NOTE: The DS-3 tape drive is configured to detect DDS Media Recognition System cartridges. Any cartridge that is not a DDS Media Recognition System cartridge is treated as write-protected; you can read it but not write to it.

1. Insert the tape cartridge with the exposed tape side of the cartridge entering the tape drive first.

2. The tape drive mechanism threads the tape and rewinds it completely. The tape drive then reads the first reference area on the tape to find the tape format and number of partitions.

   NOTE: If the tape is blank and rewound, the drive awaits the next command.

   NOTE: If the error rate is high, the amber Clean LED flashes to indicate that the tape heads need cleaning or that the cartridge is worn out.

3. To remove the tape cartridge, wait for the tape read or write activity to stop (as indicated by the flashing green Tape LED on the front panel), then press the unload button on the front panel.

Refer to the OEM manual for the DS-3 tape drive for more information about unloading tapes when the drive is busy or about forcing the tape cartridge to eject.

Write-protecting a Cartridge

You can write-protect cartridges by sliding the tab on the back of the cartridge so that the hole is open. With the tab in this position, data can be read from the tape but not written to it.

Maintenance

The only maintenance required for the DS-3 is periodic tape head cleaning. The tape heads should be cleaned at the following times:

- After every 24 hours of use
- When the Clean LED flashes on the front panel

Clean the DS-3 tape heads using a cleaning cartridge (Cray Research, Inc. part number 90341400). Insert the cleaning cartridge into the tape drive. The tape drive automatically loads the cartridge and cleans the tape heads. At the end of the cleaning cycle, the tape drive automatically ejects the tape cartridge. Record the date on the cleaning cartridge label so there is a record of how many times it has been used. Discard the cleaning cartridge after 25 uses or when the Clean LED remains illuminated after the tape drive has ejected the cleaning cartridge.

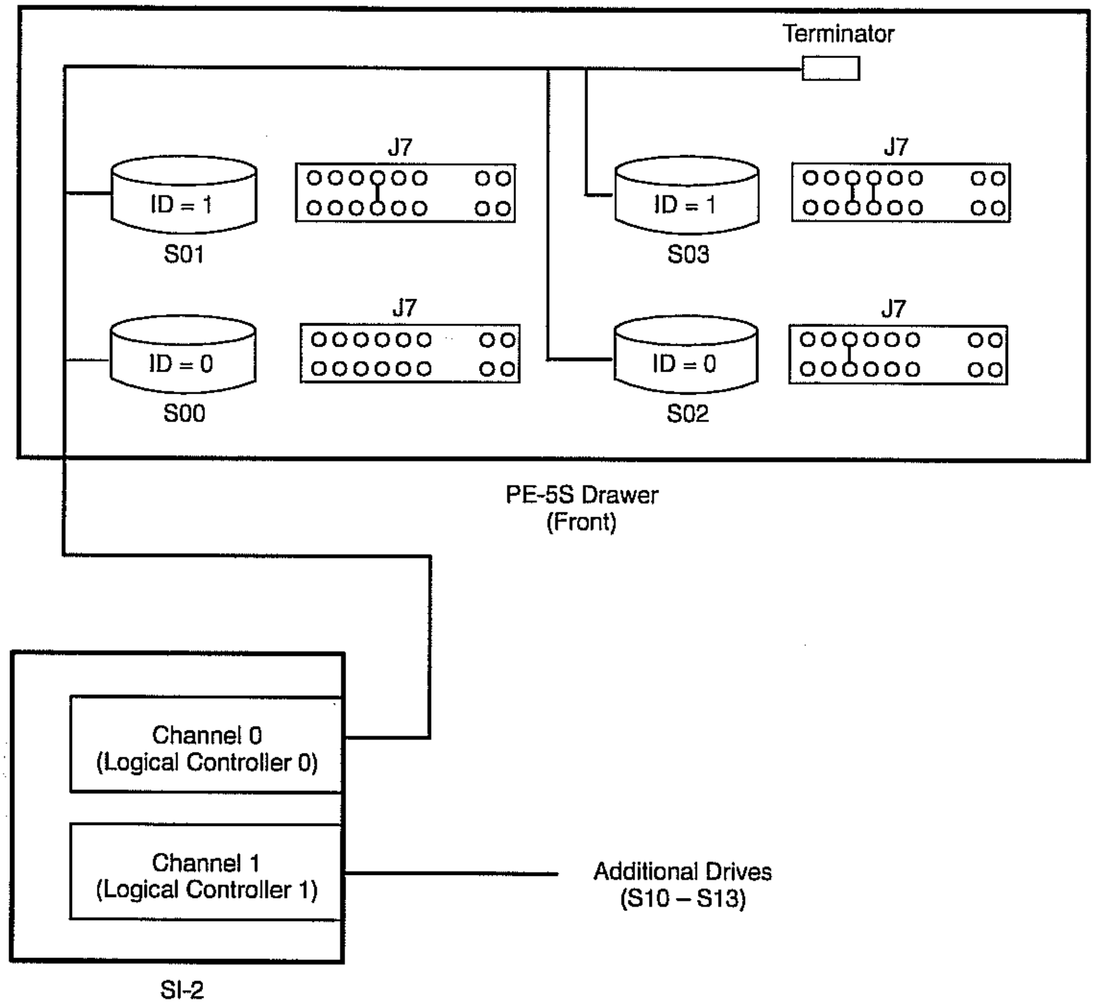

NOTE: A second SI-2 will control drives S20 – S23 on Channel 0 and S30 – S33 on Channel 1, and so on, for additional SI-2s.

Figure 5-42. Four DD-5S Disk Drive Connections on Each Channel

DD-5I Disk Drive

The DD-5I disk drive is a Seagate 5.25-in. hard disk drive that consists of a sealed disk assembly and supporting electronics. The DD-5I is connected to the DC-5I VME controller board.

The following characteristics apply to the DD-5I disk drive:

- 5.25-in. disk drive with 3.4 Gbytes of storage
- 12.4-Mbyte/s transfer rate (peak)
- No preventive maintenance is required

Each disk drive in the computer system must have a unique system address. Figure 5-43 and Figure 5-44 show the jumper locations and settings for the DC-5I controller board. Figure 5-45 and Figure 5-46 show the jumper locations and settings for the DD-5I disk drive.

Diagnostics

The following diagnostics and utilities can be used to maintain the DD-5I disk drive and DC-5I controller:

- `dd5itest` – This IOS based diagnostic can be run offline or run concurrently with UNICOS. If the test is run online, the onboard SCSI controller diagnostics are not invoked.

- `iob2test` – This IOS based diagnostic can be used to test any configured disk device.

- `/etc/errpt` – This UNICOS command processes disk error reports for a specific device.

- `olhpa` – This UNICOS utility (online hardware performance analyzer) can be used to format disk error reports.

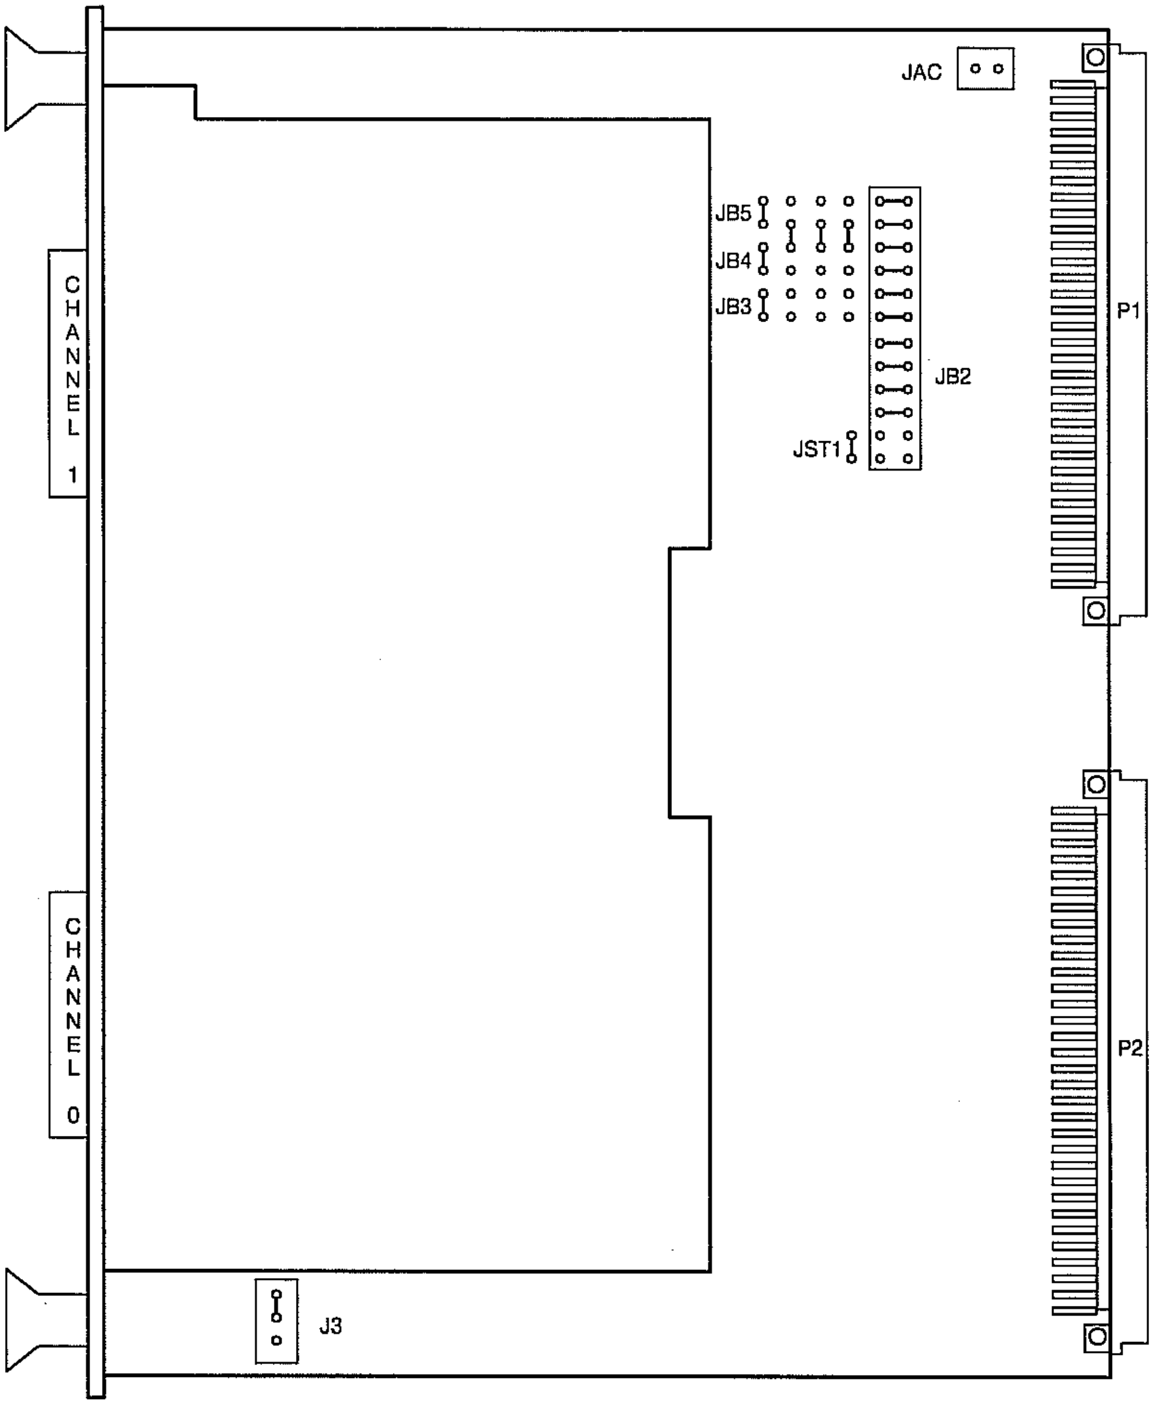

Figure 5-43. DC-5I Board Jumpers (Controller 0)

NOTE: Refer to Figure 5-44 for additional DC-5I board jumpering information.

VME Controller	0	1	2	3	4	5	6	7
Address	C000	C100	C200	C300	C400	C500	C600	C700
JB2 Jumpering	o—o	o—o	o—o	o—o	o—o	o—o	o—o	o—o
	o—o	o—o	o—o	o—o	o—o	o—o	o—o	o—o
	o—o	o—o	o—o	o—o	o—o	o—o	o—o	o—o
	o—o	o—o	o—o	o—o	o—o	o—o	o—o	o—o
	o—o	o o	o—o	o o	o—o	o o	o—o	o o
	o—o	o—o	o o	o o	o—o	o—o	o o	o o
	o—o	o—o	o—o	o—o	o o	o o	o o	o o
	o—o	o—o	o—o	o—o	o—o	o—o	o—o	o—o
	o—o	o—o	o—o	o—o	o—o	o—o	o—o	o—o
	o—o	o—o	o—o	o—o	o—o	o—o	o—o	o—o
	o o	o o	o o	o o	o o	o o	o o	o o
	o o	o o	o o	o o	o o	o o	o o	o o

Figure 5-44. JB2 Jumper Settings (DC-5I)

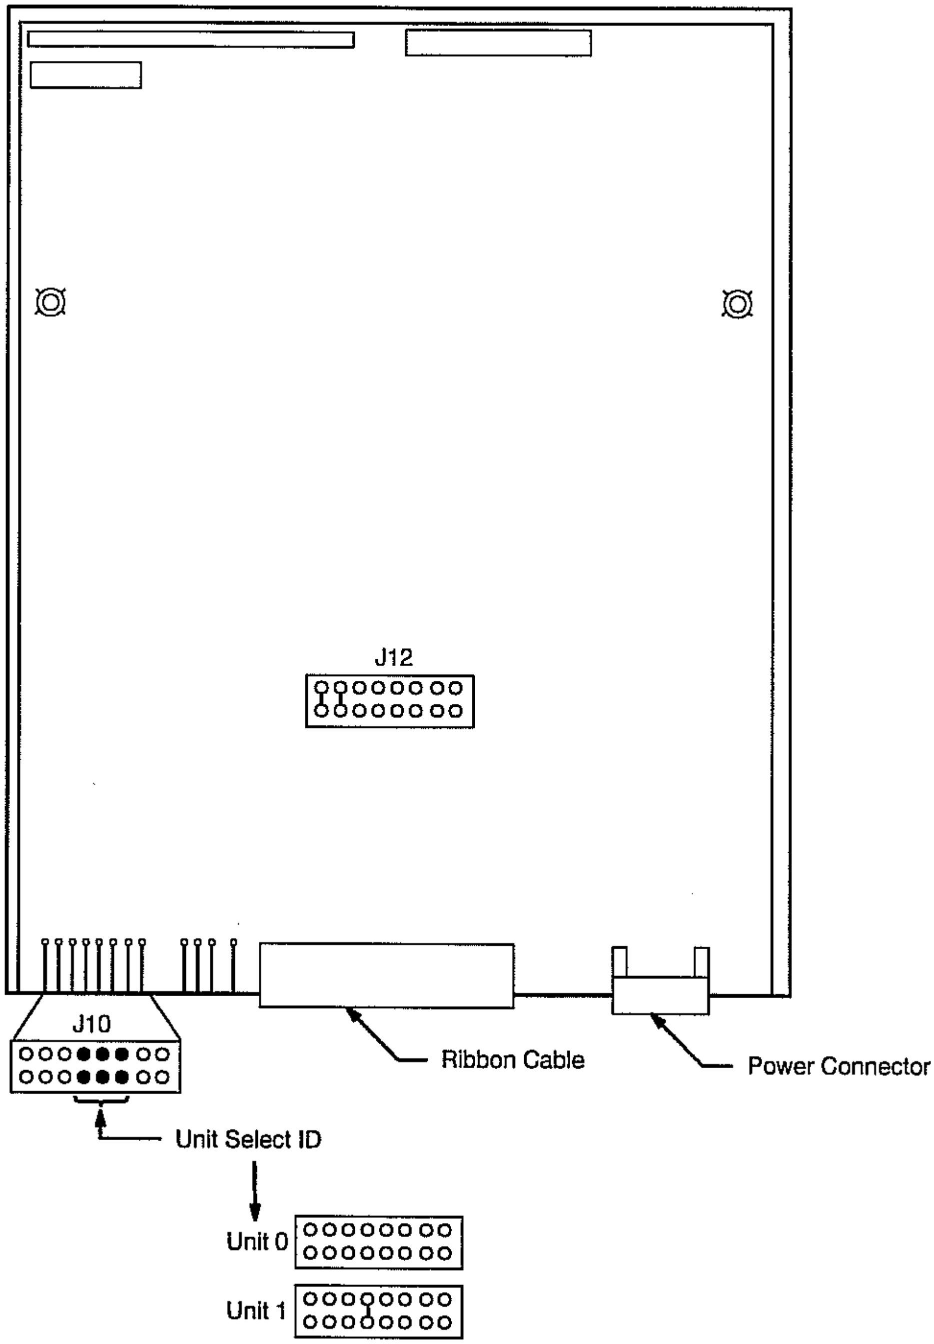

KEY
- Jumpered
- Not jumpered

Jumper	Notes	Setting
J10	Set Unit ID to Unit 0 or Unit 1	See Above
J12	Master Spindle Sync Enable Local Spin-up Enable (all other pins open)	On On

Figure 5-45. DD-5I Jumper Settings (Top View)

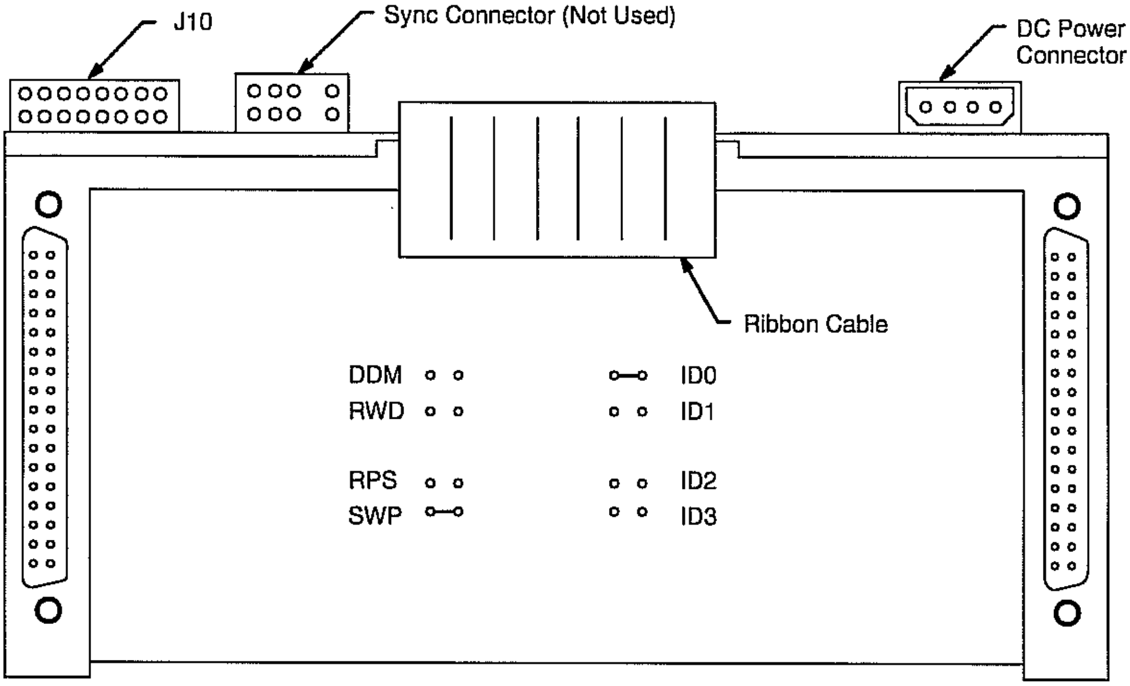

KEY
o-o Jumpered
o o Not jumpered

Jumper	Notes	Setting
DDM	Reserved	Off
RWD	Disable R/W Diagnostics	Off
RPS	Enable Short Rotational Position Sensing	Off
SWP	Enable Position Calibration on Seek	On
ID0	Microcode ID	On
ID1	Microcode ID	Off
ID2	Microcode ID	Off
ID3	Microcode ID	Off

Figure 5-46. DD-5I Jumper Settings (End View)

Two channels connect the DD-5I disk drives (refer to Figure 5-47) to the DC-5I controller. The channel number (bit 0) and the drive number (bit 1) form the drive unit number (B00, B01, and so forth). The even-numbered drives reside on channel 0; the odd-numbered drives reside on channel 1. For example, a DD-5I drive with a unit number of 0 that is connected to channel 0 is designated B00. A drive with a unit number of 0 that is connected to channel 1 is designated B01.

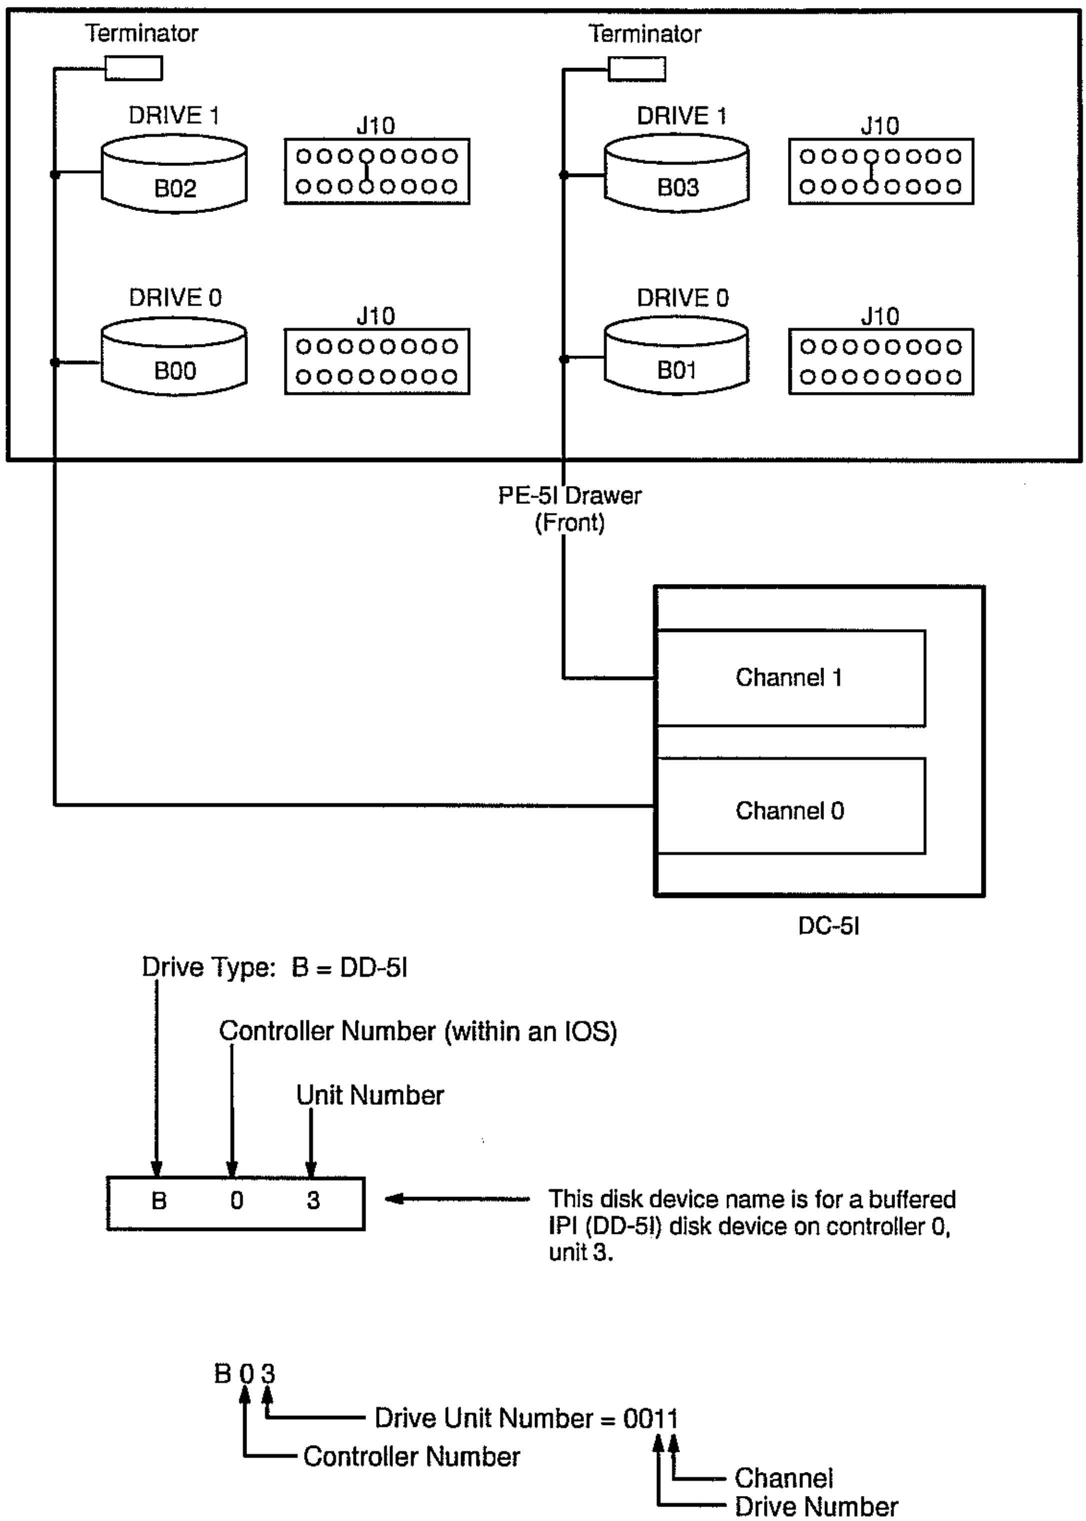

Figure 5-47. Two DD-5I Disk Drive Connections on Each Channel

DS-3 Tape Drive

Newer models of the CRAY EL series systems use the DS-3 digital tape drive instead of the 1/4-in. streaming cartridge tape drive (DS-2). The DS-3 uses digital audio tape (DAT) technology that has been adapted for computer data storage. The DS-3 is connected to the VME CPU by the SCSI daisy chain cable and is normally assigned SCSI ID number 1 on the VME CPU SCSI bus. The IOS hard disk drive (HD-1) is SCSI ID number 0.

The following characteristics describe the DS-3 tape drive:

- Data storage capacity is determined by the type of data cartridge used (up to 4 Gbytes)
- 510-Kbyte/s sustained transfer rate

Figure 5-48 shows the SCSI ID number jumper location and switch settings for the DS-3.

Diagnostics

The `tapetest` diagnostic is an IOS based confidence test for the DS-3 tape drive.

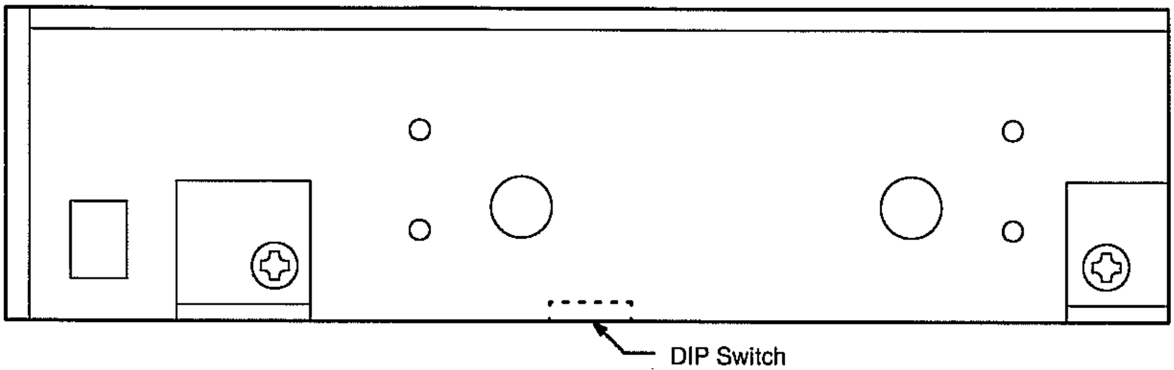

Front View with Front Faceplate Removed

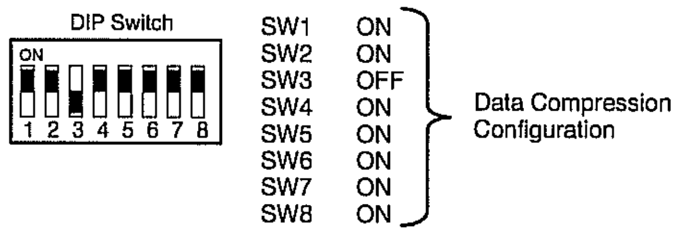

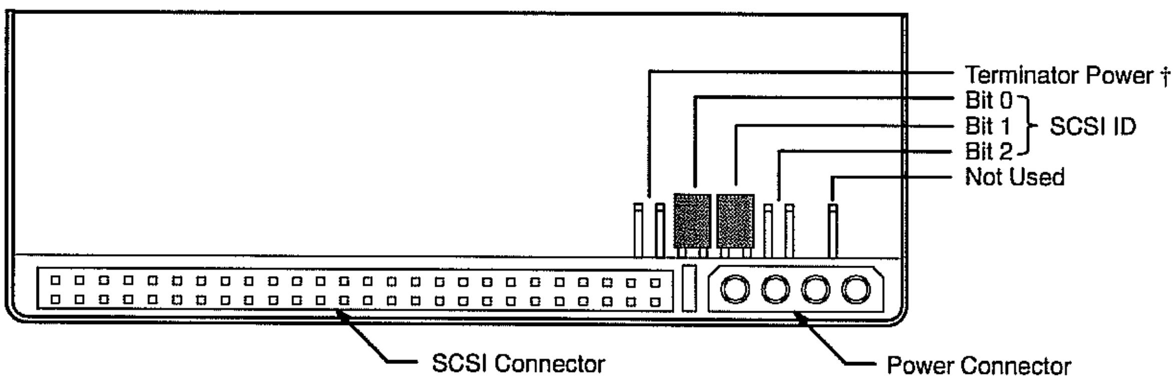

† The terminator power jumper should be installed if the DS-3 is the last device on the SCSI daisy chain cable. Usually, the DS-3 is not the last device; the HD-2 is usually the last device.

	SCSI ID Number (X = Jumpered)							
Jumper	0	1	2	3	4	5	6	7
Bit 0		X		X		X		X
Bit 1			X	X			X	X
Bit 2					X	X	X	X

NOTE: Jumper SCSI ID address 3 by installing a jumper on bit 0 and bit 1. Leave bit 2 jumper off.

Figure 5-48. DS-3 Jumper and Switch Settings

Front Panel Indicators

The front panel indicators display the drive status. The front panel has two LEDs: the Tape LED (green) shows tape activity such as tape loading or unloading and read or write activity; the Clean LED (amber) indicates either fault conditions or a need for maintenance. When the amber LED is flashing, the heads need cleaning or the cartridge is worn out.

When the amber LED illuminates steadily, either a hard fault has occurred or the self-test diagnostic has detected an error. During power-up, the drive runs a diagnostic self-test sequence; the green LED on the front panel flashes at a rate of 1 cycle per second when the diagnostic sequence is active. If the self-test fails, the amber LED illuminates continuously.

Inserting and Removing Cartridges

CAUTION

Use only cartridges labeled "DDS" in the DS-3 tape drive. Never use DAT cartridges because the media is not certified. DAT cartridges also have a different mechanical specification, which can cause them to jam in the mechanism.

Ensure that only one label is applied to the label area of the cartridge. Never use nonstandard labels, and never stick anything onto the cartridge other than in the label area; doing so may damage the tape drive.

NOTE: The DS-3 tape drive is configured to detect DDS Media Recognition System cartridges. Any cartridge that is not a DDS Media Recognition System cartridge is treated as write-protected; you can read it but not write to it.

1. Insert the tape cartridge with the exposed tape side of the cartridge entering the tape drive first.

2. The tape drive mechanism threads the tape and rewinds it completely. The tape drive then reads the first reference area on the tape to find the tape format and number of partitions.

   NOTE: If the tape is blank and rewound, the drive awaits the next command.

   NOTE: If the error rate is high, the amber Clean LED flashes to indicate that the tape heads need cleaning or that the cartridge is worn out.

3. To remove the tape cartridge, wait for the tape read or write activity to stop (as indicated by the flashing green Tape LED on the front panel), then press the unload button on the front panel.

Refer to the OEM manual for the DS-3 tape drive for more information about unloading tapes when the drive is busy or about forcing the tape cartridge to eject.

Write-protecting a Cartridge

You can write-protect cartridges by sliding the tab on the back of the cartridge so that the hole is open. With the tab in this position, data can be read from the tape but not written to it.

Maintenance

The only maintenance required for the DS-3 is periodic tape head cleaning. The tape heads should be cleaned at the following times:

- After every 24 hours of use
- When the Clean LED flashes on the front panel

Clean the DS-3 tape heads using a cleaning cartridge (Cray Research, Inc. part number 90341400). Insert the cleaning cartridge into the tape drive. The tape drive automatically loads the cartridge and cleans the tape heads. At the end of the cleaning cycle, the tape drive automatically ejects the tape cartridge. Record the date on the cleaning cartridge label so there is a record of how many times it has been used. Discard the cleaning cartridge after 25 uses or when the Clean LED remains illuminated after the tape drive has ejected the cleaning cartridge.