DSM-11 V3.3
Release Notes

AA-5461I-TC

October 1986

This is a revised document.

OPERATING SYSTEM:	DSM-11 V3.3
SOFTWARE:	DSM-11 V3.3

Software and manuals should be ordered by title and order number. In the United States, send orders to the nearest distribution center. Outside the United States, orders should be directed to the nearest DIGITAL Field Sales Office or representative.

Northeast Mid-Atlantic Region

Digital Equipment Corporation
PO Box CS2008
Nashua, New Hampshire 03061
Telephone:(603)884-6660

Central Region

Digital Equipment Corporation
Accessories and Supplies Center
1050 East Remington Road
Schaumburg, Illinois 60195
Telephone:(312)640-5612

Western Region

Digital Equipment Corporation
Accessories and Supplies Center
632 Caribbean Drive
Sunnyvale, California 94086
Telephone:(408)734-4915

digital equipment corporation • marlboro, massachusetts

First Printing, October 1978
Revised, March 1979
Revised, October 1980
Revised, October 1981
Revised, August 1982
Revised, April 1984
Revised, October 1985
Revised, May 1986
Revised, October 1986

The postpaid READER'S COMMENTS form on the last page of this document requests the user's critical evaluation to assist us in preparing future documentation.

Contents

Figures

Tables

Preface

The *DSM-11 V3.3 Release Notes* are intended for all users of DSM-11 V3.3. The *DSM-11 V3.3 Release Notes* provide:

- Information about upgrading a DSM-11 V3.0, V3.0A, V3.1, or V3.2 system to DSM-11 V3.3
- Information about enhancements to the DSM-11 system
- Information about DIGITAL services available to DSM-11 users

Version 3.3 is a maintenance release of DSM-11. All patches to DSM-11 Version 3.2 published in Software Dispatch issues up to and including December 1986 are installed in Version 3.3. The descriptions of functional changes and enhancements for Version 3.2 apply to Version 3.3 as well. There are no functional enhancements to DSM-11 for Version 3.3.

Manual Organization

The *DSM-11 V3.3 Release Notes* are divided into seven chapters.

Chapter 1 **Upgrading Your DSM-11 V3.0, V3.1 or V3.2 System to DSM-11 V3.3**

Describes how to upgrade your DSM-11 V3.0, V3.0A, V3.1, or V3.2 system to V3.3.

Chapter 2 **DSM-11 V3.2 and V3.1 Enhancements**

Describes the enhancements made to DSM-11 in V3.2 and V3.1.

Chapter 3 **Setting Up and Managing DSM-11 Volume Sets**

Explains how to set up and maintain DSM-11 volume sets.

Chapter 4 **DSM-11 V3.2 Distributed Data Processing**

Discusses Distributed Data Processing (DDP) on DSM-11 V3.2.

Chapter 5	**Transfers between DSM-11 and VAX DSM** Describes how to transfer routines and globals between DSM-11 and VAX DSM.
Chapter 6	**Patching Instructions** Provides patching instructions for DSM-11.
Chapter 7	**Digital Services** Describes services that DIGITAL provides.

Manual Conventions

These release notes use the following conventions.

Feature	Meaning
Italics	Introduces new terms.
Bold face	Emphasizes important information. In examples, shows user input.
RET	Represents the key labeled RETURN.

Related Documents

In addition to these release notes, the DSM-11 documentation set includes:

- *DSM-11 Language Reference Manual* (Order # AA-H797B-TC)
- *DSM-11 User's Guide* (Order # AA-H799B-TC)
- *DSM-11 Summary* (Order # AV-H798B-TC)
- *DSM-11 XDT Reference Manual* (Order # AA-J701A-TC)
- *Introduction to DSM* (Order # AA-K676B-TK)
- *DSM-11 BISYNC Programmer's Guide* (Order # AA-V602A-TC)

Chapter 1
Upgrading Your DSM-11 System to DSM-11 V3.3

This chapter describes how to upgrade your DSM-11 V3.0, V3.0A, V3.1, or V3.2 system to V3.3. You cannot convert a DSM-11 Version 2 system to V3.3 directly.

> **Note:** Before performing the installation, read this chapter and Chapters 1 and 10 in the *DSM-11 User's Guide*. Chapter 1 in the *DSM-11 User's Guide* contains a complete example of installation with autoconfiguration. Chapter 10 contains an example of system generation without autoconfiguration.

Take the following steps to upgrade your DSM-11 V3.0, V3.0A, V3.1, or V3.2 system to DSM-11 V3.3:

1. Back up your DSM-11 3.0, 3.0A, V3.1, or V3.2 system disk. Do not proceed until you have done this. If the upgrade fails for any reason, you could need a copy of the original to start again.
2. Follow the installation instructions in Chapter 1 of the *DSM-11 User's Guide* for magnetic tape and disk distribution.
3. Answer the installation questions from DSM-11 that begin after you boot the distribution medium.

   The following is an example of the upgrade dialogue during installation:

```
Begin DSM-11 Version 3.3 system installation

Answer with a question mark (?) any time you wish more information:

Please enter today's date [ DD-MMM-YY ] ? > 4-MAR-86

                and time [ HH:MM:SS ] ? > 16:10

Install DSM-11 on which disk unit ? > DB0
```

The upgrade procedure automatically determines what version of DSM you presently have.

```
DBO now holds a DSM-11 V3.x system disk.
Do you wish to upgrade your DSM-11 Version 3.x system to DSM-11 Version 3.3
? [Y/N] >Y

RP06 Unit 0 [ Master Label = "SYS-Y BUP-0 TUE" (23-DEC-85) ]

Loading the DSM-11 Version 3.3 system utilities onto the system disk:
```

where *3.x* is the version of DSM-11 you currently use.

(At this point the installation procedure continues and is the same as that described in Chapter 1 of the *DSM-11 User's Guide*.)

4. Answer the remaining installation questions. To obtain additional information about any installation question, type a question mark in response to that question.

   You can then do an autoconfiguration (as described in Chapter 1 of the *DSM-11 User's Guide*) or do a full system generation (as described in Chapter 10 of the *DSM-11 User's Guide*). **Using the autoconfiguration option causes a default configuration of most system generation questions.**

   At the completion of the installation, your system is a DSM-11 V3.3 baseline system.

5. The following question appears at the end of the installation procedure.

   ```
   Do you wish to proceed directly to SYSGEN ?
   ```

   Answer **Y** to run SYSGEN now. You can use any of your defined V3.1 or V3.2 configurations for V3.3 start up. You cannot use any of your defined DSM-11 V2 configurations for V3.3 start up.

   Unlike DSM-11 V3.0 (which uses 1-based device numbering), DSM-11 V3.1, V3.2, and V3.3 SYSGEN and AUTOCONFIGURE use 0-based device numbering. For example, DSM-11 V3.3 calls the first DZ11 terminal controller DZ-0. DSM-11 V3.0 calls the first DZ11 terminal controller DZ-1.

   > **Note:** Answer **N** to the question if you want to run SYSGEN later. (You can then reboot the system disk, and run SYSGEN at your convenience.)

6. You can start up your new configuration when SYSGEN is complete. To customize the configuration (by changing terminal parameters, global defaults, and other parameters), log in to the system manager's account and type the following to get access to the system definition utilities:

   ```
   >DO ^SYSDEF
   ```

   These utilities are described in Chapter 10 of the *DSM-11 User's Guide*.

7. **Be sure to use ^DISKPREP to reinitialize your back-up disks for DSM-11 V3.3.** Failure to upgrade your back-up disks could mean that future restore operations will fail or will lose data.

Chapter 2
DSM-11 V3.2 and V3.1 Enhancements

This chapter describes the enhancements made to DSM-11 in V3.2 and V3.1.

DSM-11 V3.2 Enhancements

The following sections discuss the enhancements made to DSM-11 for V3.2.

Fast Data Base Integrity Checker (^FASTIC)

DSM-11 implements global variables as *multiway trees*. These multiway trees are stored in 1024-byte disk blocks within DSM-11 volume sets. The ^FASTIC routine determines whether the pointers that connect the blocks that form the multiway trees are intact. ^FASTIC also checks intra-block integrity. If there are defects in your data base, ^FASTIC reports sufficient information for an experienced operator to repair the errors using the ^FIX utility.

^FASTIC requires that you first mount the volume set (using ^MOUNT or the start-up procedure). **Note that ^FASTIC write-locks your data base (just as ^BACKUP does) to insure that the data base remains consistent during its scan, which requires approximately 1 minute for every 20 megabytes of data.** This means that a 200 megabyte data base scan makes the data base read-only for 10 minutes.

In most cases, a single scan is sufficient. However, if ^FASTIC finds that some disk blocks are multiply allocated, it will require a second scan to make a complete report.

^FASTIC reports the following types of errors:

- Duplicate block pointer (both "down" and "right")
- Missing block pointer (both "down" and "right")
- Illegal block pointer (both "down" and "right")

- Illegal block type
- Incorrect byte offset within a directory, data or pointer block
- Corrupted map block
- Invalid character count within data or pointer block

All errors are not of equal severity. The most urgent problem would be duplicate block pointers. Duplicate block pointers should be fixed as soon as possible, and you should be careful never to kill any globals involved until the duplicate pointers have been eliminated. You should also resolve map block errors immediately.

Global Data Base Replication

Global replication is a feature that allows you to maintain multiple copies of globals within separate namespaces (UCIs). These namespaces can be on different volume sets, and on different DSM-11 systems (accessed with DDP).

Routines running in each UCI access the local copy for read-type global accesses, but perform set and kill operations on all copies of the globals. The LOCK command operates on the designated "lock master" copy of the globals.

You enable replication by first using ^REPTAB to define up to four *replication schemas*. Each replication schema defines up to 7 UCI and volume sets where globals are to be replicated. After your schemas are defined, you then use ^TRANTAB to mark individual globals for replication, and to define the "lock masters" for those globals.

DSM loads the replication schema table and UCI translation/replication table into memory automatically at start-up, and reloads them whenever you mount and dismount volume sets, both locally and remotely (with DDP).

It is the system manager's responsibility to insure the integrity and accuracy of the individual global copies. DSM cannot insure that the globals copied started out the same or that the multiple global copies remain identical. Globals can drift "out of step" when systems or communications lines go out of service, and from a variety of operator activities, such as accidental modification of replication tables.

TK50 Streaming Cartridge Tape

DSM-11 V3.2 supports the TK50 5.25 inch form-factor tape drive as one of the four DSM-11 tape drives. The TK50 is the first of DIGITAL's tape drives to use the Mass Storage Control Protocol (MSCP) controllers. All future DIGITAL products will be MSCP based.

You can install DSM-11 from TK50 tape cartridges and perform all operations that DSM-11 allows on magnetic tape with the TK50. Under most conditions, the Back-up Utility (^BACKUP) and the Fast Global Copy Utility (^%FGC) will stream.

DPV-11 BISYNC Controller

DSM-11 V3.2 supports the DPV-11 controller for binary synchronous communications (BISYNC) on Q-bus PDP-11 systems. The DPV-11 is interchangeable with the DUV-11 controller.

UCI Translation Logic Change

For flexibility and consistency, DSM-11 V3.2 has changed the way in which library (%name) globals are treated for UCI translation. DSM-11 V3.2 first attempts translation using the log-in UCI number.

If DSM-11 V3.2 finds a translation (or replication) entry in the translation table for the log-in UCI number, it does not attempt further translation. If DSM-11 does not find a translation (or replication) entry in the translation table for the log-in UCI number, it makes a second translation attempt using the library UCI number.

This feature allows you to translate "%" globals for all UCIs with a single translation table entry. You can also specify exception UCIs that translate "%" globals differently.

Non-Modem Terminal Line DTR Signal Control

You can now use parameters 4 and 5 of the OPEN and USE commands to raise and lower the Data Terminal Ready signal on terminal lines, regardless of whether the lines are marked as modem-controlled. Previously, DSM-11 ignored parameters 4 and 5 on terminal lines that were not marked as modem controlled.

Bootable and Mountable Multivolume Back-Up Sets

You can now boot and start up a back-up data base that is comprised of multiple volumes without changing the disk labels to "master."

DSM-11 V3.1 Enhancements

The following sections discuss the enhancements made to DSM-11 for V3.1. **All of these enhancements also apply to DSM-11 V3.2.**

Hardware Support

DSM-11 supports the following hardware devices or features:

- Multiple disk controllers - DSM-11 now supports multiple disk controllers of the same type. You can connect as many as three MSCP-type disk controllers and two RH-type disk controllers to a DSM-11 system.

  Note: DSM-11 requires that all disk controllers be configured at DIGITAL standard CSR addresses. Second and third MSCP disk controllers must be configured with a floating CSR address which has been determined using the DIGITAL standard convention for computing floating device CSR addresses.

- DHU11 and DHV11 terminal multiplexers - see the following section.
- DEUNA and DEQNA - DSM-11 supports the DEUNA (UNIBUS systems) and DEQNA (Q-bus systems) Ethernet controllers. See Chapter 4 for more information about DSM-11 and Ethernet.

DHU11 and DHV11 Multiplexer Support

DSM-11 supports the following multiplexer devices in addition to multiplexers already supported:

- DHU11 multiplexer - 16-line multiplexer terminal controller for UNIBUS systems.
- DHV11 multiplexer - 8-line multiplexer terminal controller for Q-bus systems, such as the MICRO/PDP-11.

These multiplexer devices have their own line parameter register bit assignments, as shown in the Table 2-1. You use these assignments in *param8* (parameter 8) of the OPEN command to open a terminal through a multiplexer device.

Table 2–1 DHV11/DHU11 Line Parameter Register Bit Assignments

Bits	Function	Assignments
0,1,2	Not used	
3,4	Character length	Bit 3 / Bit 4 0 0 = 5 bits 0 1 = 6 bits 1 0 = 7 bits 1 1 = 5 bits
5	Parity	0 = Disable 1 = Enable
6	Type parity	0 = Odd 1 = Even Ignored if bit 5 = 0
7	Number of stop bits	0 = 1 stop bit 1 = 2 stop bits for 6 - 8 bit characters 1.5 stop bits for 5 bit characters
8,9,10,11	Receiver speed	Same as transmitter assignments
12,13,14,15	Transmitter speed	Bit 12 / Bit 13 / Bit 14 / Bit 15 0 0 0 0 = 50 baud 0 0 0 1 = 75 baud 0 0 1 0 = 110 baud 0 0 1 1 = 134.5 baud 0 1 0 0 = 150 baud 0 1 0 1 = 300 baud 0 1 1 0 = 600 baud 0 1 1 1 = 1200 baud 1 0 0 0 = 1800 baud 1 0 0 1 = 2000 baud * 1 0 1 0 = 2400 baud 1 0 1 1 = 4800 baud 1 1 0 0 = 7200 baud * 1 1 0 1 = 9600 baud 1 1 1 0 = 19200 baud 1 1 1 1 = 38400 baud *

Generally, you can use the interactive utility routine ^MUX to set the line speed and other parameters where necessary. See the *DSM-11 User's Guide* for more information about the OPEN command parameters for multiplexers and terminal devices.

Escape Sequences and Type-Ahead Control

DSM-11 has always allowed type-ahead except with READ commands that include a prompt. Programmers who wished to employ type-ahead used WRITE commands to display prompts and READ commands (without prompts) for user input.

DSM-11 offers a new method of type-ahead control. To use it, issue a USE command with the format:

```
USE TTY:(::::67108864)
```

where *TTY* can be 0 or a terminal number.

This USE command disables the flushing of the type-ahead buffer by a READ prompt. Instead, this USE command lets you set a flag to indicate whether type-ahead is to be flushed or not. You set this secondary flag by issuing the following USE command:

```
(USE TTY:(::::33554432)
```

You clear this secondary flag by issuing the following USE command:

```
USE TTY:(:::::33554432)
```

If your operating environment requires no type-ahead at all, you can begin your application with:

```
USE TTY:(::::67108864+33554432)
```

This USE command indicates the new mode of type-ahead control with type-ahead now to be flushed at the start of each READ command.

If your operating environment requires type-ahead generally, but occasional flushing of type-ahead, you can begin your application with:

```
USE TTY:(::::67108864:33554432)
```

This indicates the new mode of type-ahead control with type-ahead now enabled.

On those occasions when you do not want to allow type-ahead, use the following sequence:

```
USE TTY:(::::33554432) READ X USE TTY:(:::::33554432)
```

Type-ahead is flushed at the start of the READ command, but future READ commands do not flush type-ahead. (You could actually postpone the second USE command until just before to the execution of the next READ command.)

When DSM-11 is started, these two new flags are clear. The old mode of type-ahead control is in effect by default.

After either new flag is set, that flag remains set until you explicitly clear it with an OPEN or USE command. When a terminal is logged in, these flags are in the state left by the last previous user of the terminal.

The flushing of type-ahead can interfere with the receipt of escape-sequences sent by the terminal's keyboard.

Because escape-sequences (generated by the striking of special keys) are more than one character in length, it is possible for part of an escape sequence to already have been received at the instant the in-memory buffer is flushed. As a result, the remainder of the escape-sequence is then accepted as individual characters that don't make sense to the application program.

You can avoid this problem by using the new control mode to always allow type-ahead.

VT200 Function Key Support

In addition to the support provided for the VT100 special function keys (as documented in the *DSM-11 User's Guide*), DSM-11 supports the VT200 extended function keys. Take the following steps to use the extended function keys of the VT200:

1. Use terminal set-up to Set your VT200 to VT200-mode
2. Set the DSM-11 device to "VT100" and "escape enabled" (with

```
USE 0:(:::65536+64))
```

The following table shows the codes returned in the low byte of $ZB when you press the various VT200 extended function keys:

Table 2–2 VT200 Extended Function Key Codes

Key	$ZB/256 Value	Key	$ZB/256 Value
F6	37	F17	51
F7	37	F18	52
F8	39	F19	53
F9	40	F20	54
F10	41	FIND	21
F11 (ESC)	43	INSERT HERE	22
F12 (BS)	44	REMOVE	23
F13 (LF)	45	SELECT	24
F14	46	PREVIOUS SCREEN	25
HELP	48	NEXT SCREEN	26
DO	49		

Partition Size with Routine Mapping

DSM-11 now supports 16 -Kbyte partitions with routine mapping. Formerly, partition size was restricted to 8 Kbytes when routine mapping was used. The size of each mapped routine is still limited to 8 Kbytes.

Disk-tape Cache Blocks (Buffers)

The maximum number of 1-Kbyte disk-tape block buffers you can allocate for a disk-tape cache is 2047. This means that you can allocate a total of 2 megabytes of memory for the disk-tape cache.

Fast Global Copy and Restore Utilities

DSM-11 has the following new utilities for copying globals to tape and restoring them to disk:

- %FGC - fast global copy
- %FGC - fast global tape dump
- %FGR - fast global restore

These routines are similar to %GTO and %GTI. However, you should be aware of the following differences between the new routines and %GTO and %GTI:

- %FGR is fast enough to be used as a global back-up facility. %FGR is about 10 times faster than %GTO (used in continuous mode) and about three times faster than %GTI.
- %FGC, %FGC, and %FGR do not operate on subtrees. (%GTO can operate both on entire globals and on subtrees.)
- %FGC, %FGC, and %FGR always restore the blocks of a global to disk with as much contiguity as possible.
- Unlike %GTI, the new routines do not compact global data within blocks.
- You can use the %FGR routine to change the name of a global. The new name you select must be the same length as the original name.
- You can use the %FGC routine to examine the tapes %FGC produces prior to using them for restoring a global.
- %FGC, %FGC, and %FGR do not change the global efficiency of restored globals as do %GTO and %GTI.

The following sections discuss these routines in more detail.

%FGC

The %FGC utility is an interactive utility that writes one or more globals to tape. If you are ever unsure of how to respond to any of the prompts, enter a question mark (?). The utility then provides a short explanation of your expected responses.

%FGC places the blocks for each global in the following logical order:

1. Top pointer block
2. Subsequent pointer block levels
3. Data block level

%FGC places a header before the blocks of each level, and also at the beginning of each tape.

When you execute the routine, you can specify multiple tape drives. %FGC uses those tape drives as necessary, and in the order you specify.

During execution, CTRL/Y is disabled and the globals are write-protected while being copied. The number of blocks in the pointer levels and in the data level for each global copied is reported.

Tape streaming at 100 ips (inches per second) is achieved with this utility using RP04 disk and TU80 tape drives. The speed of the copy depends on the contiguity of the global blocks on disk.

%FGC

You use %FGC to examine the tapes %FGC produces. You can display or print (on either a display terminal or line printer) any block on the tape, including the header blocks.

The %FGC utility has fast forward and backward commands and a short and full display mode. You can get a full explanation of these if you enter a question mark (?) in response to the "FGT> " prompt.

In display mode, %FGC prints the various fields within pointer blocks and data blocks in columns. If any column contains more characters than the width of the column, %FGC truncates the data and places an asterisk (*) at the end of the line to indicate truncation.

If you type CTRL/C during the display of a block, %FGC terminates that display, but does display any subsequent block output.

%FGR

The %FGR utility restores a global onto disk from a tape produced by %FGC. If you are ever unsure of how to respond to any of the prompts, enter a question mark (?). %FGR then provides a short explanation of the expected responses.

You can use %FGR to restore multiple globals from one or more tapes prepared by a single execution of %FGC. You can restore selected globals from a multiple global tape by skipping those not wanted.

%FGR requires that a dummy of each global to be restored already be present in the directory. Each dummy must have one top pointer node, and one data node. Such dummy globals can have the same names as the globals stored on tape, or another name of the same length. You can place dummy globals with a single SET operation, or with the %GLOMAN utility. You can use %GLOMAN to place the first pointer node and the first data node.

%FGR reads the header from the tape, and uses the information which it contains to adjust the pointers in the blocks as it restores them to their new positions on disk. %FGR restores the blocks in the logical order in which they were written to tape.

If you use %GLOMAN to place the dummy global, you can leave extra room between the pointer level blocks and the data level blocks to allow for global growth. The blocks are placed as close together as possible, and are contiguous if no blocks from other globals are already in the replacement space.

Before restoring the global, %FGR chains the necessary number of free blocks into linked lists using the corresponding map block entries. %FGR disables CTRL/Y so that it does not leave behind corrupted map blocks. However, a CTRL/C always causes an orderly termination of the routine.

The global is write protected during the restore. If a system crash occurs during the restore, the map blocks will be cleaned up during the subsequent mounting of the disk.

Because of pointer readjustment, %FGR runs slowly (at a speed comparable to %GTI) when it is restoring the pointer level blocks. However, during data-level restoration, %FGR goes 10 times faster. Using RP04 and TU80 tape drives, 25 ips streaming has been achieved.

Again, as with the %FGC utility, the contiguity of the global blocks on disk determines the speed at which the global can be restored. In this case, however, this contiguity is under your control. You simply need to place the global, using %GLOMAN, in a free area on disk.

Menu-Driven Utility Changes

The following sections discuss changes to the menu-driven utilities.

Utilities No Longer in the Menu

The following table lists the utilities no longer accessible through the menu and the sections of the *DSM-11 User's Guide* in which they are documented.

Utility	Documented in
Global Selector (^%GSEL)	7.2.8
Routine Selector (^%RSEL)	7.3.10
I/O Device Selector (^%IOS)	7.4.8
Header Formatter (^%HDR)	7.4.9
Menu Manager (^%MENU)	7.4.10
Modem Autodialer (^DIAL)	7.4.11

Renamed Utilities

The Extended Routine Directory Utility (^%ERD) documented in Section 7.3.12 of the *DSM-11 User's Guide* has been renamed ^%RDX.

The Autopatch Utility (^AUPATCH) documented in Section 11.3.2 of the *DSM-11 User's Guide* has been renamed ^PATCH.

Unsupported Utilities

The Global Subscript Filter Utility (^%FIND) documented in Section 7.2.9 of the *DSM-11 User's Guide* no longer exists.

Modified Utilities

You can now interactively list globals on remote DDP nodes through Global List Utility (^%G) by using a fully bracketed global reference.

New Utilities

The following sections discuss new utility routines.

New Transfer Utilities

DSM-11 has two new transfer utilities:

- Global Transfer (GLO^%GRT)
- Routine Transfer (ROU^%GRT)

Global Transfer transfers globals between processors over slaved terminal lines. Global Transfer has the following characteristics:

- Only transfers whole globals
- Does no KILLs during transfer
- Does not affect unreferenced global nodes
- Does not allow you to rename global during the transfer

Routine Transfer transmits routines between processors over slaved terminal lines. Routine Transfer does not allow you to rename the routine during the transfer.

New Routine Editing Utility

The Routine Change Every Utility (^%RCE) is a new utility that you can use to replace a common string of text in a group of specified routines.

New Remote System Status Utility

The Remote System Status Utility (^%SYR) is a new utility you can use to display the system status of remote DSM-11 nodes. The remote system must be connected to your local system by an Ethernet or DMR/DMC-11 DDP link.

On-line Help

DSM-11 contains a new on-line help facility providing information about DSM-11 commands and other aspects of the DSM-11 language and operating system.

To access the utility, type a question mark (?) after the Programmer Mode prompt (>) as follows:

```
> ?
```

DSM-11 responds with:

```
Help is available on the following DSM syntax elements:

        COMMANDS        FUNCTIONS
        OPERATORS       UTILITIES       VARIABLES

DSM help>
```

You can then enter one of the five categories listed to get a list of specific elements (such as commands). Next, enter the name of a specific syntax element (such as a command) to obtain a specific description of that element.

Modifications to the BACKUP Utility

Several modifications have been made to the BACKUP utility. These modifications provide the capability to reorganize one large MASTER disk volume into multiple smaller disks to produce a new MASTER multivolume data base volume set. The new MASTER volume set may be of a different disk type.

A new utility, ^DISKSIZ, allows you to:

- Decrease the logical disk size of a newly prepared DSM-11 disk volume.
- Increase the logical disk size of a volume that was previously decreased by ^DISKSIZ.

You can only perform these operations on disks volumes prepared as single-volume data base volume sets, SDP, or journal disks.

As an example, the following procedure describes how to reorganize a single-volume data base volume set made up of one RA81 disk volume into a two-volume data base volume set made up of two RA60 disk volumes. You can perform similar operations on other types of disk volumes.

1. Prepare the RA81 disk volume as a single-volume data base volume set using the [C]reate option of the ^DISKPREP utility.

   This operation produces a single-volume data base volume set on the RA81 with 1113 disk maps. Since there are 400 1K byte DSM-11 disk blocks per map, this is equivalent to 445,200 Kbytes of usable disk space.

2. Prepare two RA60 disk volumes as back up disks using the [B]ack up option of the ^DISKPREP utility. This operation produces two RA60 back-up volumes each with 500 disk maps (or 200000 Kbytes of usable disk space).

3. Use the ^DISKSIZ utility to reduce the RA81 logical volume size to 1000 disk maps. This produces an RA81 volume which is exactly equivalent in size to two RA60 volumes.

4. Use the ^BACKUP utility to back up the RA81 onto the two RA60 volumes. To have BACKUP produce a physical copy of the RA81 on the RA60 volumes, specify the ALL option (as opposed to the IN-USE only) when you create the back up command file.

5. Use the ^LABEL utility to relabel the RA60 volumes when back up is complete.

6. Load the RA60s, and mount them as a two-volume data base volume set using the ^MOUNT utility.

If the RA81 was originally a bootable system volume, this sequence of operations produces a bootable system volume set on the RA60 volumes. Also, you could use the ^DISKSIZ utility to decrease the logical size of the RA81 to 500 maps. This allows physical back up to a single RA60 disk volume. Later, when storage demands require more disk space, you can expand the logical size of the RA81 to its physical limits.

MSCP Disk Family Automatic Bad Block Replacement

DSM-11 incorporates support for the Bad Block Replacement (BBR) feature of certain MSCP disks, including:

- RA80
- RA60
- RA81

BBR is separate from the DSM-11 bad block replacement performed by the utility ^BBTAB. Both BBR and ^BBTAB replacements can occur on the same disk. However, starting with DSM-11 V3.1, you should have no further need of ^BBTAB for those MSCP disks that use BBR.

When a disk drive notifies your running DSM-11 system that a particular block read has had an excessive number of bits in error, DSM-11 automatically moves the data to a replacement block. DSM-11 records the replacement block number on the disk's Replacement and Caching Table (RCT) so that future transfers use the new block location.

Except for a diagnostic message sent to the error log as an audit trail, the BBR activity is not noticeable from your application. The original disk transfer may even have succeeded, since the disk's Error Correcting Code (ECC) can detect and correct up to eight 10-bit errors on each 512-byte transfer.

When there are too many bit-errors to be corrected by the disk ECC, DSM-11 still copies the block to a new location. DSM-11 then sets a forced error flag to indicate that the block holds incorrect data. This forced error flag causes normal disk read attempts to fail and generate a <DKHER> error. You can clear these forced errors only by taking the following steps:

1. Reading the offending block with ^FIX
2. Finding and correcting the bit-errors within the block
3. Filing the block

The Replacement and Caching Table (RCT) must be valid for DSM-11 BBR to work correctly. Thus, DSM-11 performs an RCT integrity check the first time the disk is brought on-line to verify that the bad block information on the disk is valid. If the RCT is corrupted, DSM-11 logs a CARETAKER error.

You may find this CARETAKER error occurring immediately after upgrading to DSM-11 since the RCT corruption could have existed during operation of your DSM-11 V3.0 system. If you encounter the error, back up your disk data and reformat the MSCP disk.

MSCP Disk Family Error Processing

The MSCP series of disks includes such disks as RD51, RD52, RA80, RA81, and RA60. DSM-11 provides more detailed error reporting and automatic printouts of error problems for these disks. These error reports now include BBR diagnostic messages where applicable.

Increased Number of Jobs

DSM-11 V3.1 and V3.2 allow as many as 100 job partitions. Previous versions allowed as many as 63 job partitions.

Increased Number of Terminals

DSM-11 V3.1 and V3.2 allow as many as 160 multiplexer (DZ11, DHV11, and so forth) terminal lines. Previous versions allowed as many as 128 multiplexer lines.

Spooling Pauses during Backup

The spooling feature now recognizes the disk read/write inhibit flags set during disk backup. Spooling automatically pauses when these flags are set.

Delete Key as Terminator

You can now select the Delete key (ASCII 127) as a terminator for terminal READ commands. To select the Delete key, specify bit 17 in a terminal device OPEN *param9* (parameter 9).

Although bit 17 logically selects CTRL/Q as a terminator, CTRL/Q cannot practically be a terminator. The DSM-11 terminal driver intercepts CTRL/Qs for flow control. Thus, the CTRL/Q bit position in parameter 9 has been preempted for use as the logical equivalent of delete.

Virtual Terminals

You can specify certain terminal lines that are outgoing connections to other computers as "virtual terminals." These terminal lines can be modem controlled. A new DSM-11 utility, ^SETHOST, lets you use these lines. If you use DIGITAL modems, you can use autodialing on the lines.

XDT Debugger Changes

Because of changes to the in-memory structure of DSM-11, the system debugger, XDT, has been modified. As a result, you must now specify the addresses of the following system modules as offsets:

- KIOD
- MTD
- KSPOOL
- SDP
- DMC
- JOBCOM
- DDP
- USRDRV

To specify offset addresses for these system modules, use the following syntax:

```
#cc,offset
```

where: cc = the first two characters in the name of the system module.

For example, you open location KIOD+42 by entering:

```
#KI,42
```

You set breakpoint 0 to that location by typing:

```
#KI,42;0B
```

While they are executing, each of the system modules is mapped to Kernel Page 4. Therefore, when you are executing within one of these system modules, XDT shows locations as the offset + 100000. For example, when XDT encounters the breakpoint shown in the previous example, it prints:

```
0B;100042
```

Selective Collating Sequence for DDP

You can now use the ^TRANTAB utility to specify the collating sequence to be used by global references through DDP. Previous versions of DSM-11 require the default collating sequence for global references through DDP.

Chapter 3
Setting Up and Managing DSM-11 Volume Sets

This chapter explains how to set up and maintain DSM-11 volume sets. The *DSM-11 User's Guide* contains additional information about DSM-11 volume sets.

DSM-11 Volume Sets

With DSM-11 V3.1 and V3.2, you can organize data into volume sets made up of one or more physical disks (disk volumes). Although, in practice, a volume set is often made up of only one volume on one disk, a volume set can be made up of disks of different types and sizes.

DSM-11 volume sets have the following characteristics:

- A DSM-11 system can have up to eight volume sets.
- Each volume set has a 3-character volume set name such as AAA or SYS.
- The first volume set on any DSM-11 system is the system disk.
- Each volume set can have its own set of UCIs with up to 30 UCIs per volume set. By using volume sets you can separate different groups of users into different volume sets.
- The system sees a volume set as a unified data base, even if it is composed of several disk volumes.

- Within a volume set, globals can grow throughout the volume set. A global can grow from one disk to another, so that one global can be contained on two different disks. Thus, a volume set approach can be useful if you expect to have a very large global that cannot be contained on one disk.
- You can set up a DSM-11 system so that any UCI can automatically and invisibly (to the user) access globals in other UCIs in other volume sets. These volume sets can be located on the same system or other systems connected through Distributed Data Processing (DDP).

To create any volume set, you must first specify a single disk volume. Later, you extend the volume set by adding disk volumes. The rest of this chapter describes how to set up, extend, and maintain volume sets. Chapter 4 describes how you can access volume sets on remote systems.

The System Volume Set (System Disk)

The DSM-11 installation procedure automatically creates the first volume set, called the system volume set. The installation procedure takes the following steps to create the system volume set:

1. Makes a back up copy of the distribution kit if you request this to be done.
2. Runs a ^DISKPREP procedure (similar to the DSM-11 utility, ^DISKPREP) on the system disk. This procedure:
   a. Prompts for a 3-character volume set name, such as (but not necessarily) SYS.
   b. Prompts for a 22-character label.
   c. Checks for bad blocks.
   d. Initializes the disk.
   e. Mounts the volume set.
   f. Installs the system manager UCI, MGR, as UCI number 1 on the volume set.
   g. Copies DSM-11 utilities.
   h. Copies the system globals.
   i. Copies the system image.

After this point, you can run a system generation (SYSGEN) procedure.

See the *DSM-11 User's Guide* for examples of the installation procedure, the autoconfiguration procedure (automatic SYSGEN), and the manual SYSGEN procedure.

System Volume Set Characteristics

The system volume set the following characteristics:

- Is initially composed of one disk, the system disk.
- Can be extended to include additional disk volumes.

  Note: The MGR UCI cannot grow beyond the first disk volume. This means that the first disk volume is always a self-contained DSM-11 system. You can boot the system from the first disk volume alone, even if the system volume set has been extended to more than one volume.

- Is always the initial volume set in the Volume Set Table. The system volume set carries a designation of S0 in this table.
- Can have UCIs and globals added as long as there is disk space available. system managers keep data base globals separate from the system disk by placing them on separate disk volumes or separate volume sets.

System Generation and Volume Set Support

The support that DSM-11 provides to volume sets depends on the system generation process you use.

In the autoconfiguration process, usually used during an installation procedure, DSM-11 creates a standard configuration that supports seven volume sets, in addition to the system volume set.

In the manual SYSGEN process, you must specifically request that support for volume sets. If you do not do this, you are restricted to one volume set, the system volume set.

The manual SYSGEN asks two relevant questions. The first is in Part 7 of SYSGEN:

```
Include support for MOUNTABLE DATA BASE VOLUME SETS [Y or N] ? <Y>
```

The second question is in Part 9:

```
Enter the number of ADDITIONAL mountable DATA BASE VOLUME SETS <7> ?
```

This number is the number of volume sets in addition to the system volume. The maximum number is seven, for a total of eight volume sets in a given configuration.

Creating Additional Volume Sets

Take the following steps to create a new volume set.

1. Run ^DISKPREP to create the volume set.
2. Run ^MOUNT to mount the new volume set.
3. Run ^%STRTAB to verify that the volume set is properly mounted.

4. Run the EXTEND option of ^DISKPREP to extend the volume set.
5. Run ^DISMOUNT and ^MOUNT to remount the extended volume set.
6. Run ^UCIADD to add UCIs to the volume set.
7. Run ^JRN and ^SDP to allocate system space.
8. Run the SPOOL option to allocate spooling space.
9. Run ^%GLOMAN to allocate globals.

The following sections describe these steps in more detail.

Running ^DISKPREP to Create the Volume Set

First, place the disk in an available disk drive. Then, run the ^DISKPREP utility to create a DSM-11 data base volume set. The ^DISKPREP utility takes the following steps:

1. Prompts for the 3-character name of the volume set.
2. Checks for bad blocks by taking the following steps:
   - Runs a test for bad blocks if you request the utility to do so. To do the test, the utility reads and writes bit patterns to the blocks.
   - Looks for any bad blocks contained in the factory-written bad block table.
   - Prompts you if you know of any bad blocks.
3. Prompts you for a 22-character label for the volume set.
4. Writes the home block, which includes such information as:
   - Volume set name
   - Volume set label
   - Number of disk volumes
   - For each disk volume:
     - Type of disk
     - Number of maps
     - Starting map
5. Initializes all blocks to empty.

Note: You can also use ^DISKPREP to create specialized volume sets used for Sequential Disk Processing (SDP), spooling, or journaling. Normally, you do not need to do this because you can allocate separate SDP, spooling, or journaling space on a system volume or on a normal data base volume set.

The following is a sample ^DISKPREP dialog:

```
>DO ^DISKPREP

Prepare a disk on which disk unit? > DRO

Select usage for this disk:

[B]ackup
[E]xtend a currently mounted volume set
[C]reate volume 1 of a new volume set
[S]DP/SPOOL/JOURNAL mountable disk

Select one ? > C

Every Volume set must have a unique 3-character identifier.
Enter the 3-character name for the Volume set > AAA

Prepare disk unit DRO

Do you wish to run a comprehensive test for bad blocks
on this disk? [Y/N] > N

Adding to DSM bad block table the following blks from factory-written
table:

         DSM relative blk # 19164
         DSM relative blk # 19244
         DSM relative blk # 41829

RM02 Unit 0 Bad Block Table

DSM relative block #     19164
                         19244
                         41829

Do you know of any other bad blocks on this disk? [Y/N] N

What would you like the new label of this disk to be?
(up to 22 characters, enclosed in quotes) ? > "XX1"

Now initializing DRO for use as a DSM-11 volume...
DRO is now a mountable volume. D ^MOUNT to make the volume accessible.
```

Running ^MOUNT to Mount the New Volume Set.

You can use ^MOUNT to mount a volume set for the following purposes:

- As a DSM-11 data base volume set (containing UCIs and globals)
- As a Sequential Disk Processor (SDP) volume set that you can use only for SDP purposes
- For view-only purposes in which you can look at or move blocks on the disk using the VIEW command

The following is a sample ^MOUNT dialog:

```
>DO ^MOUNT

Mount a volume on which disk unit ? > DRO
DRO has a DSM V3.2 MASTER label: "XX1"
DRO is volume 1 of volume set AAA which has 1 disk.

Mount DRO for access via: [S]DP, [V]IEW, or [D]ATABASE ? D
What name do you wish to use for this volume set <AAA> ? RET

Mounting AAA as Volume Set number S1
Volume 1 on DRO has 65600 blocks 65434 available
Total in volume set: 65600 blocks 65434 available

Volume set AAA has now been mounted.
```

Running ^%STRTAB to Verify the Mount

You can now verify that the volume set is properly mounted and named by running the Volume Set Table utility, ^%STRTAB. The following is an sample ^%STRTAB dialog:

```
>Do ^%STRTAB

Output Device ? > 0

                      Volume set      Disk tape       No. of
                      Name            and unit        maps

Volume set S0         JHM             DL0             25

Volume set S1         AAA             DR0             164

Volume set S2         no volume set mounted

Volume set S3         no volume set mounted
```

Using ^DISKPREP to Extend the Volume Set

After you create a volume set, you can use ^DISKPREP to extend the volume set by adding disk volumes to it. DSM-11 considers the extended volume set as one unit. The additional disk volumes can be different kinds of disks from the first disk volume of the volume set.

First, make sure that the volume set you want to extend has been created and mounted. Then, physically place a disk in an available disk drive, and run ^DISKPREP. The ^DISKPREP utility:

1. Initializes the disk volume, including checking for bad blocks.
2. Displays the Volume Set Table.
3. Adds the disk volume to the volume set as volume two of the volume set.

The following is a sample ^DISKPREP dialog:

```
>DO ^DISKPREP

Prepare a disk on which disk unit? > DR1

Select usage for this disk:

[B]ackup
[E]xtend a currently mounted volume set
[C]reate volume 1 of a new volume set
[S]DP/SPOOL/JOURNAL mountable disk

Select one ? > E

DSM-11 Mounted Volume Set Descriptor Table

                        Volume set      Disk tape       No. of
                        Name            and unit        maps

Volume set S0           JHM             DL0             25

Volume set S1           AAA             DR0             164

Volume set S2           no volume set mounted

Volume set S3           no volume set mounted

Add this disk to which Volume set ? > AAA

Prepare disk unit DR1

Do you wish to run a comprehensive test for bad blocks
on this disk? [Y/N] > N

RM02 Unit 1 Bad Block Table

The Bad Block Table is empty.

Do you know of any other bad blocks on this disk? [Y/N] > N

What would you like the new label of this disk to be?
(up to 22 characters, enclosed in quotes) ? > "XX2"

Now initializing DR1 for use as a DSM-11 volume...
DR1 is now volume 2 of structure S1, however, you must
dismount and remount the structure to make the volume accessible.
```

Using ^DISMOUNT, ^MOUNT, and ^%STRTAB to Remount the Set

You use ^DISMOUNT, ^MOUNT, and %STRTAB to dismount and remount the volume set after you add a disk volume to it.

Use the ^DISMOUNT utility to dismount the volume set. ^DISMOUNT first lists all of the mounted volume sets. In this case there is only one mounted volume set, AAA.

The ^DISMOUNT utility then asks which volume set to dismount.

The following is a sample ^DISMOUNT dialog:

```
>DO ^DISMOUNT

Dismount disk volume or Volume Set:

AAA - mounted Volume Set S1 (DRO)

Dismount which ? > AAA

Attempting to dismount AAA...
Volume Set AAA is dismounted
```

You use the ^MOUNT utility to remount the volume set. When you run ^MOUNT you need only specify the disk drive for the first volume of the volume set.

The following is an example of the ^MOUNT utility.

```
>DO ^MOUNT

Mount a volume on which disk unit ? > DRO
DRO has a DSM V3.2 MASTER label: "XX1"
DRO is volume 1 of volume set AAA which has 2 disks.

Mount DRO for access via: [S]DP, [V]IEW, or [D]ATABASE ? D
What name do you wish to use for this volume set <AAA> ? RET

Mounting AAA as Volume Set number S1
Volume 1 on DRO has 65600 blocks 65434 available
Volume 2 on DR1 has 65600 blocks 65434 available
Total in volume set: 131200 blocks 130869 available
```

You use ^%STRTAB to list the remounted volume set. By running ^%STRTAB, you can see how the two-volume volume set is listed in the Volume Set Table.

The following is an example of ^%STRTAB:

```
>DO ^STRTAB

Output Device ? > 0

DSM-11 Mounted Volume Set Descriptor Table

                        Volume set      Disk tape       No. of
                        Name            and unit        maps
Volume set S0           JHM             DL0             25
Volume set S1           AAA             DR0             164
                                        DR1             164
Volume set S2           no volume set mounted
Volume set S3           no volume set mounted
```

Using ^UCIADD to Add UCIs to a Volume Set

To add UCIs, run ^UCIADD from the system manager's UCI. Each volume set can have as many as 30 unique UCIs.

The following is a sample ^UCIADD dialog:

```
>DO ^UCIADD

Inspect/Add UCI data for which volume set ? > AAA

*********Inspect/Add UCIs for Volume Set AAA********

Enter disk data as DDU:M:BL

         Where DD = disk type
               U  = unit no.
               M  = map no.
               BL = block no. within map
or hit  CR   to accept default values.

For "UCI name" enter 3 alphabetic chars., or hit <CR> if done adding
UCIs.

             GLOBAL      NEW GLOBALS  ROUTINE     NEW ROUTINE  NEW GLOBALS
UCI_# NAME   DIRECTORY   POINTER AREA DIRECTORY   GROWTH AREA  DATA AREA
----------   ---------   ------------ ---------   -----------  -----------
1     JSS    DR0:0:200   DR0:0:201    DR0:0:202   DR0:0:203    DR0:0:204
             200:S1      201:S1       202:S1      203:S1       204:S1

1 UCI added to disk,
1 UCI added to memory.
```

In this example, the carriage return was pressed under each column. The values shown are the ones assigned by the system. You can enter specific values under each column to place UCIs at particular points on the disk volumes in the volume set.

The *DSM-11 User's Guide* discusses some issues involved in placing UCIs and globals on disks. Placing UCIs and globals on a volume set involves many of the same issues. Remember that globals can grow continuously from one disk to another if the two disks are in the same volume set.

Using ^JRN and ^SDP to Allocate System Space

After creating a volume set, you can allocate system space, including:

- Spooling space - usually used for printing queues.
- Journaling space - used for journaling.
- SDP space - used for sequential files.

You cannot use any system space for globals.

Each space has several utilities you can use to:

- Allocate space
- Deallocate space
- See what space is currently allocated
- Perform space-dependent functions

You can access these utilities as follows:

- For Spooling - through the spool option of the System Utilities Menu.
- For Journaling - through the journal option of the System Utilities Menu.
- For SDP - through the SDP option of the System Utilities Menu or by typing:
  ```
  >DO ^SDP
  ```

Keep the following points mind about allocating system space:

- You allocate space by using the appropriate utilities.
- You allocate space in disk units and 400-block maps, such as DR1:20 (space beginning at map 20 on disk unit DR1).
- You can allocate only one space for spooling. Spooling space is often allocated on the system disk (if space is available).
- You can allocate more than one space for SDP or journaling space.
- You can allocate system space on a data base volume set that also contains UCIs and globals.
- You can use the ^DISKPREP utility to create specialized disks to be used only for system space.

The following sample dialog shows how to allocate space for SDP. In this example, space is allocated on DR1 (the second volume of the volume set AAA defined earlier in this chapter).

```
>DO ^SDP

Allocate, deallocate, or show SDP spaces [A, D, or S] ? > A

Ready to allocate contiguous disk maps for SDP

Do you wish a fast disk-block tally [Y/N] ? > N

Enter the disk type, unit number, and starting map number in format:
    DDU:M
    where DD = Disk type (DK,DM,DR,DB,DL,DU)
          U = Unit number (0 to 7)
          M = Map number on that unit (1st map on
              unit is map number 0)

Enter > DR1:20

How many maps in a row ? > 10

Disk space allocation for SDP space #1 has been completed.

Allocate, deallocate, or show SDP spaces [A, D, or S] ? > S

Space  Starting  Ending  Disk  First   #
index  DSM blk   DSM blk unit  map     maps
-----  -------   ------- ----  -----   -----
    1    8000    11999   DR1   20      10
```

Generally, the process is similar for allocating spooling or journaling space. Space is allocated in disk units and maps.

Using ^%GLOMAN to Allocate Globals

You can place globals on a volume set by using the global management utility, ^%GLOMAN. Keep in mind the following points when you place globals:

- Allocate the starting point for the global in terms of the disk unit, map, and block.
- Start the global with one map for the global pointers.
- Allocate space in the next map for the first data block and the new data growth area.
- Allow plenty of space for the global to grow, if you know in advance that the global will be a large global. You do this by not allocating any other globals, UCIs, or system space directly after the global. Allow a space of several maps or more.

The following is a sample ^%GLOMAN dialog:

```
>DO ^%GLOMAN

Global Management Utility

Volume Set < JHM>  AAA

UCI > JSS

GLOBAL > ^BIG

^BIG is currently not defined.

Do you wish to create and place ^BIG [Y or N] ? Y

Address of the 1st GLOBAL POINTER BLOCK <DRO:0:0> DRO:50:0

Address of the 1st GLOBAL DATA BLOCK <DRO:0:0> DRO:51:0

^BIG placed

Global Characteristics for: ^[JSS,AAA]BIG

Collating: Numeric Journaling: Disabled Encoding: 8 - bit

New Data        1st Data       1st Pointer
Growth Area     Block          Block
-----------     --------       -----------
*DRO:0:0        DRO:51:0       DRO:50:0
* = Current UCI default for DATA GROWTH AREA will be used.

Global Access Privileges:

        System = RWD     World = None    Group = None    User = RWD

1. Show GLOBAL Characteristics
2. Change Access Privileges
3. Change Journaling Privileges
4. Change DATA GROWTH AREA
5. Change Collating Sequence
6. Change GLOBAL Encoding

Enter option > 4

Global DATA GROWTH AREA <DRO:0:0> DRO:51:0

- Characteristics updated

The space for the global has now been updated to:

New Data        1st Data        1st Pointer
Growth Area     Block           Block
-----------     --------        -----------
DRO:51:0        DRO:51:0        DRO:50:0
```

System Startup with Volume Sets

After you create additional volume sets, make sure that the system recognizes these volume sets when it is started up. You do this by building a start-up command file using the ^STUBLD utility.

Each configuration has one start-up command file that is run each time the configuration is started.

To run ^STUBLD, type the following from the system manager's account:

```
> DO ^STUBLD
```

The important question for volume sets is:

```
Mount additional disk volumes [Y or N] ? Y

Enter the disk mounting information. Type ? for help

DISK      DATA BASE       LABEL/VOLUME
UNIT      VOLUME SET      SET NAME
----      ----------      ------------
DK0       Y               JHM
DR0       Y               AAA
```

To designate a volume set, specify:

- The first disk of the volume set
- the 3-character volume set name.

For example, if volume set JHM contains disk volume one on disk drives DK0 and disk volume 2 on DK1, you only need to indicate the disk drive for volume one, DK0, and the volume set name, JHM.

Once built, the start-up command file performs the following each time you start up the configuration:

- Mounts the indicated volume sets
- Initializes the Volume Set Table

Chapter 4
DSM-11 V3.2 Distributed Data Processing

This chapter discusses Distributed Data Processing (DDP) on DSM-11.

Overview of Distributed Data Processing

Using Distributed Data Processing (DDP), you can construct point-to-point or Ethernet DSM-11 networks. Through communications devices and DDP software, you can access data bases across systems; and, using job control (the JOB command), control data processing on remote systems.

DDP is based on two types of connections between DSM-11 systems:

- Physical links - hardware connections that provide a physical data path between two or more DSM-11 systems.
- Logical circuits - software connections that provide a software circuit (based on the physical cables and connections) between two DSM-11 systems. Each pair of DSM-11 systems must have a separate logical circuit established between those two systems.

Both connections must be made and in operation before network communications can occur.

The relevant software structures and utilities used in DDP are:

- The local Volume Set Table for locally mounted volume sets.
- The local Circuit Table for remotely mounted volume sets.
- The System Generation utility (SYSGEN) for setting up DDP support and physical links.

- The start-up utility (^STU) for starting up DDP and enabling the physical links.
- The DDP utilities for establishing logical circuits.
- The DDP server (DDPSRV), a job that handles DDP requests from remote systems.

Extended Global Syntax

You access data across a DDP line through an extended global syntax. You can use the extended syntax in normal DSM-11 command lines, as in the SET command in the following example:

```
SET ^[JSS,XXX]GLO(6,8)=3.5
```

where:

JSS is the UCI in the remote data base volume set.

XXX is the name of remote data base volume set. You do not need to give the name of the remote system (for example, BBB) in the extended syntax.

GLO is the name of a global in the UCI JSS on the mounted volume set XXX.

See the *DSM-11 User's Guide* for more discussion of the extended syntax.

You can access globals across DDP lines by using the UCI Translation Table, set up by the utility, ^UCITRAN. See the *DSM-11 User's Guide* for more information about this utility.

Network Nodes

Each DSM-11 system connected to a network is called a node. The name of each node is the volume set name of the system volume set. Changing or renaming the system volume set changes the name of the node.

Thus, node BBB is called BBB because its system volume set is named BBB. If you change the name of the system volume set to CCC, the name of the node changes to CCC.

Volume Sets and DDP

Access across a DDP network is usually done through volume sets. This access has several characteristics:

- A data base volume set must be defined and mounted in the remote system. If the volume set on node BBB is not mounted, you receive a NOSYS error from BBB.
- The physical link between two nodes must be established for communications to work.
- Data base volume sets on nodes can be accessed by referring to the volume set's 3-character name, such as XXX. The volume set XXX can reside, for instance, on node BBB. The system volume set, BBB, can also be accessed.

DSM-11 follows a specific search algorithm when it attempts to resolve a global reference. DSM-11 takes the following steps:

1. Looks in the local Volume Set Table to see if the volume set is mounted locally.
2. Looks to see if the volume set is mounted on a remote node by looking in the Circuit Table (which lists volume sets on remote nodes).
3. Returns a <NOSYS> error if the volume set is not mounted locally or remotely.
4. Accesses the volume set over the DDP lines if it is mounted remotely.

DDP Line Access Types

You can do the following operations across DDP lines:

- Use the lock commands, including:
  - LOCK
  - UNLOCK
  - ZALLOCATE
  - ZDEALLOCATE

  For example:
  ```
  ZALLOCATE ^[JSS,XXX]GLO(4,10)
  ```
- Access globals with such commands and functions as:
  - SET
  - KILL
  - $DATA
  - $ORDER

  For example:
  ```
  SET ^[JSS,XXX]GLO(4,8)=3.5
  ```
- Run the JOB command to start a job on a remote node. The routine that you start with the JOB command must already exist on the remote node, for example:
  ```
  JOB ^ROUT[JSS,XXX]
  ```

DMC11/DMR11 Controller Point-to-Point Networks

DSM-11 systems can be connected in point-to-point networks. Each system must have a direct line connecting it to each node in the network.

Network Hardware

The hardware for point-to-point network consists of a DMC11 or DMR11 controller and cables. The DMC11 or DMR11 controller is a communication device. Both controllers are printed circuit boards that are mounted inside the PDP-11 system.

The cable used in a point-to-point arrangement connects a node to only one other node.

Physical Layout

You can use point-to-point arrangements to connect several nodes. The following figure shows a fully configured point-to-point arrangement of three nodes.

Figure 4–1 Three Node Point-to-Point Network

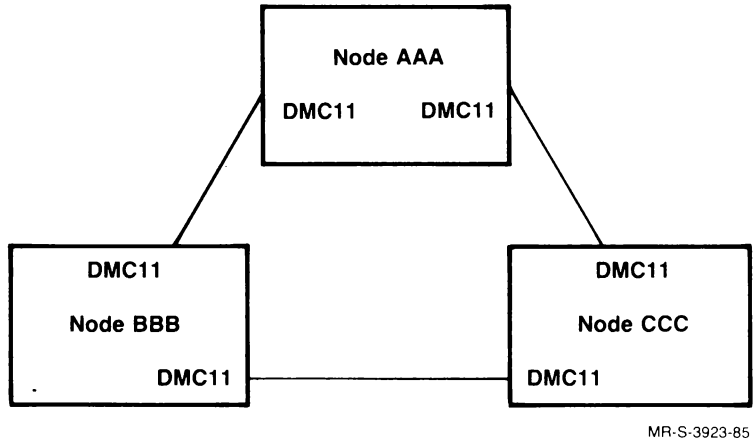

In this arrangement, each node (AAA, BBB, or CCC) must have two DMC11s or DMR11s. **A node cannot communicate through another node.** For instance, node BBB cannot reach node CCC by communicating through AAA. Node BBB can only reach CCC by using the direct link between BBB and CCC.

Ethernet Multipoint Networks

You can connect DSM-11 systems in multipoint networks. In the Ethernet network, each node connects to an Ethernet cable. Each node can access any other node connected to the cable.

Ethernet Hardware

The hardware for Ethernet consists of an Ethernet controller, a connecting cable to a transceiver, and the Ethernet cable which also connects to the transceiver.

The DEUNA Ethernet controller is used to connect PDP-11 UNIBUS systems to an Ethernet. The DEQNA controller is used to connect PDP-11 Q-bus systems to an Ethernet.

A cable runs from a PDP-11 to an Ethernet controller. The transceiver provides the connection between the cable from the PDP-11 and the Ethernet cable.

Ethernet Layout

The following figure shows a multipoint arrangement for three nodes.

Figure 4–2 Three-Node Ethernet Network

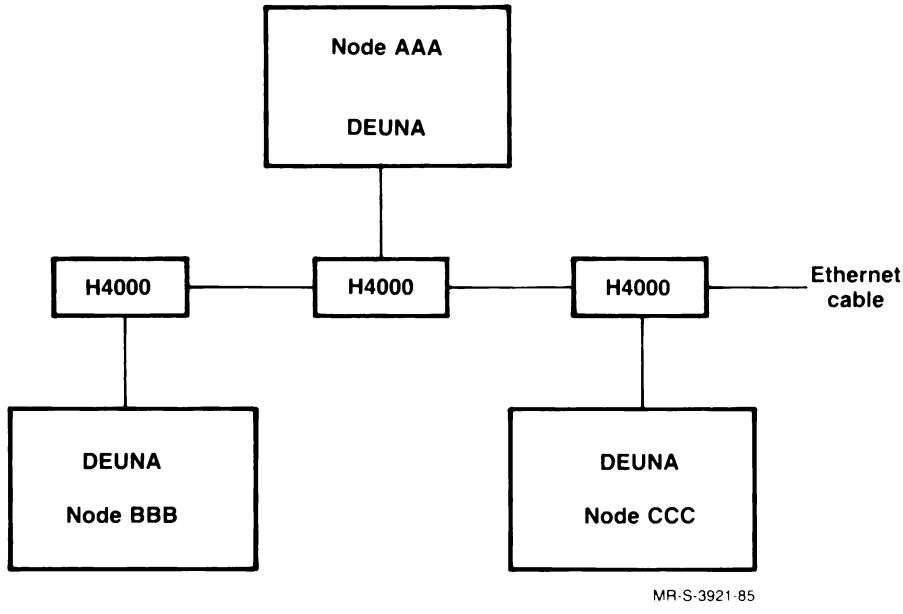

In this arrangement, each node must have one DEUNA or DEQNA, and one H4000 transceiver connecting it to the Ethernet cable. Any DSM-11 node in this arrangement can communicate with any other node connected to the same Ethernet cable.

Transceiver Address

The address of a node in an Ethernet network is based on the address of the H4000 transceiver. This defines a unique address for each system, because there is only one transceiver for each node. This address is in the form of six hexadecimal bytes such as the following:

AA-00-03-01-1B-41

Each Ethernet controller has its own address that is unique and never changes. The address for each transceiver must be correctly entered into the DSM-11 software for the network to work.

Mixed Network

You can also connect a system to more than one kind of network as shown in the following figure:

Figure 4–3 Mixed Network

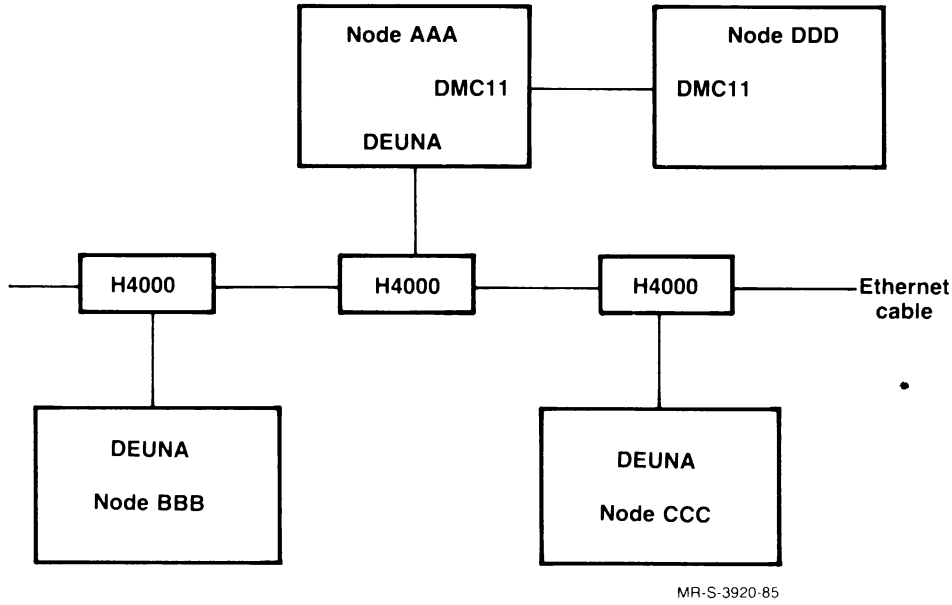

Physical Link

The physical link consists of the hardware components which connect the nodes. The software recognizes a link in terms of the communications devices, such as a DMC11 or DEUNA, attached to the node.

These devices must be physically connected to the node and connected by cables to other nodes or, in the case of an Ethernet, must be connected to an H4000 transceiver.

During SYSGEN you must identify the devices used for DDP communications to the node. At this time each DDP device is assigned a name by SYSGEN such as:

- XM0 - a DMC11 controller, the XM indicates a DMC controller.
- XE0 - a DEUNA controller, the XE indicates an Ethernet controller.
- XH0 - a DEQNA controller, the XH indicates an Ethernet controller.

When you start up DDP, the software assigns each device a link number. This link number serves to identify the physical link connected to one physical device.

For instance, Link 1 can refer to XM1, a DMC11 device. Link 2 can refer to XE0, a DEUNA device.

The link numbers do not change; they remain associated with the same physical device (as long as the same configuration is running).

Logical Circuit

The logical circuit is a software connection between two nodes that is layered on the physical link between those nodes. The logical circuit consists of data constructs which describe the status, address, and physical link where the node can be accessed.

A circuit must be established between the local node and each remote node. For a point-to-point network, there is one circuit for each physical link. For a multipoint network, there is a circuit for each remote node that exists on the Ethernet.

DSM-11 DDP supplies a network autoconfiguration facility that is initiated at DDP start up. This autoconfiguration facility automatically establishes the logical circuit connection between DSM-11 nodes.

The following figure shows the concept of a logical circuit in an Ethernet network. Note that the logical circuit is not a physical cable but a logical connection using the already existing physical connecting cables, transceivers, and Ethernet cable. The figure shows only the logical circuits from BBB to AAA and BBB to CCC.

Figure 4–4 Logical Circuits in an Ethernet System

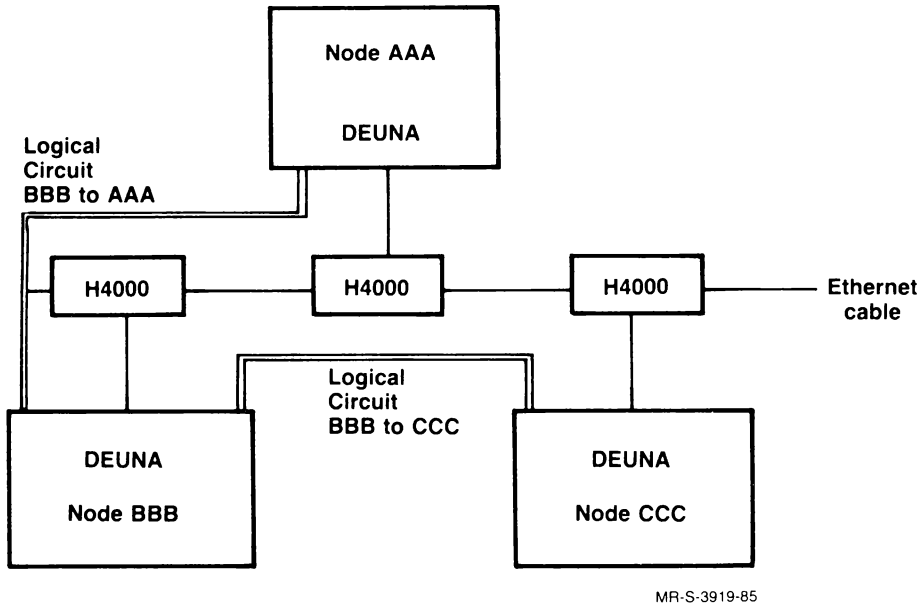

The Circuit Table

The Circuit Table provides logical circuit information about the local system to other nodes. You provide information for the circuit table by using the DDP utilities.

The Circuit Table contains information about:

- System volume set
- Mounted volume sets
- Ethernet address (if applicable)
- Link Number
- Status

Figure 4-5 shows the Circuit Table for node AAA shown in Figure 4-4 on mixed networks. Note that the mounted volume sets are indicated in this table.

Figure 4–5 Circuit Table

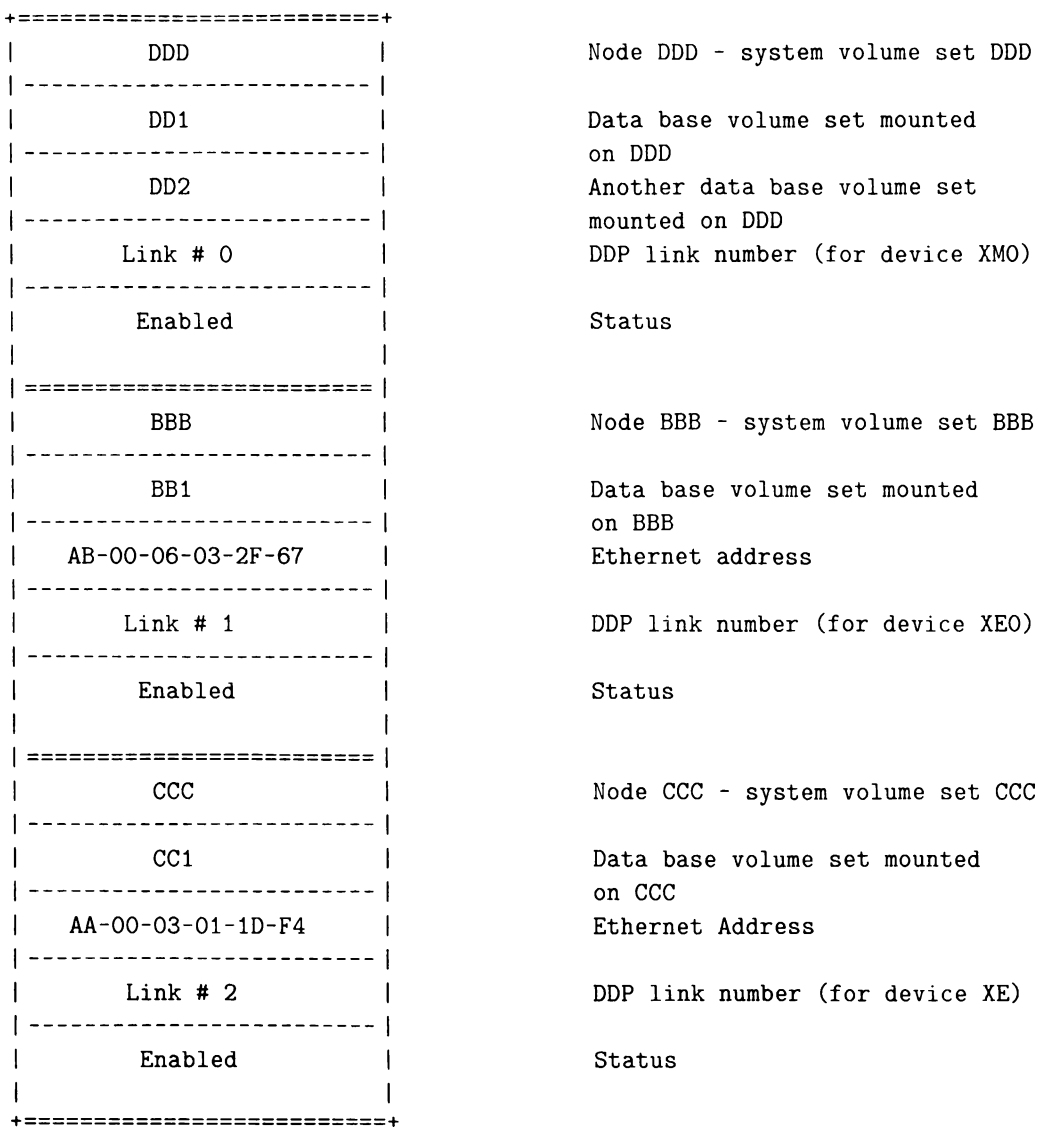

The DDP Server Job (^DDPSRV)

The DDP server job is a routine, ^DDPSRV, that handles DDP requests from remote nodes. The DDP server is a separate job that runs in its own 8K-byte partition. You can specify the number of DDP servers when you do SYSGEN. The DDP servers are activated when DDP is started. The default value is one DDP server for each circuit. You can specify a larger number of servers, but this requires additional partition space.

DDP servers serve a common queue of received requests. In situations where requests require physical disk accesses, additional DDP servers can hinder performance.

The DDP server and DSM-11 handle requests as follows:

1. DSM-11 places the request in a queue when it comes in over a DMC11 or Ethernet link.
2. DSM-11 assigns the next available DDP server to the request.
3. The DDP server processes the request and performs the requested action (such as a SET in the local data base).
4. The DDP server determines where the request originated from the request header information.
5. The DDP server looks up the remote system in the Circuit Table to determine specific information needed to return the reply.
6. The DDP server sends the reply.

The header information with a request contains a Message ID number (MID) that indicates in what order the request was sent in. Because requests can be resent and thus not received in the correct order, the DDP server uses the MID to insure that requests are executed in the intended order.

Setting up DDP on DSM-11 V3.2

Take the following steps to set up DDP:

1. Connect the systems physically using point-to-point or Ethernet hardware.
2. Run SYSGEN or autoconfiguration so that the system recognizes the communications devices.
3. Answer SYSGEN questions appropriately so DDP is initialized and supported and the DDP communications devices are set up for DDP. You must specify the total number of DDP nodes to allow sufficient space to describe all possible DDP circuits in the Circuit Table.
4. Start up the configuration, starting up DDP and initializing the physical links. The DDP autoconfiguration process determines what nodes exist on the network and enters the appropriate information in each Circuit Table on each node.

5. Physically load all the disk volumes needed for your volume sets.
6. Create (if necessary) all of the volume sets that you need.
7. Use the ^MOUNT utility to mount each volume set on its own system.
8. Use an extended global syntax or the UCI Translation Table (established with ^UCITRAN) to access globals in remote volume sets).
9. Add all DSM-11 V3.0 nodes to the Circuit Table with the ^DDPCON utility.

DSM-11 V3.2 SYSGEN for DDP

The manual SYSGEN procedure asks several questions about DDP. Keep the following points in mind about the DDP questions in SYSGEN:

- Questions about UCI Translation Table support are included because you can use this table to provide transparent access to globals on a remote system (using DDP).
- Questions about mountable volume set support are included because you probably will use volume sets in a DDP situation.
- Questions that are not shown in this SYSGEN are similar to those in the Version 3.0 SYSGEN. Refer to the *DSM-11 User's Guide* for more information on other parts of SYSGEN.
- You can also use the system autoconfiguration procedure in a DDP system. This procedure is mostly automatic and is the one used in the first-time installation procedure. Most questions are answered automatically in the system hardware autoconfiguration procedure. It does ask you the specific questions on DDP devices indicated in Part 4 of the manual SYSGEN.

The following shows the relevant DDP questions from the DSM-11 SYSGEN:

```
PART 3   SYSTEM DEVICES


3.5 How many DMC11's are there (max = 4) ?   0   2

Enter the VECTOR address, in OCTAL, for DMC11 controller #0 >300
Enter the CSR address, in OCTAL, for DMC11 controller #0 >160070

Enter the VECTOR address, in OCTAL, for DMC11 controller #1 >310
Enter the CSR address, in OCTAL, for DMC11 controller #1 >160100

3.6 How many DEUNA's are there (max = 2) ? 0  1

Enter the VECTOR address, in OCTAL, for DEUNA controller #0 >123
Enter the CSR address, in OCTAL, for DEUNA controller #0 > 160120



PART 4 CONFIGURE DMC/DMR-11'S AND DDP

DMC/DMR-11 controller #0

4.1 Is this device HALF-DUPLEX [Y or N] ? <N> [RET]
4.3 Will this device be used for DDP [Y or N] ? Y
4.4 Is this device connected to a Version 3.0 system [Y or N] <N> [RET]

DMC/DMR-11 controller #1

4.1 Is this device HALF-DUPLEX [Y or N] ? <N> [RET]
4.3 Will this device be used for DDP [Y or N] ? Y
4.4 Is this device connected to a Version 3.0 system [Y or N] <N> [RET]

How many DSM nodes are connected to DEUNA #0 > 3



PART 7 SOFTWARE OPTIONS



7.7 Include support for UCI TRANSLATION TABLES [Y OR N] ? Y

7.8 Include support for MOUNTABLE DATABASE VOLUME SETS [Y OR N] ? Y



PART 9: SYSTEM DATA STRUCTURES

Space allocated for DISK-MAP and BAD BLOCK TABLE: 15872 Bytes
Space allocated to UCI TRANSLATION TABLE: 1024 Bytes
Space allocated for DDP Structures:    13248 Bytes


9.2 Enter the number of ADDITIONAL mountable VOLUME SETS <3> [RET]
```

DDP Autoconfiguration

This section describes the DDP autoconfiguration process for automatically configuring DDP. This process does the following:

- Determines what nodes exist over a network.
- Enters the nodes in the local Circuit Table.
- Enters the mounted volume sets on each node in the local Circuit Table.
- Enables each circuit.

Do not confuse the DDP autoconfiguration process with the hardware autoconfiguration procedure (automatic SYSGEN). The hardware autoconfiguration procedure is described in the *DSM-11 User's Guide*.

The DDP autoconfiguration process occurs and the local Circuit Table is updated whenever:

- You start up DDP.
- A node mounts a new volume set.
- A node starts up or shuts down DDP.
- You run the update option of the CONTROL UTILITY (^DDPUTL).

You can still use the DDP utilities to manually edit the Circuit Table, and to enable or disable links or circuits.

DDP and Physical Link Start-Up

Starting up DDP involves the following steps:

1. Defining a start-up command file that starts up DDP.
2. Running the command file to start up the system as well as DDP. The command file does the following:
   - Starts the DDP servers.
   - Enables the physical links.
   - Runs the DDP autoconfiguration process.

If you previously defined and saved values for the Circuit Table, DSM uses those values when you start up DDP. You can, however, override those saved values with the DDP autoconfiguration process. When there are no permanently saved values for the Circuit Table, the DDP autoconfiguration process obtains information on the remote systems and volume sets and enters it into the local Circuit Table.

Defining a Start-Up Command File with DDP

You define the command file by running ^STUBLD:

```
>DO ^STUBLD
```

This utility asks you several questions. See the *DSM-11 User's Guide* for a complete description of this utility.

The relevant questions are:

```
Start Up Distributed Database Processing [Y or N] ? <N>Y
Enter the number of DDP servers to start <2> RET
```

In this example, the default answer 2 was accepted.

Sample DDP Start-Up

You start up DDP by running the DSM-11 Start-Up utility (^STU). You must have defined a start-up command file that requests DDP be started.

Run ^STU by typing:

```
>DO ^STU

Start Up the default system (11/44) [Y/N] ? Y

Re-configuring memory...
Memory re-configured

Mounting JHM as Volume Set number SO
Volume 1 on DL0 has 10000 blocks 2136 available.
Total in Volume Set 10000 blocks 2136 available.

Building terminal control blocks...

Caretaker is now running as job number 2
```

The following lines appear only if you have constructed a Circuit Table previously for this configuration and have saved the values for the table. In this case only one circuit (to BBB) had been previously defined and saved.

```
Building DDP Circuit Table

Circuit to node BBB created
```

The rest of this example occurs whether or not you have a Circuit Table previously defined.

```
DDP will start in 5 seconds.

DSM-11 Version 3.2 11/44B is now up and running!
Exit

DDP start-up . . . 2 DDP servers started ... DDP startup complete
15:04 DDP server 4 - Circuit to node CCC through link #2 created

DSM-11 Version 3.2 Device #1 UCI:
```

At the end of the process, the system gives you the log in prompt and you are ready to log in. The physical links for DDP have been initialized and enabled. The DDP autoconfiguration process sets up circuits to any existing (and running) node, in this case a circuit to system CCC is established.

Using the DDP Utilities to Start Up or Shut Down DDP

You can start up or shut down DDP by selecting options from the menu of DDP utilities. The following example shows how to shut down DDP:

```
>DO ^DDP

DDP Utilities:

1. CIRCUIT MANAGEMENT            (^DDPCIR)
2. CONTROL UTILITY               (^DDPUTL)
3. LINK MANAGEMENT               (^DDPLNK)
4. SHUTDOWN DDP                  (STOP^DDPLNK)
5. STARTUP DDP                   (START^DDPLNK)

Enter Option > 4

DDP Shutdown

Ready to shutdown [Y or N] ? <N> Y

DDP shutting down

DDP servers stopping

17:58    DDP Server 1 - DDP communications shutdown - ^DDP server
                halted
17:58    DDP Server 4 - DDP communications shutdown - ^DDP server
                halted

DDP shutdown complete
```

If you select the STARTUP DDP option of the DDP utilities menu, you see this dialog:

```
Start up DDP

Enter the number of DDP servers to start <2> RET

DDP startup ... DDP servers starting

2 servers started
```

Using the ^DDPUTL Utility

You can use the ^DDPUTL utility as another way of controlling circuits. This utility enables you to control circuits without going through the menu structure.

You see the following dialog when you select the ^DDPUTL utility:

```
>DO ^DDPUTL

DDP Control Utility

[S]tatus [E]nable [D]isable [V]erify [U]pdate

Option >
```

Select the U option to force a DDP autoconfiguration. Whenever you choose U, the local system initiates a configuration update. The local system sends information (including what volume sets it has mounted) to remote systems and requests that the remote systems return similar information.

Using the DDP Physical Link Utilities

The physical links for DDP are created during the SYSGEN process. The links are initialized and enabled in the DSM-11 start-up procedure (if the DSM-11 start-up command file requests DDP to be started).

You can disable, enable, change the service state, or check the status of the physical links by using the DDP utilities.

You can access and run these utilities as follows:

```
>DO ^DDP

DDP Utilities:

1. CIRCUIT MANAGEMENT              (^DDPCIR)
2. CONTROL UTILITY                 (^DDPUTL)
3. LINK MANAGEMENT                 (^DDPLNK)
4. SHUTDOWN DDP                    (STOP^DDPLNK)
5. STARTUP DDP                     (START^DDPLNK)

Enter Option> 3

DDP Link Management:

1. DISABLE LINK                    (DISAB^DDPLNK)
2. ENABLE LINK                     (ENAB^DDPLNK)
3. LINK STATUS                     (STA^DDPLNK)
4. SET LINK SERVICE                (SERV^DDPLNK)

Enter Option>
```

Link Status

You run the link status utility by selecting the LINK STATUS option of the DDP LINK MANAGEMENT utility. An example follows:

```
DDP Link Status

                                                    Ethernet
Link # Device    Type    State           Errors          Address
0        XM0     DMC11   Enabled                 2
                         In Service
1        XM1     DMC11   Enabled                 0
                         In Service
2        XE0     DEUNA   Enabled                 7     AA-00-03-01-1B-41
                         Out of Service
```

Enabling and Disabling Links

All links are enabled whenever DDP starts up. If you want to disable a link permanently, you must put it in an "out of service" state, by using the SET LINK SERVICE option of the LINK MANAGEMENT menu.

If you select the DISABLE LINK option of the LINK MANAGEMENT menu, you see the following dialog:

```
Disable DDP link

Link # > 2

Link  2 - disabled
```

If you select the ENABLE LINK option you see the following dialog:

```
Enable DDP link

Link # > 2

Link #2 - enabled
```

Link Service

You can use the SET LINK SERVICE option to put a link in service or out of service. The service state of the link is permanent and is not changed when DDP is started up. This is different from the ENABLE LINK or DISABLE LINK options. All links are enabled whenever DDP is started up.

The SET LINK SERVICE option can be particularly useful if you have two DMC11 or DMR11 links between two systems. In this case one of these links is treated as a back up or reserve link and must be placed out of service. In fact, you must place of the links out of service in such a situation or the DDP autoconfiguration process will not work properly.

You can change the service state by selecting the SET LINK SERVICE option of the LINK MANAGEMENT menu. When you select the SET LINK SERVICE option, you see the following dialog:

```
Set Link Service

Link # > 2

DDP Link #2 is currently In Service

Set In Service or Out of Service [I or O] > O
```

Using DDP Utilities to Establish or Modify a Logical Circuit

You use the CIRCUIT MANAGEMENT option of the DDP utilities menu to establish or modify a logical circuit. The DDP autoconfiguration process automatically establishes circuits, so in most cases, you do not need to do this manually. (You must, however, manually create and preserve circuit entries to DSM-11 V3.0 systems.)

The following is an example of this option:

```
>DO ^DDP

1. CIRCUIT MANAGEMENT              (^DDPCIR)
2. CONTROL UTILITY                 (^DDPUTL)
3. LINK MANAGEMENT                 (^DDPLNK)
4. SHUTDOWN DDP                    (STOP^DDPLNK)
5. STARTUP DDP                     (START^DDPLNK)

Enter Option> 0

DDP Circuit Management:

1. CIRCUIT STATUS                  (STA^DDPCIR)
2. DISABLE CIRCUIT                 (DISAB^DDPCIR)
3. ENABLE CIRCUIT                  (ENAB^DDPCIR)
4. MODIFY CIRCUIT INFORMATION      (^DDPCON)
5. RESET CIRCUIT COUNTERS          (RESET^DDPCIR)
6. SHOW CIRCUIT COUNTERS           (COUNT^DDPCIR)

Enter option > 4

Modify Circuit Information

Node name >
```

Enter the 3-character uppercase node name of the system you are referencing. The node name is the name of the booted SYSTEM VOLUME set.

```
Type * if you wish to list or reference all KNOWN nodes

Node name > BBB

                                              Mounted      Ethernet
Circuit Link #   Link State Circuit State Volume Sets  Address
BBB      0       Enabled            Enabled       BBB     AA-00-03-01-1D-F5
                           Reachable

Edit node BBB circuit information [Y or N] Y

DDP link   > 2

Mounted Volume Sets > BBB

Ethernet Physical Node Address > <AA-00-03-01-1D-F5> AA-00-03-01-1D-F4

Node BBB Circuit information updated

Update the permanent database [Y or N] ? <N>Y

Node BBB permanent database updated.
```

To see the local Circuit Table, select the CIRCUIT STATUS option of the CIRCUIT MANAGEMENT menu:

```
Circuit Status

                                              Mounted      Ethernet
Circuit Link #   Link State Circuit State Volume Sets  Address
BBB      0       Enabled            Enabled       BBB     AA-00-03-01-1D-F4
                           Reachable
```

A circuit is in an unreachable state before any attempts are made to reach the node or when a request to the node fails. Once a request succeeds in reaching a node, the circuit state is reachable.

To enable or disable a link, use the LINK MANAGEMENT option of the DDP utilities menu. To enable or disable a circuit, use the ENABLE CIRCUIT or DISABLE CIRCUITS options of the CIRCUIT MANAGEMENT menu.

The COUNTERS options of the CIRCUIT MANAGEMENT menu can be helpful in debugging procedures.

DDP Error Procedures

DSM-11 DDP monitors and reports DDP software and hardware errors. DDP retries requests that have not been successfully sent. Repeated errors can disable the link. Errors are also logged in the hardware error log which you can accessed by using the utility ^KTR.

Link Errors

DDP link errors may be detected and reported in the following situations:

- During the hardware initialization procedure initiated when a DDP link is enabled
- Whenever an attempt to transmit or receive a message results in a hardware error condition.

In either case, DDP reports these errors to the system CARETAKER. If the CARETAKER is running, the device error status is logged and a hardware report can be generated using the ^KTR utility. (See the *DSM-11 User's Guide* for more information about the CARETAKER and the ^KTR utility.)

When a hardware error is detected, the DDP Link is automatically disabled, and a DDP server is signaled to attempt to reestablish the DDP link. These attempts are reported to the system logging console.

If the hardware error recurs, and four attempts by the DDP servers fail to reinitialize the device, the DDP link is disabled and no further attempts are made to restart the link. The only method of restarting the link is to use the ENABLE LINK option in the ^DDPLNK utility.

Any DDP requests that were outstanding during this error process are requeued to the link. If the link is successfully enabled, the requests are sent. If the link is not reestablished after four attempts, all jobs with outstanding DDP requests queued to that DDP link receive a DSM-11 DSTDB error.

You can display the current number of hardware errors and the DDP link status by using the ^DDPLNK utility or the LINK MANAGEMENT option on the ^DDP utility menu. This utility reports:

- Received message errors
- Transmitted message errors
- Device errors

Received Message Errors

Received message errors occur during the reception of a DDP request. The received message error counter is relevant primarily for Ethernet-based DDP Links. The error counter tallies the number of messages with corrupted data or "runt" data packets, usually the result of a packet collision on the Ethernet.

While it is normal to experience Received Message Errors during periods of very high traffic (more than 250 messages/second), high rates of Received Message Errors at low traffic levels can indicate a cabling problem or poor Ethernet tap.

Transmitted Message Errors

Transmitted message errors occur during DDP request transmission. The transmitted message error counter is relevant primarily for Ethernet-based DDP Links. The error counter tallies the number of failed message transmissions, usually because of excessive packet collision on the Ethernet. It is normal to experience some transmitted message errors during periods of high message traffic.

Device Errors

Device errors are hardware failures, spurious interrupt events, and unexpected conditions. High rates of device errors can indicate device failures.

Circuit Errors

DSM detects circuit errors when a local job's DDP request fails to get a response from the remote DDP server. The DSM system retries a DDP request. If DSM is not able to get a response after four tries at five-second intervals, it declares the node unreachable. DSM then takes the following action:

- Reports a < DSTDB > error to the requesting job
- Sends an error message to the system logging device

DSM does not change the circuit status to reachable until a successful DDP request is completed (a response from the remote node received).

The following table shows possible causes of circuit errors and their remedies:

Problem	Remedy
Local DDP link state is disabled.	Use the ^DDPLNK utility's ENABLE DDP LINK option on the local node.
Remote DDP link state is disabled.	Use the ^DDPLNK utility's ENABLE DDP LINK option on the remote node.
Local DDP circuit state is disabled.	Use the ^DDPCIR utility's ENABLE DDP CIRCUIT on the local node
Remote DDP circuit state is disabled.	Use the ^DDPCIR utility's ENABLE DDP CIRCUIT option on the local node.
Local Circuit Table has incorrect Ethernet address for remote node.	Use the the ^DDPUTL utility's U option to force a DDP network configuration update.
Local Circuit Table has wrong DDP link number for remote node.	Use the the ^DDPUTL utility's U option to force a DDP network configuration update.
Remote node Circuit Table has incorrect Ethernet address for the local node.	Use the the ^DDPUTL utility's U option to force a DDP network configuration update.
Remote node Circuit Table has wrong DDP link number for local node.	Use the the ^DDPUTL utility's U option to force a DDP network configuration update.
Remote node Circuit Table contains no entry for local node.	Use the the ^DDPUTL utility's U option to force a DDP network configuration update.

You can examine error activity by using SHOW CIRCUIT COUNTERS option for the ^DDPCIR utility. The following is an example:

```
Circuit Activity Counters

Node name > BBB

          Requests    Requests Requests     Requests Received
Node      Received     Sent    Retried      Out-of-sync
BBB       15          10       2              2
```

Note that DSM V3.2 does not attempt retries to DSM-11 V3.0 DDP nodes. In all other situations, you should expect retired requests on Ethernet links. Retried requests reflect lost requests or bit errors within a message. Normally, the number of errors should be a relatively small percentage of the total number of requests transmitted. A high retry rate can indicate a faulty or damaged Ethernet connection.

Chapter 5
Transfers between DSM-11 and VAX DSM

This chapter describes how to use magnetic tape to transfer routines and globals between DSM-11 and VAX DSM. The VAX DSM system must be V2.2 or later.

Transferring Routines from DSM-11 to VAX DSM

Take the following steps to transfer routines between DSM-11 and VAX DSM. You can only transfer one file (one routine save) in a single operation.

1. Create a tape on DSM-11 using ^%RS. Use magnetic tape mode AVLn where n is 4 for a 1600-BPI magnetic tape or 3 for an 800-BPI magnetic tape.
2. Load the magnetic tape onto a VAX tape drive.
3. Log in to the VAX and use the DCL MOUNT command:
   ```
   $ MOUNT/OVERRIDE=OWNER MTAu:/DENSITY=n MUMPS1
   ```
   where:
   - u is the unit number of the tape drive.
   - n is 800 for 800-BPI tapes and 1600 for 1600-BPI tapes.
4. Run VAX DSM:
   ```
   $ DSM
   ```
5. Run ^%RR to restore the routines. Use MTAn:MUMPS.SRC as the input file name when the utility asks for the input file.
   ```
   >D ^%RR
   Routine Restore
   Input File ? MTAO:MUMPS.SRC
   ```

Transferring Globals from DSM-11 to VAX DSM

Take the following steps to transfer globals from a DSM-11 system to a VAX DSM system. In this procedure, you can transfer only one file (one global save) in a single operation.

1. Create a tape on DSM-11 using ^%GTO. Use magnetic tape mode AVLn where n is 4 for a 1600-BPI magnetic tape and 3 for an 800-BPI magnetic tape.
2. Load the tape on the VAX.
3. Log in to the VAX and use the DCL MOUNT command:
   $ MOUNT/OVERRIDE=OWNER MTAu:/DENSITY=n MUMPS1

   where:

   u is the unit number of the tape drive.

   n is 800 for 800-BPI tapes and 1600 for 1600-BPI tapes.
4. Run VAX DSM:
   $ DSM
5. Use the %GTI11 utility supplied with VAX DSM to read the tape.

Transferring Routines From VAX DSM to DSM-11

Take the following steps to transfer routines between VAX DSM and DSM-11. You can only transfer one file (one routine save) in a single operation.

1. Load the tape on the VAX.
2. Log in to the VAX and initialize the tape:
   $ INIT MTAu:/DENSITY=n VAXDSM

   where:

   u is the unit number of the tape drive.

   n is 800 for 800-BPI tapes and 1600 for 1600-BPI tapes.
3. Mount the tape using the DCL MOUNT command:
   $ MOUNT/OVERRIDE=OWNER MTAu:/DENSITY=n VAXDSM
4. Run VAX DSM:
   $ DSM
5. Run the VAX DSM utility, ^%RS, selecting MTAu:ROUTINES.DAT as the output file.
   >DO ^%RS

   ```
   Routine Save
   Output File ? MTAO:ROUTINES.DAT
   ```

6. Dismount the tape from the VAX, using the DCL DISMOUNT command.
   ```
   $ DISMOUNT MTAu:
   ```
7. Load the tape on the DSM-11 system.
8. Run the ^%RR utility. Specify magnetic tape mode AVLn, where n is 4 for a 1600-BPI tape and 3 for an 800-BPI tape. Specify a block size of 2048.

Transferring Globals from VAX DSM to DSM-11

Take the following steps to transfer globals between VAX DSM and DSM-11. You can only transfer one file (one global save) in a single operation.

1. Load the tape on the VAX.
2. Log in to the VAX and initialize the tape:
   ```
   $ INIT MTAu:/DENSITY=n VAXDSM
   ```
   where:

   u is the unit number of the tape drive.

   n is 800 for 800-BPI tapes and 1600 for 1600-BPI tapes.
3. Mount the tape using the DCL MOUNT command:
   ```
   $ MOUNT/OVERRIDE=OWNER MTAu:/DENSITY=n VAXDSM
   ```
4. Run VAX DSM:
   ```
   $ DSM
   ```
5. Run the VAX DSM utility, ^%GS, selecting MTAu:GLOBALS.DAT as the output file.
   ```
   >DO ^%GS
   Global Save
   Output File ? MTAO:GLOBALS.DAT
   ```
6. Dismount the tape from the VAX, using the DCL DISMOUNT command.
   ```
   $ DISMOUNT MTAu:
   ```
7. Load the tape on the DSM-11 system.
8. Run the ^%GR utility. Specify magnetic tape mode AVLn, where n is 4 for a 1600-BPI tape and 3 for an 800-BPI tape. Specify a block size of 2048.

Chapter 6
Patching Instructions

This chapter provides patching instructions for DSM-11.

Overview of Patching

Patches are published in the *DSM Software Dispatch*. Read this publication for up-to-date patches that apply to your system.

There are two types of DSM patches:

1. Patches that do not require patch storage space because they use the same number of, or fewer, bytes of storage as the original program.
2. Patches that do require patch storage space because they are larger than the corresponding portion of the original program.

The *DSM Software Dispatch* describes patches and indicates whether a patch requires additional space.

You have the responsibility of maintaining information about the number of patches installed and their locations in memory. When you install patches, remember that all patch numbers are listed in *octal* notation.

Applying Patches

You can apply a patch either:

1. To the system image in memory
2. To the system image on the system disk

Using the Autopatch utility (^PATCH), you can specify which method to use.

Any patch you install on the system image is permanent. **Thus, you should test a new patch in memory before applying it to the system image.**

Chapter 7
DIGITAL Services

This chapter describes services that DIGITAL provides to help you use DIGITAL products effectively.

Software Problem Reporting

You can report software problems, inadequacies, and suggestions for improvements with *Software Performance Reports* (SPRs). Whenever DIGITAL receives an SPR, it sends an acknowledgment to the reporter. DIGITAL also publishes any SPRs of wide-spread significance in the *DSM Software Dispatch.*

> **Note:** Documentation errors and inadequacies should be reported on the READER'S COMMENTS page at the end of each manual.

Before sending an SPR to DIGITAL, make certain that the problem is reproducible. Then, check the *DSM Software Dispatch* to determine that a correction has not already been published.

If the problem is new, fill out a Software Performance Report and send it to your local SPR center. You can find the address of your local SPR center on the back of the SPR form.

When you fill out an SPR, provide as much information as possible to help describe and isolate the problem. Include configuration information, software version numbers, and any examples, tapes, or listings that might be needed to investigate a problem or suggested change. In general, the response time is shortened when you enclose complete and accurate information.

SPRs are also useful for reporting suggestions and comments on DSM-11. SPRs are monitored by DIGITAL management and are considered by the development groups when DSM-11 changes are made.

Blank SPR forms are included in software kits, and additional forms are available from the Software Distribution Center. Replacement forms are included with each answer.

To make the best use of the SPR form, follow these steps:

- Complete each SPR form with customer name and address, and statement of the problem. Type or print these items clearly.
- Complete a separate SPR form for each problem or suggestion.
- Report problems with patches as soon as possible. Do not use the SPR for problems concerning software policy, software distribution, or hardware. These problems should be reported to your local office.
- Label the tapes you submit with SPRs. Tape labels should indicate tape density and other recording characteristics.
- Mark the SPR to indicate how you want it handled.
  - SPRs that you mark DO NOT PUBLISH are not published in the *DSM Software Dispatch*. Use this designation for security SPRs.
  - DSM-11 SPRs that you mark PROBLEM/ERROR are answered.
  - SPRs that you mark FYI or SUGGESTION are forwarded to the pertinent software group for information purposes. They are answered at the discretion of the software group.

The DSM-11 Software Dispatch

Announcements of new and revised software, as well as programming notes, software problems with their proposed solutions, and documentation corrections, are published in the *DSM Software Dispatch*. Contact your local DIGITAL sales office to receive this publication for one year.

You should test and verify software changes off line before you install them in your operating system.

Software and Document Distribution

The *PDP-11 Software Price List* contains a list of software and documents currently available. You can order items directly from the Software Distribution Center by using the software order form enclosed in the price list. As noted above, new software is announced in the *DSM Software Dispatch*.

Digital Equipment Computer Users' Society

DIGITAL created and supports the Digital Equipment Computer Users' Society (DECUS), a voluntary, nonprofit, users' group. The purpose of DECUS is to advance the effective use of Digital Equipment Corporation's computers and peripheral equipment. Its objectives are to:

- Advance the art of computation through mutual education and the interchange of ideas and information.
- Establish standards and provide channels to facilitate the free exchange of computer programs among members.
- Provide feedback to the manufacturer on equipment and programming needs.

DECUS sponsors technical symposia twice a year (Spring and Fall) in the United States, and once a year in Europe, Canada, and Australia. It maintains a Program Library and publishes a library catalog, a booklet on the proceedings of symposia, and a periodic newsletter (DECUSCOPE).

DIGITAL formed a DECUS-Europe organization in 1970 to assist in the servicing of European members.

If you are interested in joining DECUS, obtain and complete a registration form. Forms can be obtained from the nearest DIGITAL sales office or from the appropriate administrative office.

The main administrative office address is:

```
DECUS
MRO2-1/C11
Digital Equipment Corporation
One Iron Way
Marlboro, Massachusetts 01752
```

All correspondence should be directed to the attention of the DECUS Executive Director.

The European administrative office address is:

```
DECUS EUROPE
Digital Equipment Corporation International
12, Avenue Des Morgines
Case Postale 510
CH - 1213 Petit-Lancy 1 (Geneva)
Switzerland
```

MUMPS Users' Group

The MUMPS Users' Group (MUG) is an organization of current and potential MUMPS users. Its objective is to acquaint new and potential users with the numerous MUMPS applications. It also provides a forum for current users to exchange information and remain abreast of new ideas in MUMPS development.

Because DSM-11 is DIGITAL's version of Standard MUMPS, you may be interested in joining the MUMPS Users' Group, or in obtaining additional information about the organization. If so, please contact:

```
MUMPS Users' Group
P.O. Box 37247
Washington, D.C. 20013
```

READER'S COMMENTS

NOTE: This form is for document comments only. DIGITAL will use comments submitted on this form at the company's discretion. If you require a written reply and are eligible to receive one under Software Performance Report (SPR) service, submit your comments on an SPR form.

Did you find this manual understandable, usable, and well organized? Please make suggestions for improvement.

_____

_____

_____

_____

_____

_____

_____

_____

Did you find errors in this manual? If so, specify the error and the page number.

_____

_____

_____

_____

_____

_____

_____

_____

Please indicate the type of user/reader that you most nearly represent:

- [ ] Assembly language programmer
- [ ] Higher-level language programmer
- [ ] Occasional programmer (experienced)
- [ ] User with little programming experience
- [ ] Student programmer
- [ ] Other (please specify) _____

Name _____ Date _____

Organization _____

Street _____

City _____ State _____ Zip Code _____
or Country

— — Do Not Tear - Fold Here and Tape — — — — — — — — — — — — — — —

digital

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO.33 MAYNARD MASS.

POSTAGE WILL BE PAID BY ADDRESSEE

SOFTWARE PUBLICATIONS
200 FOREST STREET MRO1–2/L12
MARLBOROUGH, MA 01752

— — Do Not Tear - Fold Here — — — — — — — — — — — — — — —