IBM

Storage Subsystem Library

GC26-4493-0

Using the IBM 3380 Direct Access Storage in a VM Environment

306
NEW

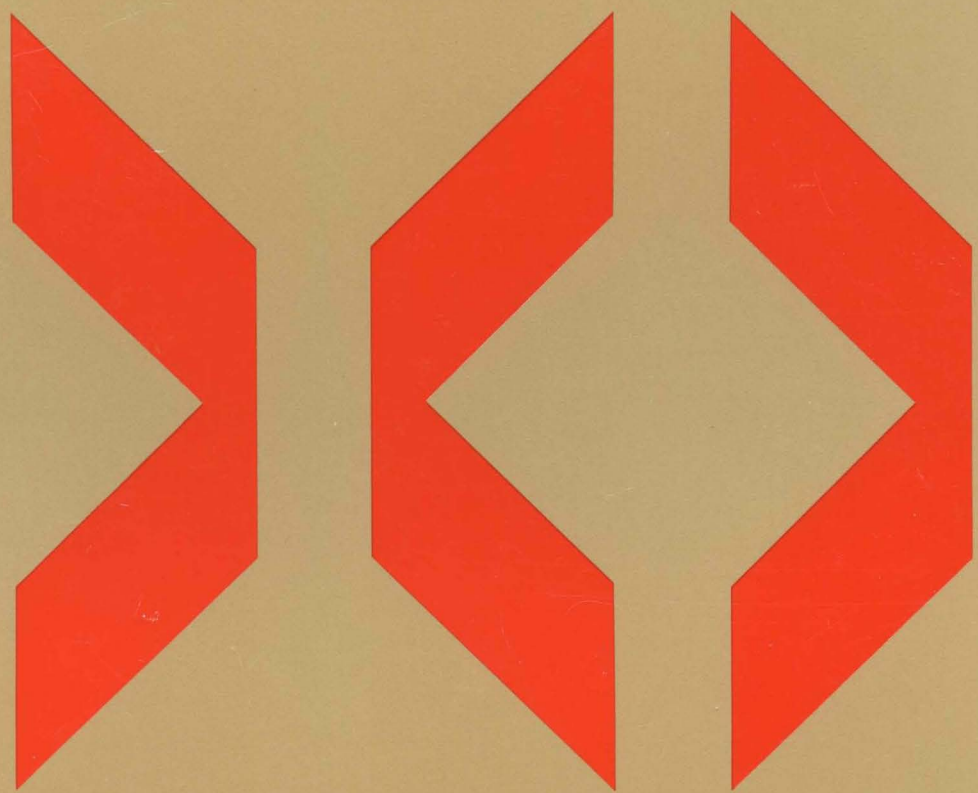

Storage Subsystem Library

GC26-4493-0

Using the IBM 3380 Direct Access Storage in a VM Environment

Federal Communications Commission (FCC) Statement

Warning: This equipment generates, uses, and can radiate radio frequency energy and if not installed and used in accordance with the instructions manual, may cause interference to radio communications. It has been tested and found to comply with the limits for a Class A computing device pursuant to Subpart J of Part 15 of FCC Rules, which are designed to provide reasonable protection against such interference when operated in a commercial environment. Operation of this equipment in a residential area is likely to cause interference in which case the user at his own expense will be required to take whatever measures may be required to correct the interference.

In addition to the FCC statement above, the reader of this manual should be aware that the referenced statement applies only to devices manufactured after January, 1981, and used in the United States.

First Edition (September 1987)

This edition applies to all models of the IBM 3380 Direct Access Storage. It incorporates the information in *IBM 3380 Direct Access Storage: Planning and Use*, GC26-4208, *IBM 3380 Direct Access Storage: Migration*, GC26-4197, and *VM/SP IBM 3380 Direct Access Storage Device Models AE4/BE4 User's Guide*, SC24-5281, which are now obsolete.

Changes are made periodically to this publication; before using this publication in connection with the operation of IBM systems, consult the latest *IBM System/370, 30xx, and 4300 Processors Bibliography*, GC20-0001, for the editions that are applicable and current.

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM licensed program in this publication is not intended to state or imply that only IBM's program may be used. Any functionally equivalent program may be used instead.

Requests for IBM publications should be made to your IBM representative or to the IBM branch office serving your locality. If you request publications from the address given below, your order will be delayed because publications are not stocked there.

A form for readers' comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM Corporation, P.O. Box 50020, Programming Publishing, San Jose, California, U.S.A. 95150. IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Preface

This manual is part of the Storage Subsystem Library (SSL)—a set of manuals that provides information about the hardware components of IBM disk storage subsystems. Although the SSL includes both direct access storage devices (DASD) and storage control publications, this manual is part of the SSL subset that is concerned primarily with 3380 DASD.

To use this manual productively, you should plan to read the SSL manuals, *IBM 3380 Direct Access Storage Introduction*, for a description of all 3380 models except the 3380 Model CJ2, *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference*, for a description of the 3380 Model CJ2, and *Maintaining IBM Storage Subsystem Media*, for information on 3380 maintenance. The manuals describing the storage controls to which you will attach the 3380 should be available for reference. You should also have a complete set of publications for your VM operating system available for reference.

About This Manual

This manual is written for the VM system programmer or hardware specialist who is responsible for installing and using the 3380 under any of the following VM operating systems:

- Virtual Machine/System Product (VM/SP), 5664-167
- Virtual Machine/System Product High Performance Option (VM/SP HPO), 5664-173
- Virtual Machine/Extended Architecture Systems Facility (VM/XA SF), 5664-169

In addition to device-specific information for the various models of the 3380, this manual also illustrates techniques for more efficient storage management under VM. It offers guidance on managing system performance, availability, and space through effective use of the DASD storage subsystem.

This manual contains:

Chapter 1, "Introducing the IBM 3380 Direct Access Storage" on page 1, provides an overview of 3380 functions and describes how 3380 units can be attached to storage controls and processor channels.

Chapter 2, "Planning for Installation" on page 5, describes how to plan for physical installation, determine software requirements, and develop a device installation project plan.

Chapter 3, "Understanding Your Data and Hardware Configuration" on page 17, explains how to identify data in your environment, where to look for device-dependent data and programs, and how to use software tools to measure the effectiveness of your hardware configuration.

Chapter 4, "Planning the Hardware Configuration" on page 35, describes considerations for hardware capability, capacity, performance, availability, shared DASD, and guest systems when planning the hardware configuration. It also discusses how to build and document a hardware configuration.

Chapter 5, "Planning the Data Configuration" on page 53, describes how to map data to volumes, and how to schedule data movement during device migration.

Chapter 6, "Installing the 3380 under VM" on page 61, explains how to define devices to VM, how to install 3380 units, and how to prepare volumes for use.

Chapter 7, "Operating the 3380 under VM" on page 79, describes the system commands you can use to determine or change the status of 3380 volumes.

Chapter 8, "Moving Data onto 3380 Volumes" on page 83, discusses the various tools available for moving data, and provides examples of how to move volumes, minidisks, files, and system data areas.

Chapter 9, "Monitoring and Maintaining the 3380" on page 93, outlines the techniques you can use to make effective use of the 3380, including performance tuning, space management, and media maintenance.

Appendix A, "Sample DASD Installation and Migration Project Plan" on page 105, provides a sample device installation and migration project plan that you can use as a model for your data processing center.

Appendix B, "Migration Aids" on page 115, contains a sample exec to automate the migration of minidisks and a sample exec to merge minidisks to larger capacity 3380s.

"Glossary" on page 117 lists the terms and abbreviations used in the Storage Subsystem Library manuals.

"Bibliography" on page 127 lists the related publications that you can reference for further information on related topics and hardware.

Notes for Guest Systems

Boxes like this, labeled **Notes for Guest Systems**, are used throughout this manual to isolate special considerations for guest operating systems running under VM. If you are planning to use the 3380 in a guest operating environment, look for these notes in each chapter. These notes have been indexed for easy reference.

Terminology

A comprehensive glossary is provided at the back of this manual. This glossary contains terms used not only in this manual but also terms, abbreviations, and acronyms from other DASD and storage control manuals in the Storage Subsystem Library.

Before reading further, be sure you understand how the following terms are used within this manual:

3380 refers to all models of the IBM 3380 Direct Access Storage, unless otherwise indicated.

Controller refers to the part of the 3380 A-unit or C-unit that controls the transfer of data between the devices and the storage control. The controller is sometimes referred to as the head-of-string.

Device refers to a uniquely addressable part of the 3380 unit that includes a set of access arms, their associated disks, and the electronic circuitry needed to locate, read, and write data.

Storage Control refers to the hardware that handles interactions between the processor channel and the controller. In the past, this has been called a control unit.

Volume refers to the set of disk surfaces associated with a device.

VM refers to all VM operating systems, including VM/SP, VM/SP HPO, and VM/XA SF, unless otherwise indicated.

The Storage Subsystem Library

The Storage Subsystem Library describes characteristics, capabilities, and features of the hardware and provides instructions for installing, using, and maintaining storage subsystem components effectively in the various operating environments. The library is designed to provide both hardware- and software-related information for both DASD and storage controls.

The DASD subset of the Storage Subsystem Library contains the following manuals:

1. *IBM 3380 Direct Access Storage Introduction*, GC26-4184

   Provides a complete description of the various models of the 3380, including characteristics, features, and capabilities. In addition, the configuration and attachment options are described along with other information that helps in designing a storage subsystem to meet your needs. This manual does *not* cover the 3380 Model CJ2.

2. *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference*, GC26-4193

   Provides a complete description of the 3380 Direct Channel Attach Model CJ2 characteristics, features, capabilities, and string configuration options.

3. *Using the IBM 3380 Direct Access Storage in an MVS Environment*, GC26-4192

   Provides specific guidance for using the 3380 in an MVS/XA or MVS/370 operating environment. This manual provides detailed instruction for planning the addition of new 3380 devices from a logical and physical point of view, installing devices, moving data to new devices, and performing ongoing activities to maintain a reliable storage subsystem.

4. *Using the IBM 3380 Direct Access Storage in a VM Environment*, GC26-4191

   Provides specific guidance for using the 3380 in a VM/SP, VM/SP HPO, or VM/XA SF operating environment. This manual provides detailed instruction for planning the addition of new 3380 devices, installing devices, moving data to new devices, and performing ongoing storage management activities to maintain reliable performance and availability. In addition, storage considerations related to guest systems are addressed.

5. *Using the IBM 3380 Direct Access Storage in a VSE Environment*, GC26-4494

   Provides specific guidance for using the 3380 in a VSE operating environment. This manual provides instruction for planning the addition of new 3380 devices, installing devices, moving data to new devices, and performing ongoing storage subsystem management.

6. *Maintaining IBM Storage Subsystem Media*, GC26-4495

   Describes how the storage subsystem and the various operating systems handle disk storage errors and provides instruction on using the Environmental Record Editing and Printing (EREP) program and the Device Support Facilities (ICKDSF) program to diagnose and correct disk media errors. Recovery procedures are provided for the various device types. In addition, background material on DASD storage concepts is included.

7. *IBM 3380 Direct Access Storage Reference Summary*, GX26-1678

   Provides a summary of 3380 capacity, performance, and operating characteristics.

The Storage Control subset of the Storage Subsystem Library contains the following manuals:

1. *IBM 3990 Storage Control Introduction*, GA32-0098

   Provides a complete description of the various models of the 3990 Storage Control, including its data availability, performance, and reliability improvements over previous storage controls. In addition, this manual provides descriptions of the configuration attachment options, optional features, performance characteristics, and software support of the 3990 Storage Control.

2. *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*, GA32-0100

   Provides a functional description of the 3990 Storage Control. This manual describes the planning, program installation, and storage management tasks used in typical environments. Configuration examples as well as sample programs for controlling the various functions of the 3990 Storage Control are provided.

3. *IBM 3990 Storage Control Reference*, GA32-0099

   Provides descriptions and reference information for the 3990 Storage Control. This manual contains channel commands, error recovery, and sense information.

The *Storage Subsystem Library Master Index*, GC26-4496, provides a central source for information related to storage subsystem topics. Manuals for IBM 3380 Direct Access Storage, 3380 Direct Channel Attach Model CJ2, and 3990 Storage Controls are indexed in this publication. An overview of the material in the Storage Subsystem Library is provided with this index.

Figure 1 shows the relationships among the Storage Subsystem Library manuals in terms of high-level tasks described in each manual.

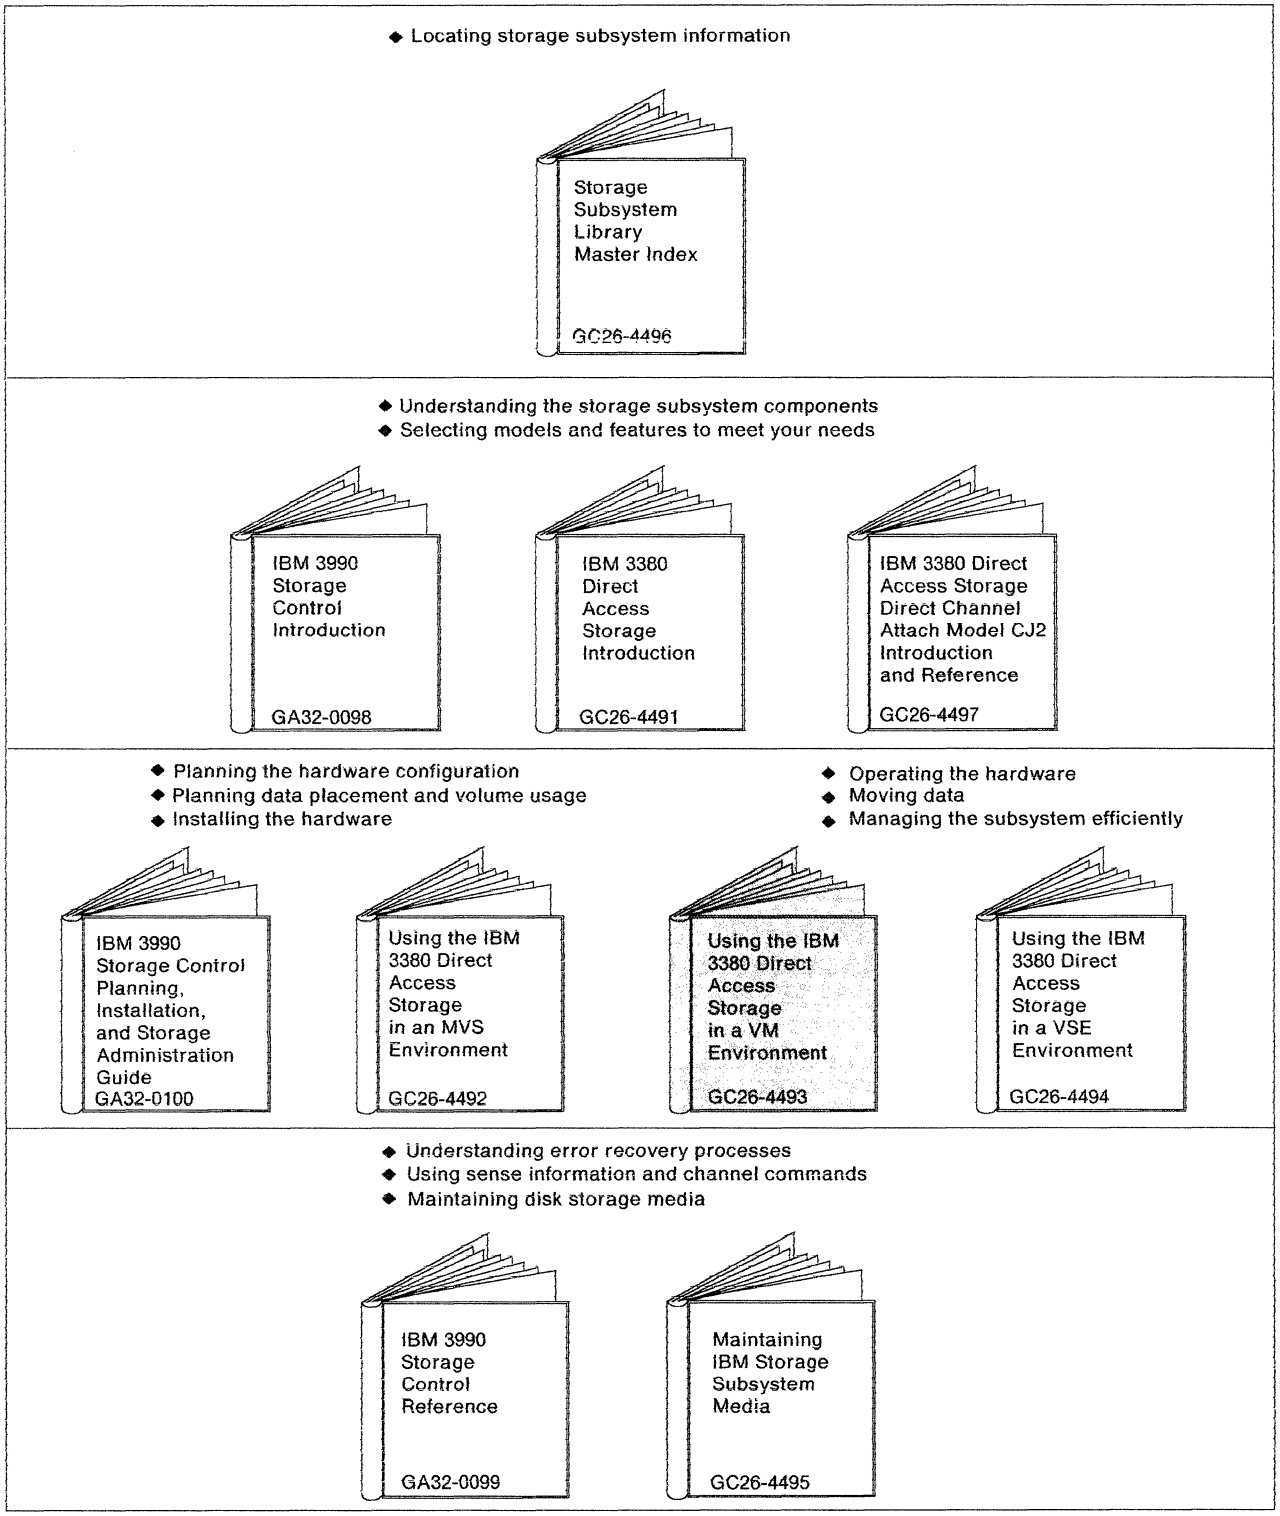

Figure 1. The Storage Subsystem Library

SSL Ordering Information

You can obtain a copy of **every manual** in the SSL using one General Bill of Forms (GBOF) number, **GBOF-1762**. Select one of the following GBOF numbers to obtain subsets of the SSL that provide information for specific environments and equipment. The GBOFs in the highlighted columns contain manuals for the VM environment. To obtain an individual SSL manual, use its order number.

Title	GBOF-1756	GBOF-1757	GBOF-1758	GBOF-1759	GBOF-1760	GBOF-1761	GBOF-0366
IBM 3380 Direct Access Storage Introduction, GC26-4491	X	X	X				
IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference, GC26-4497				X	X	X	
Using the IBM 3380 Direct Access Storage in an MVS Environment, GC26-4492	X			X			
Using the IBM 3380 Direct Access Storage in a VM Environment, GC26-4493		X			X		
Using the IBM 3380 Direct Access Storage in a VSE Environment, GC26-4494			X			X	
*Maintaining IBM Storage Subsystem Media,** GC26-4495	X	X	X	X	X	X	
IBM 3380 Direct Access Storage Reference Summary, GX26-1678	X	X	X	X	X	X	
IBM 3990 Storage Control Introduction, GA32-0098							X
IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide, GA32-0100							X
IBM 3990 Storage Control: Reference, GA32-0099							X
Storage Subsystem Library Master Index, GC26-4496	X	X	X	X	X	X	X

* *Device Support Facilities: Primer for the User of IBM 3380 Direct Access Storage*, GC26-4498, is distributed with this manual.

Related Publications

These publications (not part of the Storage Subsystem Library) describe the IBM 3880 Storage Control models to which the 3380 device can also attach:

- *IBM 3880 Storage Control Models 1, 2, 3, and 4 Description Manual*, GA26-1661
- *Introduction to IBM 3880 Storage Control Model 13*, GA32-0062
- *IBM 3880 Storage Control Model 13 Description*, GA32-0067
- *IBM 3880 Storage Control Model 23 Introduction*, GA32-0082
- *IBM 3880 Storage Control Model 23 Description*, GA32-0083

Device Support Facilities: Primer for the User of IBM 3380 Direct Access Storage, GC26-4498, is a new publication intended for use with 3380 manuals in the Storage Subsystem Library.

Other publications referenced in this manual or providing additional related information are included in "Bibliography" on page 127. To help you assess the potential usefulness of each reference, the bibliography includes a short description of the relevant contents of each publication.

Summary of Changes

Library Structure

This is a new manual. Some of the material in this manual was formerly contained in the following manuals, which are no longer available:

- *IBM 3380 Direct Access Storage: Planning and Use*, GC26-4208
- *IBM 3380 Direct Access Storage: Migration*, GC26-4197

The reformatting of this material is part of an overall restructuring of disk storage documentation, as shown below.

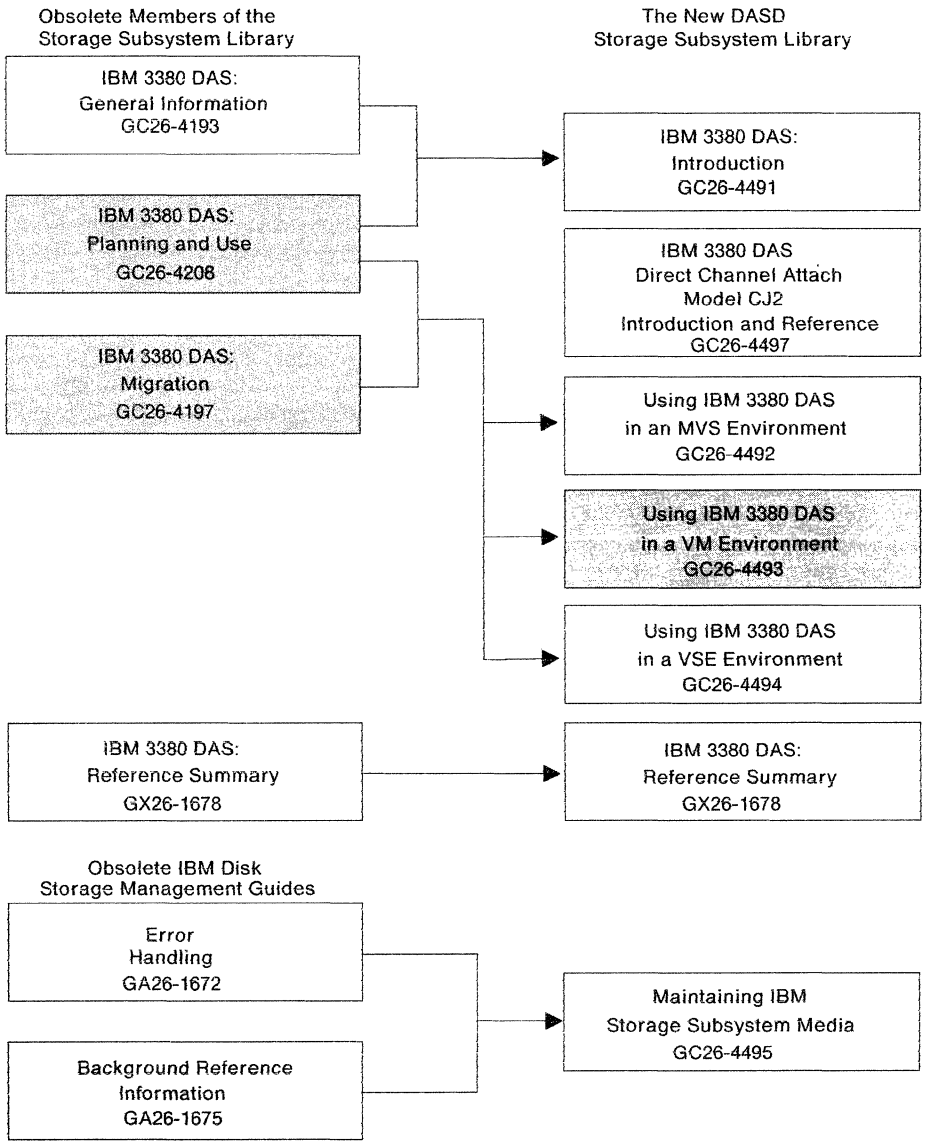

Technical Updates

Technical additions have been made to the original material to describe the following new hardware and hardware features:

- IBM 3380 Direct Access Storage Models AJ4, BJ4,
- The IBM 3380 Direct Channel Attach Model CJ2
- The 3380 AJ4/AK4 Support features #3005 for the IBM 3880 Storage Control Model 3 and #3010 for the IBM 3880 Storage Control Model 23
- IBM 3990 Storage Control Models 1 and 2

Contents

Figures

Chapter 1. Introducing the IBM 3380 Direct Access Storage

The IBM 3380 Direct Access Storage is a high-speed disk storage device designed for online data storage. It offers quick and reliable access to data and is a cost-effective way to expand your online storage capabilities.

This chapter provides the following overview information for using 3380 disk storage in the VM environment:

- A brief description of the 3380 family
- Hardware attachment requirements for 3380 models

3380 Direct Access Storage Overview

The twelve models of 3380 Direct Access Storage form four model groups. Several models offer increased storage capacity. However, all models of the 3380 have the same number of bytes per track and the same number of tracks per cylinder.

Volume Size	Standard Model Group	Extended Capability Model Group	Enhanced Subsystem Model Group	Direct Channel Attach Model
Single Capacity	A04 B04 AA4	AD4 BD4	AJ4 BJ4	CJ2
Double Capacity		AE4 BE4		
Triple Capacity			AK4 BK4	

Compared to the Standard models, the **Extended Capability** model group offers improved performance and availability and can provide concurrent transfer of data on two paths, from any two devices in the string. In addition, Models AE4 or BE4 provides twice the capacity of a Standard model.

The members of the **Enhanced Subsystem** model group are the most recent members of the 3380 family and provide further performance and availability enhancements. For example, the Enhanced Subsystem models have a faster seek time than other 3380 models, and when attached to the 3990 Storage Control Models 2 or 3, these 3380s can provide higher performance and availability for your data with concurrent data transfer on four paths. Furthermore, the triple capacity of the AK4 and BK4 models offers the 3380 family's lowest total cost of ownership on a per-megabyte basis.

The new 3380 **Direct Channel Attach** Model CJ2 provides both 3380 direct access storage and storage control functions in a single unit called a "C-unit". The direct access storage functions of the Model CJ2 provide two paths for concurrent data transfer, and have improved seek characteristics over 3380 Standard and Extended Capability Models. The inclusion of a two-path storage control function means that this 3380 model can be directly attached to a host processor channel.

As many as three 3380 B-units can be attached to an A-unit to form a string. As many as three Enhanced Subsystem B-units may be attached to a Model CJ2. The

Model A04 provides a single path string, while the Model AA4 and Extended Capability models form 2-path strings. The Enhanced Subsystem models can be configured as either 2-path or 4-path strings. A 4-path string consists of a minimum of two A-units, and can have as many as 3 B-units attached to each A-unit. The Model CJ2 and its attached B-units form 2-path strings.

This manual describes how to use A-units, B-units, and C-units in a VM environment. A complete description of the characteristics, features, capabilities, and string configurations of the various models of the 3380 can be found in the appropriate *Introduction* manual:

3380	**A04, AA4, B04, AD4, BD4, AE4, BE4, AJ4, BJ4, AK4, BK4**	*IBM 3380 Direct Access Storage Introduction*
3380	**CJ2**	*IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference*

Attaching the 3380 to Storage Controls and Processors

Figure 2 on page 3 shows how the 3380 can be attached to IBM processors through various models of the 3880 and 3990 storage controls. Note that the table reflects only those devices (3380 models, storage controls, and processors) currently supported in one or more of the VM operating environments. For information on all the possible DASD-storage control-processor combinations, see *IBM 3380 Direct Access Storage Introduction*. For a list of 3380 models supported in various VM operating environments, see "Identifying Programming Requirements" on page 7.

In some cases, it may be necessary to install features on your 3380s or storage controls. These features are explained in *IBM 3380 Direct Access Storage Introduction*.

String Type	IBM Storage Control Model Attachment	IBM Processors
A04	3880 Model 2 (Storage Director 2) 3880 Model 3	Standard Processors[1] 4381, 308x, and 3090 Processors[3] 4361 Processor[4] System/370 Models 158, 158-3 System/370 Models 168, 168-3 9375 and 9377 Processors
AA4	3880 Model 2 (Storage Director 2) 3880 Model 3 (AA4 string(s) only)	Standard Processors[1] 4381, 308x, and 3090 Processors[3] 4361 Processor[4] System/370 Models 158, 158-3 System/370 Models 168, 168-3 9375 and 9377 Processors
	3880 Model 3 (with AD4 or AE4 string)	Standard Processors[1] 4381, 308x, and 3090 Processors[3] 4361 Processor[4] 9375 and 9377 Processors
	3880 Model 13, 23	Standard Processors[1] 4381, 308x, and 3090 Processors[3] 9375 and 9377 Processors
	3990 Models 1, 2[2]	4381, 308x, and 3090 Processors[3] 9375 and 9377 Processors
AD4 AE4	3880 Model 3	Standard Processors[1] 4381, 308x, and 3090 Processors[3] 4361 Processor[4] 9375 and 9377 Processors
	3880 Model 23	Standard Processors[1] 4381, 308x, and 3090 Processors[3] 9375 and 9377 Processors
	3990 Models 1, 2	4381, 308x, and 3090 Processors[3] 9375 and 9377 Processors
AJ4 AK4	3880 Models 3	4381, 308x, 3090 Processors[3] 4361 Processor 9375 and 9377 Processors
	3880 Models 23	4381, 308x, 3090 Processors[3] 9375 and 9377 Processors
	3990 Models 1, 2	4381, 308x, 3090 Processors[3] 9375 and 9377 Processors
CJ2	Not required, attaches directly to channel	4381, 308x, 3090 Processors[3] 9375 and 9377 Processors

Figure 2. IBM Processors and 3880 Storage Control Models Compatible with 3380 Strings

Notes to Figure 2:

1 The standard processors include 3031, 3032, 3033, 4341, and the 3042 Attached Processor Model 2

2 Attachment to a 3990 requires the Model AA4 have a serial number of 15000 or greater for 60 Hz units or X0300 or greater for 50 Hz units

3 The 3090 Models 300 and 600 run VM/XA or MVS/XA only

4 Supported on 3-megabyte channel

Chapter 2. Planning for Installation

Adding new devices to your system requires adequate planning to ensure the process runs smoothly. Device installation *can* be streamlined, but you must be willing to spend the time up front to develop a detailed project plan.

Many of the problems found during device installation are typically things that could and should have been identified long before the device arrived on the loading dock. These kinds of "surprises"—inadequate cable lengths, missing prerequisite features, user program device dependencies, or insufficient software support—cause delays that inhibit both user and system productivity. With careful and comprehensive planning, you can reduce the surprises at installation time to a minimum.

The major stages of installation planning are described in this chapter:

- Physical planning
- Software planning
- Administrative planning

All of these stages should be addressed in your project plan.
Appendix A, "Sample DASD Installation and Migration Project Plan" on page 105
contains a sample device installation project plan that you may want to use as a model.

Physical Planning

Before you install new 3380 devices, there are a number of physical planning activities that you should complete. Use the information about physical characteristics of the 3380 in either *IBM 3380 Direct Access Storage Introduction* or *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference* to help you identify requirements for:

Floor space, electrical power, and cables
I/O addresses
Prerequisite features

Identifying Floor Space, Power, and Cable Requirements

The first tasks in physical planning are to identify and reserve floor space, arrange electrical power, and order cables for the new 3380. For more information about these tasks, see *IBM I/O Equipment: Installation—Physical Planning for S/360, S/370, and 4300 Processors* and *9370 Information System Installation Manual—Physical Planning*. Work with your installation planning representative to determine the physical requirements of 3380 installation.

Reserving I/O Addresses

You must assign I/O addresses to all 3380 devices you are going to install on your system. These I/O addresses must be defined to the 3380s, 3880s and/or 3990s, and to the operating system (and Input/Output Control Program when present) during the installation process. At this time, decide which addresses will be assigned to the 3380 units you will install. For details about how to specify I/O addresses for the 3380 units, see either *IBM 3380 Direct Access Storage Introduction* or *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference.* For details about how to specify I/O addresses for the 3880s and/or 3990s, see *IBM 3880 Storage Control Models 1, 2, 3, and 4 Description Manual* or *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide.* For information about defining I/O addresses to an operating system, see Chapter 6, "Installing the 3380 under VM" on page 61.

Nondisruptive DASD Installation

In subsystems configured for 4-path strings, the 3380 Enhanced Subsystem models can be installed without disrupting the availability of other storage resources. Models BJ4 and BK4 can be added to an existing string, or a second Model AJ4 or AK4 4-path string can be added without disrupting the availability of the existing string. Since this function is handled totally within the DASD subsystem, any operating system, Extended Architecture or System/370, that is connected to a 4-path subsystem can take advantage of this capability.

To take advantage of the nondisruptive DASD installation capability, you must pre-assign addresses for devices planned for future installation and predefine these addresses to the operating system. Defining future DASD additions in present system generations will avoid future changes and IPLs.

For more information on nondisruptive DASD installation, see *IBM 3380 Direct Access Storage Introduction* and *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide.*

Identifying Prerequisite Features

Before the 3380 arrives, you may also need to apply prerequisite features to your existing storage subsystem hardware. Check with your service representative to make sure all of your system's hardware has been updated with the necessary features. If you must apply any additional features, be sure to schedule system time to do so before installing the 3380.

For a description of prerequisite features, see *IBM 3380 Direct Access Storage Introduction*, or the appropriate storage control manual.

Software Planning

In addition to physical planning, you should do thorough software planning to ensure that your system software will support the 3380 after it is installed. If you don't have the appropriate levels of software, contact your IBM representative to order the required licensed programs and software maintenance.

Install new or changed software before the 3380 arrives, and use the remaining time to become familiar with the programs. The more comfortable you are with the software that supports the new devices, the easier device installation will be.

Identifying Programming Requirements

The 3380 is supported in three VM operating environments:

- Virtual Machine/System Product (VM/SP), licensed program 5664-167
- Virtual Machine/System Product High Performance Option (VM/SP HPO), licensed program 5664-173
- Virtual Machine/Extended Architecture Systems Facility (VM/XA SF)[1], licensed program 5664-169

The tables that follow identify the ***minimum required*** programming support for the 3380. Although these tables list only the minimum version and release levels required for support of the 3380 models, you should always use the most current level available to take advantage of the latest product enhancements. For information on the most current release levels of IBM licensed programs, contact your IBM representative.

---

[1] VM/XA SF replaces the Virtual Machine/Extended Architecture Migration Aid, licensed program 5664-169.

Programming Support in VM/SP

All models of the IBM 3380 can be used in a native VM/SP environment or on a guest operating system running under VM/SP. Minimum release levels of VM/SP and related programs for 3380 and attached storage control support are shown in Figure 3.

The programs referred to in this table are:

Virtual Machine/System Product (VM/SP), 5664-167
Device Support Facilities (ICKDSF), 5747-DS1
Environmental Recording, Editing, and Printing (EREP), 5654-260

3380 Model	Attached to	VM/SP Release	ICKDSF Release	EREP Version
A04	3880 Model 2 or 3	1.0	7.0	3.1.0
AA4	3880 Model 2 or 3	1.0	7.0	3.1.0
	3990 Model 1	4.0 with PTF	9.0	3.3.2
	3990 Model 2	4.0 with PTF	9.0	3.3.2
AD4	3880 Model 3	3.0 with PTF	7.0	3.1.0 with Feature 2
	3990 Model 1	4.0 with PTF	9.0	3.3.2
	3990 Model 2	4.0 with PTF	9.0	3.3.2
AE4	3880 Model 3	3.0 with PTF	8.0	3.1.0 with Feature 2
	3990 Model 1	4.0 with PTF	9.0	3.3.2
	3990 Model 2	4.0 with PTF	9.0	3.3.2
AJ4, AK4	3880 Model 3	4.0 with PTF	9.0	3.3.2
	3990 Model 1	4.0 with PTF	9.0	3.3.2
	3990 Model 2	4.0 with PTF	9.0	3.3.2
CJ2	Attaches directly to channel	4.0 with PTF	9.0	3.3.2

Figure 3. Minimum Programming Support in the VM/SP Environment

In addition, if you use VSAM files on 3380s under CMS, you need Release 3.0 of the Virtual Storage Extended/Virtual Storage Access Method (VSE/VSAM) program (5746-AM2), with its current modifications.

See your IBM marketing representative where program temporary fixes (PTFs) or program update tapes (PUTs) are required. For more information on PTFs, see "Applying Software Changes" on page 10.

Programming Support in VM/SP HPO

All models of the IBM 3380 can be used in a native VM/SP HPO environment or on a guest operating system running under VM/SP HPO. Minimum release levels of VM/SP HPO and related programs for 3380 support are shown in Figure 4.

The programs referred to in this table are:

Virtual Machine/System Product High Performance Option (VM/SP HPO), 5664-173
Device Support Facilities (ICKDSF), 5747-DS1
Environmental Recording, Editing, and Printing (EREP), 5654-260

3380 Model	Attached to	VM/SP HPO Release	ICKDSF Release	EREP Version
A04	3880 Model 2 or 3	1.0	4.0	3.1.0
AA4	3880 Model 2 or 3	1.0	4.0	3.1.0
	3880 Model 13 or 23	4.0[1]	4.0	3.1.0
	3990 Model 1	4.2 with PTF	9.0	3.3.2
	3990 Model 2	4.2 with PTF	9.0	3.3.2
AD4	3880 Model 3	3.2	7.0	3.1.0 with Feature 2
	3880 Model 23	4.0[1]	7.0	3.1.0 with Feature 2
	3990 Model 1	4.2 with PTF	9.0	3.3.2
	3990 Model 2	4.2 with PTF	9.0	3.3.2
AE4	3880 Model 3	3.2 with PTF	8.0	3.1.0 with Feature 2
	3880 Model 23	4.0[1]	8.0	3.1.0 with Feature 2
	3990 Model 1	4.2 with PTF	9.0	3.3.2
	3990 Model 2	4.2 with PTF	9.0	3.3.2
AJ4, AK4	3880 Model 3	4.2 with PTF	9.0	3.3.2
	3880 Model 23	4.2 with PTF	9.0	3.3.2
	3990 Model 1	4.2 with PTF	9.0	3.3.2
	3990 Model 2	4.2 with PTF	9.0	3.3.2
CJ2	Attaches directly to channel	4.2 with PTF	9.0	3.3.2

Figure 4. Minimum Programming Support in the VM/SP HPO Environment

Note to Figure 4:

[1] To support the 3880 Storage Control Model 13 or 23 you need either:

VM/SP HPO Release 4.0 or later; or

VM/SP HPO Release 3.2, 3.4, or 3.6 with the program offering, VM/SP HPO CMS Support for the IBM 3880 Model 13 and 23 (5798-DRJ).

In addition, if you use VSAM files under CMS, you need Release 3.0 of the Virtual Storage Extended/Virtual Storage Access Method (VSE/VSAM) program (5746-AM2), with its current modifications.

See your IBM marketing representative where program temporary fixes (PTFs) or program update tapes (PUTs) are required. For more information on PTFs, see "Applying Software Changes" on page 10.

Programming Support in VM/XA SF

All models of the IBM 3380 can be used in a native VM/XA SF environment or on a guest operating system running under VM/XA SF. Minimum release levels of VM/XA SF and related programs for 3380 support are shown in Figure 5.

The programs referred to in this table are:

Virtual Machine/Extended Architecture Systems Facility (VM/XA SF), 5664-169
Device Support Facilities (ICKDSF), 5747-DS1
Environmental Recording, Editing, and Printing (EREP), 5654-260

3380 Model	Attached to	VM/XA SF Release	ICKDSF Release	EREP Version
A04	3880 Model 2 or 3	1.0	7.0	3.2.0
AA4	3880 Model 2 or 3	1.0	7.0	3.2.0
	3880 Model 13 or 23	1.0 with PUT	4.0	3.2.0
	3990 Model 1	2.0 with PTF	9.0	3.3.2
	3990 Model 2	2.0 with PTF	9.0	3.3.2
AD4	3880 Model 3	1.0	7.0	3.2.0
	3880 Model 23	1.0 with PUT	7.0	3.2.0
	3990 Model 1	2.0 with PTF	9.0	3.3.2
	3990 Model 2	2.0 with PTF	9.0	3.3.2
AE4	3880 Model 3	1.0 with PUT	8.0	3.2.0
	3880 Model 23	1.0 with PUT	8.0	3.2.0
	3990 Model 1	2.0 with PTF	9.0	3.3.2
	3990 Model 2	2.0 with PTF	9.0	3.3.2
AJ4, AK4	3880 Model 3	2.0 with PTF	9.0	3.3.2
	3880 Model 23	2.0 with PTF	9.0	3.3.2
	3990 Model 1	2.0 with PTF	9.0	3.3.2
	3990 Model 2	2.0 with PTF	9.0	3.3.2
	3990 Model 3	2.0 with PTF	9.0	3.3.2
CJ2	Attaches directly to channel	2.0 with PTF	9.0	3.3.2

Figure 5. Minimum Programming Support in the VM/XA SF Environment

See your IBM marketing representative where program temporary fixes (PTFs) or program update tapes (PUTs) are required. For more information on PTFs, see "Applying Software Changes."

Applying Software Changes

If there are program temporary fixes (PTFs) outstanding for your software that affect device support, apply them before you install the 3380 devices. Work with your IBM account team to determine if there are any known problems that could affect device installation, and plan to apply the changes when you install the software. You may need to install a new program update tape to support the new devices. The advance planning you do here can save you a lot of time and effort later in problem determination.

IBM Service maintains an online *hardware install index* that lists the software programs that have been modified to support the 3380. The index also includes pointers to existing preventive service planning (**PSP**) information, describing problems previously encountered in specific operating environments. Contact your

service representative to request information from the index that might help you plan for 3380 installation.

For instructions on how to apply PTFs and PUTs, see *VM/SP HPO Installation Guide*.

Note: If you are using a 3310 or 3370 device as the VM system residence volume on a VM/SP or VM/SP HPO system, you ***must*** apply the fix for APAR VM22450 (available on PUT tape 8505 or later) to your system before you install the 3380 AD4/AE4 Support feature on any storage controls in that system.

Identifying Software Tools

There are a number of software tools available under VM to identify data, move data to new devices, and manage storage on new devices. If you are planning to use any of these tools, be sure they have been installed and tested prior to device installation. To order new tools, contact your IBM representative.

VM/Directory Maintenance (DIRMAINT)

Licensed program number: 5748-XE4

VM/Directory Maintenance is an IBM licensed program for VM/SP and VM/SP HPO systems that you can use to update user and minidisk information in the VM directory and, if necessary, to relocate minidisks. Among the tasks that VM/Directory Maintenance can perform are:

- Listing minidisks and minidisk extent allocation
- Moving data by minidisk
- Adding and deleting minidisks
- Transferring minidisk ownership
- Defining, changing, and deleting userids

For an overview of the capabilities of VM/Directory Maintenance, see *VM/Directory Maintenance General Information*.

Virtual Machine Monitor Analysis Program (VMMAP)

Licensed program number: 5664-191

VMMAP is an IBM licensed program that processes performance data generated by the VM/Monitor for VM/SP and VM/SP HPO systems. You can use VMMAP reports to help identify performance bottlenecks, monitor system utilization, and track the effects of system tuning. The historical performance data recorded by VMMAP can help you make cost-effective hardware decisions, plan for orderly and timely system changes, and project future capacity requirements.

"Using VMMAP to Review I/O Workload" on page 21 illustrates how VMMAP reports can be used to collect performance statistics and evaluate the I/O subsystem. Using the Graphical Data Display Manager (GDDM, 5748-XXH/01) and Presentation Graphics Feature (PGF, 5748-XXH/02) programs, you can also display and print color graphs of VMMAP data with the GDDM support provided by VMMAP.

For an overview of the capabilities of VMMAP, see *VMMAP General Information*.

VM Real Time Monitor (VM/RTM)

Program offering number: 5796-PNA

VM/RTM is an IBM program offering that collects and processes performance data generated by CP for VM/SP and VM/SP HPO systems. You can use VM/RTM reports to help identify performance bottlenecks, monitor system utilization, and track the effects of system tuning. VM/RTM has online interactive trace facilities, and can be timer-driven for performance measurement.

"Using VM/RTM to Review I/O Workload" on page 30 illustrates how VM/RTM reports can be used to collect performance statistics and evaluate the I/O subsystem.

For an overview of the capabilities of VM/RTM, see *RTM Program Description/Operations Manual.*

VM/XA Realtime Monitor/Systems Facility (RTM/SF)

Program offering number: 5798-DWD

RTM/SF is an IBM program offering that collects and processes performance data generated by CP for VM/XA SF systems. You can use RTM/SF reports to help identify performance bottlenecks, monitor system utilization, and track the effects of system tuning. RTM/SF has online interactive debugging facilities, and can be timer-driven for precise performance measurement.

For an overview of the capabilities of RTM/SF, see *RTM/SF Program Description/Operations Manual.*

VM Performance Planning Facility (VMPPF)

Licensed program number: 5664-179

VMPPF is an IBM licensed program used for performance measurement and capacity planning on VM/SP and VM/SP HPO systems. You can use VMPPF to create a mathematical model of your current VM system and to project the effects of changes in workload or hardware configuration. VMPPF can also model the effects of shared DASD under VM.

VMPPF provides an online user interface and extensive graphics display facilities to assist you in performance and capacity planning.

For an overview of the capabilities of VMPPF, see *VMPPF General Information.*

VMBACKUP Management System

Licensed program number: 5664-291

VMBACKUP is an IBM licensed program used to make incremental backup copies of CMS user minidisks and files on VM/SP and VM/SP HPO systems. You can use VMBACKUP to dump data to tape and to restore it to DASD when necessary. VMBACKUP also provides an interactive screen display of backup copies so that users can restore their own files without the assistance of the VM system programmer.

The VMArchive Subsystem of VMBACKUP allows you to archive unused CMS files that are occupying space on user minidisks. Files can be copied to predefined archive areas on DASD or tape. When needed, archived files can be recalled directly to a user's reader or minidisk.

For an overview of the capabilities of VMBACKUP and VMArchive, see *VMBACKUP Management System General Information.*

Additional Software Support

The following program products help you manage and share data on 3380s in a VM environment.

Resource Access Control Facility/VM (RACF/VM)

Licensed program number: 5740-XXH (base) plus Feature 4577 (VM support)

RACF/VM consists of the IBM licensed program RACF along with the supplement that supports VM. You can use RACF/VM to manage data, system, and resource security for VM/SP and VM/SP HPO systems. For an overview of the capabilities of RACF/VM, see *RACF General Information.*

VM/Inter-System Facilities (VM/ISF)

Licensed program number: 5664-376

VM/ISF allows users on two separate processors running VM/SP HPO Release 4.2 to access data files on shared minidisks, as well as share spool files and directories. Users can log on to either system and access their reader files. Certain defined files can be accessed via shared mini-disk as set up by the system administrator. For more information, see *VM/ISF General Information.*

Identifying Publication Requirements

As you begin to plan for the new 3380 devices, you may need additional reference publications. "Bibliography" on page 127 lists publications for the topics that are relevant to device migration under VM. Contact your IBM representative to order copies of IBM publications.

Administrative Planning

In addition to physical and software planning, there are a number of administrative planning tasks that are vital to making device installation and data movement a success. These tasks are summarized in the following sections.

Creating a Project Team and Project Plan

To handle the complex installation and migration process, you will need to establish an installation and migration team from among your data processing personnel. Work with your management and your IBM representatives to determine the appropriate representatives for the device installation and migration team at your data processing center. Consider including:

- VM system programmers or administrators
- Hardware configuration specialists
- Operations specialists
- Local tools specialists
- IBM service and account representatives

It's important to keep accurate records of your planning activities so that you can anticipate problems in scheduling device installation and data movement. A written project plan can help you determine what needs to be done to install and use the new devices. Make sure the project plan is available for all members of the device installation and data movement team to use as they complete their assigned tasks. The project plan should:

- Identify all tasks and responsibilities, including dependencies on other groups or events
- Assign direct responsibility for each task
- List target dates for completion of each task
- Identify meaningful checkpoints to evaluate progress
- Include statements of commitment from all involved groups

Appendix A, "Sample DASD Installation and Migration Project Plan" on page 105 contains a sample device installation project plan that you may want to use as a model.

Verifying Device Delivery Schedules

Contact your IBM representative to verify that your new devices have been ordered, and to confirm the delivery date. Use this date to plan a tentative installation schedule for the 3380 units.

Scheduling System Time for Volume Preparation and Data Movement

In most cases, you will need to schedule system time to prepare volumes and to move or re-create data. Obviously, the scheduling constraints of each installation will be different, but it's important to make these arrangements well in advance. Keep in touch with your user groups to determine the best time to prepare volumes and to move key minidisks.

Training Personnel

Although installation of a new device should be transparent to most users, you will need to train some of your data processing personnel to use the new device. Operators will need to know how to power up a 3380 string and how to vary the devices online and offline; programmers may need to learn how to use available software tools to move data onto 3380 volumes.

You should plan hands-on training sessions for these people as soon as the 3380 units are available. For more information on the tasks that operators perform, see Chapter 7, "Operating the 3380 under VM" on page 79. For more information on moving data, see Chapter 8, "Moving Data onto 3380 Volumes" on page 83.

Documenting New Procedures

In addition to training your operators and programmers in the use of the 3380, you may need to document new operating and programming procedures for your user groups. These new procedures may include guidelines for:

- Minidisk allocation
- Data movement
- Data security
- Backup and recovery

Depending on the needs of your data processing center, you may want to document these policies and procedures as a printed user's guide, online briefing, or other type of presentation.

You should also plan to document the changes you make in your system configuration and in volume or minidisk placement. Any notes you make regarding migration problems, minidisk location and ownership, and application dependencies may provide valuable input for future device migrations.

Chapter 3. Understanding Your Data and Hardware Configuration

To make the best use of your new hardware, you should understand how each device fits into your overall system configuration and how its distinctive capabilities can be used to best advantage. The key to effective use of the 3380 is to match its physical characteristics to the logical requirements of your data. Like hardware, data should follow a planned *configuration*; the way it is grouped, stored, and retrieved should depend on its logical data needs just as the placement of hardware depends on physical device needs.

Before you introduce the 3380 into your VM storage subsystem, take a comprehensive inventory of the data and storage in your data processing center. You can identify requirements for specific levels of performance, availability, and space by grouping data into categories and discussing these categories with your users. Identify device-dependencies in user programs that might constrain efficient use of your devices. Measure the effectiveness of your current hardware configuration by evaluating access paths and I/O activity. The more thoroughly you understand your current environment, the easier it will be to assess the impact of installing the 3380 on your data, programs, and users.

Identifying Data Groups and Their Requirements

VM data is usually grouped according to the tools that originally formatted it, such as CP, CMS, VSE/VSAM, or MVS/SP. Depending on the format, this data may have various requirements for performance, availability, or space. You need to evaluate all data by format to match its requirements to the capabilities of your devices, and to determine how it should be moved.

Figure 6 illustrates the types of data found in a typical VM environment.

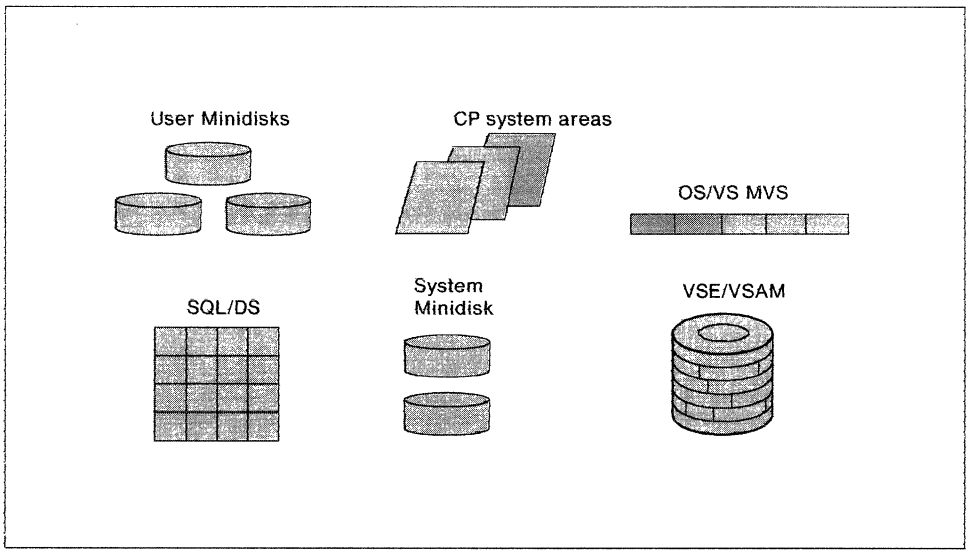

Figure 6. Types of Data in a VM Environment

CP System Areas

CP has 4K-byte pages of DASD storage allocated and formatted for specialized system data. These system areas include:

CP nucleus
Checkpoint area
Warm start area
Paging areas
Swapping areas (in VM/HPO 3.4 or later and VM/XA SF 2.0 or later)
Spooling areas
Dump areas
Override areas (in VM/SP 4.0 or later)
Saved system areas
Directory area

VM/XA does not have dump, override or saved system areas but has multiple directory areas.

Identify all system data areas in your environment and record their individual performance, availability, and space needs. Details on the special requirements of system data areas can be found in the appropriate VM *Planning* manual:

VM/SP	All releases	*VM/SP Data Areas and Control Block Logic-CP*
VM/SP	All releases	*VM/SP Planning Guide and Reference*
VM/SP HPO	All releases	*VM/SP HPO Planning Guide and Reference*
VM/XA SF	All releases	*VM/XA SF Virtual Machine Planning*

You should define dummy minidisks to represent CP system areas, using dummy userids in the directory. You can identify these areas using the VM/Directory Maintenance program, as shown in Figure 8 on page 19, or you can use the DISKMAP exec available with VM to list the contents of a specific volume. Figure 7 illustrates sample output from the DISKMAP exec. The word overlap in the comments column means that multiple minidisks are allocated to the same area.

```
 (1)      (2)      (3)  (4)    (5)   (6)   (7)    (8)

VMSY82  $ALLOC$   B01  3380   0000  0000  0001
        $SAVSYS$  B02  3380   0001  0017  0017
        CMSBATCH  195  3380   0018  0018  0001
        MAINT     296  3380   0019  0034  0016
        CMSUSER   191  3380   0035  0037  0003
        MAINT     496  3380   0038  0038  0001
        AUTOLOG1  191  3380   0039  0039  0001
                                40   326   287     GAP
        $TEMP$    C03  3380   0327  0426  0100
        DIRMAINT  195  3380   0427  0435  0009
                                436   870   435     GAP
        MAINT     295  3380   0871  0875  0005
        MAINT     124  3380   0000  0884  0885  **OVERLAP**
        SYSDUMP1  124  3380   0000  0884  0885  **OVERLAP**
```

Figure 7. Sample Listing of a CP Volume from DISKMAP. The columns indicate: (1) volume serial number, (2) userid, (3) virtual I/O address, (4) device type, (5) starting cylinder, (6) ending cylinder, (7) size in cylinders, (8) comments.

For more information on using DISKMAP, see the appropriate VM *Installation* manual:

VM/SP	All releases	*VM/SP Installation Guide*
VM/SP HPO	All releases	*VM/SP HPO Installation Guide*

Minidisks

CP divides DASD storage into variable slices called *minidisks.* Each minidisk is owned by a single CMS user, by the VM system, or by a guest running under VM. Although a minidisk contains files belonging to a single user, many other users may be allowed to access the files on that minidisk. Information about the minidisk, including the physical location, size, and owner, is stored in the VM system directory.

Minidisks take up most of the DASD space on a VM system. As part of your data inventory, you should generate a list of minidisks in your environment and determine the amount of space they occupy.

If you have the VM/Directory Maintenance program installed, you can use the DIRMAINT command to create a list of minidisks.

Note: The VM/Directory Maintenance program is not available with VM/XA.

Figure 8 shows a sample listing of minidisks created with the DIRMAINT command.

OWNERUID	CUU	DEVTYPE	DEVDESCR	START	SIZE	VLSER
CW1413	5A0	3380	3380	0	1	VM5800
CW1413	100	3380	3380	0	885	VM5800
CANFIG	191	3380	3380	62	10	VM5800
CHKCEL	191	3380	3380	72	2	VM5800
TRALLI	191	3380	3380	74	5	VM5800
TADWILL	191	3380	3380	79	20	VM5800
GDALLAS	191	3380	3380	99	2	VM5800
CW1413	510	3380	3380	101	220	VM5800
H47SOOP	193	3380	3380	321	15	VM5800
TBING	191	3380	3380	336	16	VM5800
BOSSIGER	191	3380	3380	352	4	VM5800
BIEMON	191	3380	3380	356	4	VM5800
FIENTES	191	3380	3380	360	1	VM5800
DISKUPS	192	3380	3380	361	30	VM5800
TONE	191	3380	3380	391	10	VM5800
YEATS	191	3380	3380	401	5	VM5800
SUNTAN	191	3380	3380	406	5	VM5800
H57COUP	191	3380	3380	411	10	VM5800
ITTWIG	191	3380	3380	421	5	VM5800
HAAKEPT	191	3380	3380	426	5	VM5800
KEPTALLA	191	3380	3380	435	5	VM5800
JOBMAN	191	3380	3380	440	5	VM5800
CULLEN	191	3380	3380	445	20	VM5800
BLOOM	191	3380	3380	465	3	VM5800

Figure 8. Sample Listing of Minidisks from DIRMAINT

For more information on DIRMAINT, see *VM/Directory Maintenance Installation and System Administrator's Guide*.

By talking to user groups, you should be able to identify all user minidisks in your environment and record their individual performance, availability, and space needs. While each user may have specific requirements, you should be able to group minidisks according to common needs.

Special-Purpose Minidisks

Some minidisks contain data that is critical to the smooth operation of the VM system or of key applications. The location of these minidisks has a dramatic effect on both response time and system throughput.

Performance-critical minidisks usually include:

- Data bases, such as SQL/DS, which may be very large and may also have high availability requirements
- System minidisks for CMS (for example, S- and Y-disks)
- Minidisks for licensed programs
- Service machine minidisks
- Common tools minidisks

Identify these special minidisks and list them separately from the CMS user minidisks.

Other Data: VSE/VSAM, MVS, and non-IBM

A typical VM system contains more than just CP system areas and CMS minidisks. It also includes data inherited from other systems and programs, which may be stored on minidisks or on dedicated guest operating system volumes. When identifying the data in your environment, don't forget to list separately data originally formatted by VSE/VSAM and MVS. This data will require special tools and techniques if you move it to 3380 volumes.

Don't overlook data created by non-IBM programs when making your data inventory. In many cases, this data can be moved successfully using IBM utilities, but some data may require special tools to move. Keep a list of this data to help you plan for migration to new devices.

Identifying Device-Dependencies in Data and Programs

A major consideration in device migration occur when application programs, procedures, and EXECs are tuned for a single device type and are then moved to a different device type. Sometimes these programs stop working; sometimes they continue to work, inefficiently; sometimes they appear to work, but produce incorrect output.

As part of your data inventory, you should identify the programs that contain device dependencies. Talk to application owners to determine which applications may be affected by changing device types. It's good programming practice to avoid coding device-dependencies in user programs; use this opportunity to reinforce that concept among your application programming groups.

Remember to check data and programs that run on guest operating systems under VM. The number of device-dependent applications in VM is typically quite low; device dependencies occur most often in programs that run in the guest operating environment. You can read more about guest operating systems and their effect on device migration in "Guest System Considerations" on page 45.

In addition to user programs, check for device dependencies in the VM system. Many parameters in DMKSYS and DMKSNT (for VM/SP and VM/SP HPO systems) and HCPSYS (for VM/XA systems) are device-dependent. For example, the VM/SP HPO swap set size differs depending on the device type. For details on how to modify these parameters, see the appropriate VM *Planning* manual:

VM/SP	All releases	*VM/SP Planning Guide and Reference*
VM/SP HPO	All releases	*VM/SP HPO Planning Guide and Reference*
VM/XA SF	All releases	*VM/XA SF Virtual Machine Planning*

Note: Application programs that handle their own track arithmetic should be checked to verify that they are using fullword (four-byte) arithmetic rather than halfword (two-byte). Because the 3380 Models AK4 and BK4 have more than 32K tracks, applications that rely on a two-byte, signed field to make track calculations will need to be modified.

Measuring the Effectiveness of Your Current Hardware Configuration

Part of understanding your current VM environment is identifying limitations caused by inadequate access paths and inefficient volume utilization. There are a number of tools available to help you evaluate the performance of your storage subsystem and identify the sources of I/O bottlenecks. Study a period of time that reflects the variability of the workload in your environment, and use the results to answer the following questions:

- Which devices have consistently high or low busy periods?
- Which channels have consistently high or low busy periods?
- Where are the peaks and valleys of I/O activity?

The answers to these questions can help you determine where to place new 3380 devices in your configuration.

Using VMMAP to Review I/O Workload

You can collect performance statistics on your current VM/SP or VM/SP HPO system by using the Virtual Machine Monitor Analysis Program (VMMAP) to generate reports from VM/Monitor data. VMMAP reports can help you identify performance bottlenecks in the current storage subsystem and plot trends in I/O workload.

Before running VMMAP reports, be sure you have the VM/Monitor facility turned on and collecting the appropriate statistics. "Collecting Performance Statistics with VM/Monitor" on page 93 describes how to set up VM/Monitor to collect performance data. For details on how to run VMMAP, and on the variety of

VMMAP reports and graphic facilities available, see *VMMAP User's Guide and Reference*.

Notes for Guest Systems

You should use VMMAP in conjunction with guest operating system performance tools to identify I/O bottlenecks for guest systems under VM. Use tools that have been developed specifically for the guest system to determine how it is performing in the VM environment. For example, use VSE/PT to collect information on the performance of VSE in the VM environment.

In the sections that follow, VMMAP examples illustrate a performance measurement scenario for a sample VM/SP system running on a 4381. VM/SP HPO scenarios would be different.

Reviewing System-Wide I/O Utilization

The first report you should look at to understand the level of performance in your I/O subsystem is the VMMAP **Monitor Statistical Summary**, illustrated in Figure 9.

```
FIRST DATE 04/18/87  TIME 14:58:57       MONITOR STATISTICAL SUMMARY
LAST  DATE 04/18/87  TIME 16:59:27

                                          TIME OF   TIME OF   STD
VARIABLE   AVERAGE   MINIMUM   MAXIMUM    MINIMUM   MAXIMUM   DEV   #OBS          DESCRIPTION
--------   -------   -------   -------    --------  --------  ---   -------   ----------------------------------------

LOGGED      233.00    233.00    233.00    14:59:26  14:59:26  0.00      241   # OF USERS LOGGED ONTO SYSTEM
ACTIVE      135.80    108.00    158.00    16:07:56  16:04:27  8.30      241   # OF USERS ACTIVE IN SAMPLE INTERVAL
  .           .         .         .         .         .       .        .
  .           .         .         .         .         .       .        .
  .           .         .         .         .         .       .        .
PCTCPUQ       5.60      0.00     22.22    15:01:26  15:41:56  4.64   32,728   PCT OF ACTIVE USERS IN CPU WAIT
PCTSTGQ       0.00      0.00      0.00    14:59:26  14:59:26  0.00   32,728   PCT OF ACTIVE USERS WAITING FOR MAIN STG
PCTPAGEQ      0.46      0.00      8.50    14:59:26  15:41:26  0.88   32,728   PCT OF ACTIVE USERS IN PAGE WAIT
PCTSWAPQ      0.00      0.00      0.00    14:59:26  14:59:26  0.00   32,728   PCT OF ACTIVE USERS IN SWAP WAIT
PCTIOQ        2.15      0.00      7.32    15:14:27  15:09:27  1.51   32,728   PCT OF ACTIVE USERS IN I/O WAIT
  .           .         .         .         .         .       .        .
  .           .         .         .         .         .       .        .
  .           .         .         .         .         .       .        .
VIO         118.99     78.37    155.43    15:31:26  15:11:26 14.01    7,230   # VIRTUAL I/O REQUESTS SIMULATED/SEC
IOINT       181.52    113.53    234.93    15:51:26  16:59:27 22.98    7,230   # REAL I/O INTERRUPTS PER SECOND
  .           .         .         .         .         .       .        .
  .           .         .         .         .         .       .        .
  .           .         .         .         .         .       .        .
PAGEQ         0.62      0.00     13.00    14:59:26  15:41:26  1.23      241   # USERS IN PAGE WAIT
SWAPQ         0.00      0.00      0.00    14:59:26  14:59:26  0.00      241   # USERS IN SWAP WAIT
IOQ           2.93      0.00      9.00    15:14:27  15:09:27  2.06      241   # USERS IN I/O WAIT
  .           .         .         .         .         .       .        .
  .           .         .         .         .         .       .        .
  .           .         .         .         .         .       .        .
```

Figure 9. VMMAP Report: Monitor Statistical Summary (OUTSTAT)

From the statistical summary report, you can get an overview of total system performance and of I/O performance in particular. The most interesting variables in this report are:

PCTPAGEQ The percent of active users, on average, who were waiting for completion of paging operations during the monitored period. The smaller the percentage, the less this resource is acting as a bottleneck. For example, PCTPAGEQ under 10% is OK.

PCTIOQ The percent of active users, on average, who were waiting for I/O during the monitored period. PCTIOQ is an overall measure of I/O contention. I/O can be tape I/O or DASD I/O. The higher the percentage, the less productive work your system is doing. As a general guideline, an average PCTIOQ value of 5% or higher signals a possible I/O bottleneck.

In the sample report, PCTIOQ is 2.15. Few users on this system are consistently waiting for I/O, but they are still waiting for it more often than they are waiting for main storage (PCTSTGQ) or paging (PCTPAGEQ).

VIO The number of virtual I/O requests per second. VIO is a good overall measure of I/O load; the higher the value, the more productive work your system is doing. Track the VIO measurements for your system over time and use these historical values to identify workload trends. "Plotting Trends in I/O Workload" on page 28 shows how to use VIO measurements to plot workload trends in graph form.

PAGEQ The number of concurrent users waiting for paging operations to complete during the sampled period.

SWAPQ The number of concurrent users waiting for swapping operations to complete during the sampled period.

IOQ The number of concurrent users waiting for I/O operations to complete during the sampled period. This figure includes guest virtual machines (for example, VSE or VS1) but does not include I/O queues internal to the guest.

Reviewing Channel Utilization

If the statistical summary report indicates that you might have an I/O contention problem, you should then look at more detailed reports on the I/O subsystem. Figure 10 illustrates a sample **Channel Activity Summary** report from VMMAP.

DATE 04/18/87 FROM 14:58:57 TO 16:59:37 CHANNEL ACTIVITY SUMMARY # SECONDS 7230

					----------IPL CHANNELS----------					
					I/O REQUESTS QUEUED				UTILIZATION	
CHANNEL	TOTAL I/O	RATE /SEC	MSEC/ I/O	PCT OF SYSTEM	AVG ACTIVE	AVG	MAX	PCT OBS WITH Q	NUMBER SAMPLES	PCT BUSY
1	262,895	36	3.2	21	0.0	0.0	0	0	3,614	11.6
2	324,174	45	3.0	25	0.0	0.0	0	0	3,614	13.5
3	279,845	39	3.2	22	0.0	0.0	0	0	3,614	12.2
4	384,463	53	2.9	30	0.0	0.0	0	0	3,614	15.4
5	29,530	4	5.6	2	0.0	0.0	0	0	3,614	2.3
TOTALS	1,280,907	177		100						

Figure 10. VMMAP Report: Channel Activity Summary (OUTCHAN)

From the channel activity summary report, you can get an idea of which channels are most heavily used and whether your channel activity is balanced. The most interesting variables in this report are:

RATE/SEC

The I/O rate per second handled by the channel. RATE/SEC is a measure of channel load, and should be relatively balanced across the channels used by DASD for similar kinds of activity (such as CMS). In addition, the PCT OF SYSTEM value should also be roughly balanced.

In the sample report, the load is fairly evenly distributed among channels 1 through 4, although channel 4 carries a slightly heavier load than the other three.

UTILIZATION PCT BUSY

The percent of the time the channel was busy handling I/O. PCT BUSY is a measure of channel utilization. As a general guideline, a channel that is more than 30% busy with CP or CMS activity puts a severe drain on other system resources and inhibits good response time. Channels used primarily for paging and swapping can be driven to a higher utilization rate.

In the sample report, all of the channels are well below 30% busy.

Reviewing Device Utilization

In addition to channels, you should look at the activity on your current devices. Figure 11 on page 25 illustrates a sample **Disk and Tape I/O Activity Summary** report where additional seeks analysis data was included. See "Collecting Performance Statistics with VM/Monitor" on page 93 for details on how to collect seeks analysis data.

In VMMAP Version 1 Release 1.4 , the device type column is modified to contain a field to further describe the device. For a 3380, the device type will be shown as 3380.M, where M will be:

0 to represent an A04, AA4, or B04
J to represent an AD4, BD4, AJ4, BJ4, or CJ2
E to represent an AE4 or BE4
K to represent an AK4 or BK4

```
DATE  04/18/87  FROM 14:58:57 TO 16:59:37        DISK AND TAPE I/O ACTIVITY SUMMARY     # SECONDS 7230

MONITOR HIGH FREQUENCY DEVICE SAMPLING RATE (SAMPLES PER INTERVAL) - IPL PROC: 15

                    OBSERVATIONS  <----- TOTAL I/O ----> ACTIVE  <PCT I/O> < % BUSY >  < ACTIVE I/O REQUEST QUEUE >  <SEEKS ANALYSIS>
  DEVICE                   PCT          RATE  MSEC/    I/O   ON   OF   IPL  NIPL  <FOR DEVICE>   <FOR  CTLU>    NUMBER   AVG-LEN
ADDR   TYPE   VOLSER  TOTAL ACT    ISSUED  /SEC  I/O    RATE   CH  SYS  PROC  PROC  AVG  MAX  PCT  AVG  MAX  PCT  NON-ZERO  ALL NO
----  ------  ------  ----- ---  --------  ----  -----  -----  ---  ---  ----  ----  ---  ---  ---  ---  ---  ---  --------  --- --

0140  3350    VMMD02    241 100    70,067    10   18.6     10   27    5  18.1        1.2    2    2  0.0    0    0       392   10  47
0141  3350    VMMD04    241 100    64,871     9   20.3      9   25    5  18.2        1.4    3    5  0.0    0    0       317   15  51
0142  3350    VMMD06    241 100    68,749    10   19.6     10   26    5  18.6        1.8    4    6  0.0    0    0       489   19  54
0143  3350    VMMD08    241 100    59,208     8   18.8      8   23    5  15.4        1.0    1    3  0.0    0    0       439   12  42
0280  3380.J  VMSP83    241 100    33,006     5   18.1      5   10    3   8.3        1.0    1    1  0.0    0    0       164    7  28
0282  3380.J  VMSP82    241 100    16,710     2   19.3      2    5    1   4.5        1.0    1    2  0.0    0    0        79   43 159
0283  3380.J  VMPG01    241 100   113,889    16   18.3     16   35    9  28.7        1.7    5   12  0.0    0    0     2,053    3   4
0287  3380.J  VMSY82    241 100   160,569    22   17.5     22   50   13  38.9        1.5    3   15  0.0    0    0     2,058    4   5
0380  3380.J  VMSY81    241 100   154,104    21   16.7     21   55   12  35.7        1.5    4   12  0.0    0    0    69,144    7  10
0381  3380.J  VMSP81    241 100    11,401     2   21.4      2    4    1   3.4        1.0    1    0  0.0    0    0        45   44 168
0382  3380.J  VMPG02    241 100   114,340    16   18.4     16   41    9  29.0        1.9    4    9  0.0    0    0     1,959    3   4
04C0  3375    VMPG03    241 100   111,656    15   21.5     15   29    9  33.2        2.3   10   20  0.0    0    0     2,067    5   6
04C1  3375    VMMD01    241 100    70,932    10   23.1     10   18    6  22.7        1.3    3    7  0.0    0    0       372    8  37
04C2  3375    VMMD03    241 100    67,254     9   23.9      9   17    5  22.2        1.4    3    6  0.0    0    0       490   17  58
04C3  3375    VMMD05    241 100    67,721     9   23.6      9   18    5  22.1        1.5    3    5  0.0    0    0       260    7  31
04C4  3375    VMMD07    241 100    66,900     9   23.4      9   17    5  21.6        1.5    3    4  0.0    0    0       222    7  38
----  ----    ------  ----- ---  --------  ----  -----  -----  ---  ---  ----  ----  ---  ---  ---  ---  ---  ---  --------  --- --
```

Figure 11. VMMAP Report: Disk and Tape I/O Activity Summary (OUTDASD)

From the I/O activity summary report with seeks analysis, you can get a good idea of which devices are most heavily used. The most interesting variables in this report are:

ACTIVE I/O RATE

The total I/O count divided by the number of intervals in which the device was active, expressed in seconds. This figure will show how heavily each device/VOLSER is utilized when it is in use. This tells you if the device is being used evenly (Active I/O Rate = Total I/O Rate) or if the usage is sporadic.

The I/O rate is averaged over the intervals that have I/O, intervals with zero I/O are ignored. The interval should be minimized to prevent the I/O rate from being averaged out over too long a time period.

In the sample report, the Active I/O Rate is the same as the Total I/O Rate, indicating that devices are not being used sporadically.

%BUSY IPL PROC

The percentage of time that the device was busy doing work for the IPL processor. (In a multiprocessor environment, there is a corresponding %BUSY NIPL PROC field for the non-IPL processor.) %BUSY is a measure of device utilization. A high %BUSY limits user response time for that volume.

ACTIVE I/O REQUEST QUEUE (FOR DEVICE)

The average number of I/O requests queued on this device during the monitored period. The size of the I/O queue is a measure of contention among users for this particular device. The AVG value is the average number of I/O requests queued for a device. The MAX value is the largest I/O request queue encountered. The PCT value is the frequency of request queuing for a device. The active I/O request queue determines (to a large extent) response time.

In addition to the average, check the percent of observations in which a queue for I/O was observed; but remember that *any* I/O queue is usually an indication of an I/O bottleneck. For example, an average

queue of 4.0 would be cause for concern, unless the percent column showed there was only a queue 5% of the time!

In the sample report, volume VMPG03, with an average queue of 2.3 requests 20% of the time, stands out as the most significant point of device contention in the I/O subsystem.

SEEKS ANALYSIS: NUMBER NONZERO

The number of I/O requests that generated DASD arm movement. Excessive seeks indicate that the arm is spending a lot of time crossing back and forth across the cylinders to retrieve data, adding unnecessary delay to I/O. This indicates data placement could be improved.

In the sample report, volume VMSY81 is by far the highest contributor to DASD arm movement with 69,144 nonzero seeks. To identify the data on volume VMSY81 that's causing these seeks, look at the detailed seeks analysis report for this volume (shown in Figure 12 on page 27).

SEEKS ANALYSIS: AVG LENGTH NON-0

The average number of cylinders that the DASD arm moved when handling the seeks recorded in the previous column. This value shows how sometimes even a few seeks can result in excessive I/O if the arm is consistently moving across the volume.

In the sample report, note how volumes VMSP82 and VMSP81—even though they have relatively few nonzero seeks—move the arm more than 150 cylinders on average when a seek is required. (150 cylinders is approximately 16% of a standard 3380 volume.) This may or may not indicate a seek problem, depending on the type of data and its usage.

Seeks Analysis reports can help you determine how data and minidisks could be reorganized to reduce DASD arm movement. Figure 12 on page 27 illustrates a sample **Seeks Analysis: Detail by Pack** report for the volume (VMSY81) that showed the largest number of nonzero seeks in the previous report.

Note: Seeks analysis produces large amounts of data. Route the monitor records to tape, not spool, to avoid filling the spool.

```
DATE  02/12/87   FROM 13:59:59 TO 14:20:00     S E E K S   A N A L Y S I S  DETAIL BY PACK    # SECONDS:  1078           PAGE    1
       START END                   TOTAL READ  % OF    % OF                      <---- %  O F   S Y S T E M --->
VOLSER  CYL  CYL  OWNER ID VADDR   SEEKS  PCT DEVICE  SYSTEM                    1        1        2        2        3
------  ---  ---  --------  ---   ------  --- ------  ------  0.:.:.:.:.5.:.:.:.:.0.:.:.:.:.5.:.:.:.:.0.:.:.:.:.5.:.:.:.:.0.:.:.:.:
------  ---  ---  --------  ---   ------  --- ------  ------

VMSY81    0  199  MAINT     190   11,900  100  54.72  54.727  :*************************************************************************
        200  399  MAINT     19E    6,453  100  29.67  29.677  :************************************************************
        400  599  MAINT     192        0       0.00   0.000  :
        600  799  MAINT     193    3,391  100  15.59  15.595  :*********************************
        800  999  MAINT     194        0       0.00   0.000  :
------  ---  ---  --------  ---   ------  --- ------  ------

                  *TOTAL*          21,744      100.00 100.000  :
DATE  02/12/87   FROM 13:59:59 TO 14:20:00     S E E K S   A N A L Y S I S  PACK SUMMARIES    # SECONDS:  1078           PAGE    2
       START END                   TOTAL READ  % OF    % OF                      <---- %  O F   S Y S T E M --->
VOLSER  CYL  CYL  OWNER ID VADDR   SEEKS  PCT DEVICE  SYSTEM                    1        1        2        2        3
------  ---  ---  --------  ---   ------  --- ------  ------  0.:.:.:.:.5.:.:.:.:.0.:.:.:.:.5.:.:.:.:.0.:.:.:.:.5.:.:.:.:.0.:.:.:.:
------  ---  ---  --------  ---   ------  --- ------  ------

VMSY81            DEVTOTAL         21,744  100 100.00 100.000  :
------  ---  ---  --------  ---   ------  --- ------  ------
                  SYSTOTAL         21,744      100.00 100.000  :
  .       .    .       .       .       .     .      .      .
  .       .    .       .       .       .     .      .      .
  .       .    .       .       .       .     .      .      .
DATE  02/12/87   FROM 13:59:59 TO 14:20:00     S E E K S   A N A L Y S I S  DETAIL BY USAGE   # SECONDS:  1078           PAGE    4
       START END                   TOTAL READ  % OF    % OF                      <---- %  O F   S Y S T E M --->
VOLSER  CYL  CYL  OWNER ID VADDR   SEEKS  PCT DEVICE  SYSTEM                    1        1        2        2        3
------  ---  ---  --------  ---   ------  --- ------  ------  0.:.:.:.:.5.:.:.:.:.0.:.:.:.:.5.:.:.:.:.0.:.:.:.:.5.:.:.:.:.0.:.:.:.:
------  ---  ---  --------  ---   ------  --- ------  ------

VMSY81    0  199  MAINT     190   11,900  100  54.72  54.727  :*************************************************************************
VMSY81  200  399  MAINT     19E    6,453  100  29.67  29.677  :************************************************************
VMSY81  600  799  MAINT     193    3,391  100  15.59  15.595  :*********************************
```

Figure 12. Seeks Analysis: Detail by Pack Report

With the Detail by Pack report, you can identify the data that is responsible for excessive seeks, as well as data that is a potential cache candidate. The most interesting variables in this report are:

VADDR The minidisk address, indicating the location of the data.

In the sample report, you can see immediately that minidisks MAINT 190 and MAINT 19E are the source of most of the seeks to this volume.

READ PCT

The percentage of the total seeks that were read operations. This is given only at the minidisk level and is only available with VM/HPO 4.0 or later.

Minidisks with high READ PCT values are potential candidates for placement behind storage controls with cache.

% OF DEVICE

The percentage of the seeks on this device represented by the seeks to this minidisk.

In the sample report, seeks to minidisk MAINT 190 account for about 54% of the seeks on this device; seeks to minidisk MAINT 19E account for another 29%.

% OF SYSTEM

The percentage of the seeks on the entire system represented by the seeks to this minidisk.

In the sample report, seeks to minidisk MAINT 190 account for about 54% of the seeks on the system; seeks to minidisk MAINT 19E account for another 29%.

You can see from this data that the placement of minidisks MAINT 190 and MAINT 19E on volume VMSY81 may be causing an I/O bottleneck on this system. You'll find some solutions for this particular data placement problem in "Placing the S- and Y-Disks: A Scenario" on page 54.

Plotting Trends in I/O Workload

You can also use VMMAP to draw graphs that illustrate the relationships between I/O variables. Graphs can help you identify trends more quickly and project their outcome. Graphs depicting workload, capacity, and performance trends can also be effective in convincing data processing managers to plan for and allocate additional computer resources.

For planning purposes, some graphs to examine are number of virtual I/O (VIO) requests per second versus time, and the number of active users versus time. The graphs shown in Figure 13 on page 29 illustrate the peaks and valleys of I/O activity during the sample monitoring period.

```
10/01/86 09:01:18  PLOT 3 - ACTIVE USER DISTRIBUTION BY TIME
X VARIABLE - TOD            TIME-OF-DAY
SYM Y-VARIABLE   SCALE  ORG   AVG       MAX          DESCRIPTION
--- ---------- ------- ---- ------- --------- ------------------------- -------
  *  ACTIVE       20.00       105.64    126.00 # USERS ACTIVE IN A SAMPLE INTERV
   X                                                                  1ST Y-VAR  2ND Y-VAR
 VALUE                                                             1  #OBS CUM%  #OBS CUM%
 ----- 0----1----2----3----4----5----6----7----8----9----0  ----- ---  ----- ---
 09:15 |.....................*                              |     15
 09:30 |.....................*                              |     15
 09:45 |......................*                             |     15
 10:00 |.......................*                            |     15
 10:15 |.......................*                            |     15
 10:30 |.......................*                            |     15
 10:45 |......................*                             |     15
 11:00 |.......................*                            |     15
 11:01 |........................*                           |      1
 ----- 0----1----2----3----4----5----6----7----8----9----0  ----- ---  ----- ---
 TOTAL                                                          121


10/01/86 09:01:18  PLOT 4 - VIO DISTRIBUTION BY TIME
X VARIABLE - TOD            TIME-OF-DAY
SYM Y-VARIABLE   SCALE  ORG   AVG       MAX          DESCRIPTION
--- ---------- ------- ---- ------- --------- ------------------------- -------
  *  VIO          100.00       550.12   1091.35 # VIRTUAL I/O REQUESTS SIMULATED/
   X                                                                  1ST Y-VAR  2ND Y-VAR
 VALUE                                                             1  #OBS CUM%  #OBS CUM%
 ----- 0----1----2----3----4----5----6----7----8----9----0  ----- ---  ----- ---
 09:15 |.....................*                              |    899
 09:30 |.....................*                              |    900
 09:45 |......................*                             |    900
 10:00 |.....................*                              |    900
 10:15 |........................*                          |    900
 10:30 |.......................*                            |    900
 10:45 |.......................*                            |    900
 11:00 |........................*                           |    900
 11:01 |.......................*                            |     60
 ----- 0----1----2----3----4----5----6----7----8----9----0  ----- ---  ----- ---
 TOTAL                                                          7259
```

Figure 13. VMMAP Graph: Active Users vs Time, VIO vs Time

Notice for this system that I/O activity peaked at 10:15 and 11:00 am during the monitoring period. For some environments, planning to handle the average I/O activity may be sufficient; for most others, the system must be configured to handle peak activity. The service objectives you establish for your own environment will help you determine the levels of I/O activity that you should plan to handle.

Using VM/RTM to Review I/O Workload

The VM Real Time Monitor (VM/RTM) program offering is a CMS application that displays and records I/O resource load information for VM/SP and VM/SP HPO systems.[2]

In addition to VMMAP, VM/RTM can provide assistance in identifying I/O bottlenecks and measuring system performance. Because it monitors performance characteristics in a real-time environment, VM/RTM is an effective problem determination tool.

VM/RTM runs in its own virtual machine. It has online interactive data collection facilities, and can be timer-driven to collect system or user data at regular intervals.

Among the VM/RTM reports that you can use are the following:

System Log
Records overall I/O contention and load by time interval.

State Resource Consumption
Records I/O contention (I/O wait time) by user.

Channel Display
Records channel contention, load, and utilization.

Device Display
Records DASD contention, load, and utilization.

For more information on VM/RTM, see *RTM Program Description/Operations Manual*.

In the sections that follow, VM/RTM examples illustrate a performance measurement scenario for a sample VM/SP system.

Reviewing System I/O Contention and Load

The first VM/RTM report you should look at to understand the level of performance in your I/O subsystem is the **System Log**, illustrated in Figure 14 on page 31. This report was generated in a VM/SP environment and will differ when generated in a VM/SP HPO environment.

---

2 VM/XA Realtime Monitor/Systems Facility (RTM/SF) is the equivalent program for VM/XA systems.

SYSTEM LOG

TOD VM/370 CPU4381 SERIAL 10022 8192K DATE 04/17/87 START 16:02:37 END 17:52:20

H.M	%CPU	%CP	%USR	%TWT	%PAG	%I/O	%IDL	%STO	ISEC	PSEC	XPG	PPAG	USR	IQ	WQ	ACT	Q1SE
1611	91	50	41	9	7	2	0	97	41	110	66	1604	233	69	28	211	.84
1617	83	50	33	17	17	0	0	86	35	176	67	1603	233	02	18	213	2.42
1622	85	51	34	15	15	0	0	72	56	234	71	1599	233	93	0	231	3.12
1627	82	51	31	18	18	0	0	89	47	233	73	1597	233	89	16	231	3.34
1633	80	54	26	20	20	0	0	90	42	224	79	1591	233	01	24	232	3.22
1644	83	51	32	17	17	0	0	71	42	201	80	1590	233	05	0	232	2.80
1649	84	55	29	16	16	0	0	89	44	219	81	1589	233	93	38	229	3.46
1655	82	51	31	18	18	0	0	95	46	221	83	1587	233	06	18	230	2.91
1701	78	50	28	22	22	0	0	90	38	209	86	1584	233	15	9	226	2.83
1707	81	54	27	19	19	0	0	90	41	232	87	1583	233	94	24	224	3.26
1713	79	51	28	21	21	0	0	92	41	211	87	1583	233	00	10	221	2.85
1721	83	51	32	17	17	0	0	87	40	200	89	1581	233	91	29	221	3.02
1725	84	54	30	16	16	0	0	95	47	237	88	1582	233	93	1	219	3.31
1731	81	52	29	19	19	0	0	93	46	222	90	1580	233	14	5	217	3.13
1738	81	51	30	19	19	0	0	96	38	190	92	1578	233	81	21	218	3.01
1742	88	51	37	12	11	1	0	42	51	194	95	1575	233	71	0	213	2.14
1746	95	53	42	5	3	2	0	73	58	180	95	1575	233	74	0	216	.84
1752	85	50	35	15	15	0	0	96	41	193	95	1575	233	65	18	217	1.06
AVG	88	52	36	12	12	0	0	76	47	189	94	1575	233	72	9	216	1.76

Figure 14. VM/RTM Report: System Log

From the system log, you can get an overview of total system performance and of I/O performance in particular. The most interesting fields in this report are:

%PAG The percent of time the system was waiting for paging during the monitored period. %PAG is an overall measure of paging contention; the higher the percentage, the less productive work your system is doing. As a general guideline, an average %PAG value of 10% or higher signals a possible paging bottleneck.

In the sample report, %PAG is 12, which indicates that the paging subsystem should be investigated.

%I/O The percent of time the system was waiting for I/O during the monitored period. %I/O is an overall measure of I/O contention; the higher the percentage, the less productive work your system is doing. As a general guideline, an average %I/O value of 10% or higher signals a possible I/O bottleneck.

In the sample report, %I/O is 0, so there are no obvious I/O problems on this system.

ISEC The number of non-spooled I/O requests per second. ISEC is an overall measure of I/O load; the higher the value, the less productive work your system is doing.

PSEC The number of paging requests per second. PSEC is an overall measure of paging load; the higher the value, the more productive work your system is doing.

Reviewing Channel Contention, Load, and Utilization

If the system log report indicates that you might have an I/O contention problem, you should then look at more detailed reports on the I/O subsystem. Figure 15 illustrates a sample **Channel Display** report from VM/RTM.

```
                                CHANNEL DISPLAY

VM/370   CPU4381   SERIAL 10022   8192K    DATE 04/17/87       START 12:13:56  END 12:31:03
<--- ACTUAL DATA --->   <IO SAMPLE>   <----------- 0% SAMPLED TRACE TABLE DATA ---------->
CH  TYPE IOREQST  SEC  MXB  MXW  MXQ  %DVB  %CUB  %CHB  %PC  %CUX  IOCC0  %BMX  %ERR  %CNT
00  MPX    *** CHANNEL IDLE ***
01  BMPX   19211   18    0    0    0  ....  ....     6  ...  ....     62     0  ....     6
02  BMPX   48954   47    1    0   12  ....  ....     1   67  ....    131     0  ....     1
03  BMPX   56556   55    2    0   15  ....  ....     5   80  ....    151    12  ....     5
04  BMPX   37448   36    2    0   27  ....     1    10  188     0     88    11  ....    11
05  BMPX    1818    1    0    0    0  ....  ....  ....  ...  ....     14  ....  ....  ....
```

Figure 15. VM/RTM Report: Channel Display

From the channel display report, you can get an idea of which channels are most heavily used and whether your channel activity is appropriately balanced. The most interesting fields in this report are:

SEC The I/O rate per second handled by the channel. SEC is a measure of channel load, and should be relatively balanced across channels used by DASD for similar activity (such as CMS).

In the sample report, channels 2 and 3 carry a heavier load than channel 4, and a much heavier load than channel 1. (Channel 5 is used only for tape drives, and is therefore not evaluated.)

MXW The maximum number of devices that were waiting for use of the channel or storage control during the recording period. MXW is a measure of channel contention.

In the sample report, MXW indicates no visible channel contention.

MXQ The maximum number of I/O operations that were queued on devices on this channel during the recording period. MXQ is a measure of device contention.

In the sample report, MXQ indicates slight contention for the devices on channels 2, 3, and 4.

Reviewing Device Contention, Load, and Utilization

In addition to channels, you should look at the activity on your current devices. Figure 16 illustrates a sample **Device Display** report from VM/RTM.

```
                                DEVICE DISPLAY

VM/370   CPU4381   SERIAL 10022   8192K  DATE 04/17/87    START 12:13:56  END 12:31:03
<------ ACTUAL DEVICE DATA ------> <-IO SAMPLE-> <--- SAMPLED TRACE TABLE IO DATA --->
DEV  TYPE  VOLSER IOREQST SEC  %CH %UT %WT MQ SW %DVB %CUB %CHB %PC  IOCC0  %ERR  %CNT
140  3350  VMMD02    5215   5   27   0   0  0  0 .... .... .... ...    27  ....     7
141  3350  VMMD04    5186   5   26   0   0  0  0 .... .... .... ...     4  ....  ....
142  3350  VMMD06    3992   3   20   0   0  0  0 .... .... .... ...    21  ....     9
143  3350  VMMD08    4730   4   24   0   0  0  0 .... .... .... ...    10  ....  ....
4C0  3375  VMPG03   16937  16   45 100   0 27  0 .... .... .... ...    48  ....     2
4C1  3375  VMMD01    3754   3   10   0   0  0  0 .... .... .... ...     5  ....    40
4C2  3375  VMMD03    5941   5   15   0   0  0  0 .... .... .... ...    12  ....    33
4C3  3375  VMMD05    7744   7   20  33   0  0  0 .... .... .... ...    18  ....    11
4C4  3375  VMMD07    3072   2    8   0   0  0  0 .... .... .... ...     5  ....    20
280  3380  VMSP83    2409   2    4   0   0  0  0 .... .... .... ...     4  ....    50
282  3380  VMSP82    1278   1    2   0   0  0  0 .... .... .... ...     1  ....  ....
283  3380  VMPG01   45267  44   92 100   0 12  0 .... .... .... ...   126  ....  ....
380  3380  VMSY81   23091  22   40  66   0  0  0 .... .... .... ...    60  ....    10
381  3380  VMSP81     927   0    1   0   0  0  0 .... .... .... ...  ....  ....  ....
382  3380  VMPG02   32538  31   57 100   0 15  0 .... .... .... ...    91  ....     2
```

Figure 16. VM/RTM Report: Device Display

From the device display report, you can get an idea of which devices are most heavily used. The most interesting fields in this report are:

SEC The I/O rate per second handled by the device. SEC is a measure of device load, and should be relatively balanced across devices used for similar activity (such as CMS).

In the sample report, devices 283, 380, and 382 carry a heavier load than most of the other devices. Two of these are paging volumes (VMPG01 and VMPG02), and one is a CP system volume (VMSY81). Heavy loads on these volumes may indicate possible I/O bottlenecks.

%CH The percent of I/O activity on the channel that is being handled by this device. This field will help you find out when one device is dominating a particular channel to the detriment of other devices.

In the sample report, devices 4C0, 283, 380, and 382 contribute more than 40% of the I/O to their respective channels.

%WT The percent of time a device was kept waiting for I/O because a channel or storage control was busy.

In the sample report, no devices were restricted by waits for channels or storage controls.

MQ The maximum number of I/O operations that were queued on the device during the recording period. MXQ is a measure of device contention.

In the sample report, devices 4C0, 283, and 382 (all paging volumes) show nonzero values for MXQ, indicating contention for these devices.

Using CP INDICATE to Review I/O Workload

You can use the CP INDICATE command to view snapshots of I/O performance interactively. The values recorded by INDICATE are 8-minute averages of system performance, weighed toward the most recent sampling.

CP INDICATE I/O displays a list of virtual machines that are waiting for I/O and indicates the device that each is waiting for. Figure 17 shows a sample CP INDICATE I/O display.

```
TTWILLIA 05B4 ASCHENBR 09E5 NET      003A TL69     0A0F PVM      0045
DISKACCT 05B2
```

Figure 17. Sample CP INDICATE I/O Display for a VM/SP System. You must be authorized with CP privilege class E to issue this command.

You might use this command to see if a significant number of virtual machines are waiting for a specific device.

CP INDICATE USER *userid* displays an individual user's resource consumption, including the number of non-spooled I/O requests (SIO), and the number of reader, printer, and punch records spooled (RDR, PRT, PUN). Figure 18 illustrates a sample CP INDICATE USER display.

```
PAGES: RES-0063 WS-0055 READS=000561 WRITES=000501 PG-0065 PP-0000
VTIME=031:51 TTIME=052:03 SIO=001772 RDR-00178 PRT-00000 PCH-00132
SWAPS: SWAPOUT=000064 SWAPIN=000064 SW-0000
```

Figure 18. Sample CP INDICATE USER Display. You must be authorized with CP privilege class E to issue this command for users other than yourself. This example was taken from an HPO system; the display for a VM/SP or VM/XA system is slightly different.

You can write a simple EXEC to issue CP INDICATE commands at regular intervals and record the results in a console log. For more information on CP INDICATE, see the appropriate VM manual:

VM/SP	Up through Release 4	*VM/SP System Programmer's Guide*
VM/SP	Release 5 or later	*VM/SP CP for System Programming*
VM/SP HPO	Up through Release 4	*VM/SP HPO System Programmer's Guide*
VM/SP HPO	Release 5 or later	*VM/SP HPO CP for System Programming*
VM/XA SF	All releases	*VM/XA SF CP Command and Diagnosis Reference*

Chapter 4. Planning the Hardware Configuration

Depending on your needs, you can configure a set of devices in a variety of ways. Some configurations provide faster paging rates; some improve I/O rates; and others increase data availability. Whatever your goals, you want to make informed decisions about hardware configuration so that the results of your efforts are predictable and appropriate for your environment.

The factors to consider in designing an effective hardware configuration are:

- Hardware capability
- Capacity
- Performance
- Availability
- Shared DASD
- Guest systems

The purpose of this chapter is to discuss the various aspects of hardware configuration and to give you some techniques to help you achieve more effective configurations. To provide a detailed analysis of configuration possibilities is well beyond the scope of this book; consult the "Bibliography" on page 127 for a list of publications that provide more information on configuration.

For a description of valid attachment configurations of the 3380 and IBM storage controls and more detailed information on 3380 functions, see:

IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference

IBM 3380 Direct Access Storage Introduction

Understanding Hardware Capabilities

Before designing your storage subsystem, you need to understand how the various hardware functions might assist you in (or prevent you from) achieving your configuration objectives.

Because you must frequently sacrifice one aspect of your system, such as utilization, to improve another aspect, such as performance, the flexibility of your configuration is critical. A solid understanding of the following hardware capabilities will help you make sensible configuration trade-offs among capacity, performance, and availability:

- Internal Paths
- Dynamic Path Selection (DPS)
- Device Level Selection (DLS)
- Device Level Selection Enhanced (DLSE)

Internal Paths

Each 3380 A-unit model (except A04) contains two controllers, and each controller has four paths for accessing the devices on the string. Each successive 3380 model group provides improved access path capabilities between the controllers and the devices on the string.

Standard Model Internal Paths

The 3380 Model AA4 has four internal data paths attached to the two controllers by means of a 2x4 switch array. Since Model AA4 controllers have dedicated paths to specific devices, it is important to spread high-activity volumes across the various paths. At any instant, only one controller can use an internal path to access any specific device on that path. Simultaneously, the other controller can transfer data to or from any of the other 12 possible devices using another internal path. You can maximize performance on AA4 strings by placing high-activity volumes on different paths.

An AA4 string consisting of only four devices (a single 3380 AA4 unit) has only two internal data paths; path 0 with devices 0 and 1, and path 1 with devices 2 and 3. The devices sharing internal data paths in an AA4 string are shown in Figure 19.

INTERNAL	DEVICES															
PATH	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0	S	S							S	S						
1			S	S							S	S				
2					S	S							S	S		
3							S	S							S	S

Figure 19. 3380 AA4 Internal Data Path Sharing. In this table, an S below a device indicates that it is on the internal data path listed at the left side of the chart.

This internal path structure supports the DPS function. DPS controls which of the paths will be used by the controllers. See "Dynamic Path Selection (DPS)" on page 37 for further information.

Extended Capability Model Internal Paths

Extended Capability strings provide internal pathing from the controllers to the devices on the string that allows any device on the string to be selected by either controller on any of four internal paths. Thus, when one device is busy, all other devices in the 2-path string remain accessible to the other controller via any of three remaining internal paths. ***Any two*** devices in the string may be selected concurrently, one by each controller. The 2-path capability that is provided with these models provides Device Level Selection (DLS) capability; that is, any two devices may read or write data simultaneously. See "Device Level Selection (DLS)" on page 38 for additional information.

Note that when 2-path strings of Extended Capability models are intermixed with a 4-path Enhanced Subsystem string on a 3990 Model 2 Storage Control, the storage subsystem runs in DLSE support mode, and the 2-path strings have DLS capability.

Enhanced Subsystem Model Internal Paths

Enhanced Subsystem strings provide internal path capabilities that allow any controller on the string to access any device on any of four internal controller paths. The concurrent device access capabilities for Enhanced Subsystem models depend upon how the units are configured.

- When Enhanced Subsystem models are configured as **2-path strings,** the pathing capabilities are similar to those of Extended Capability models. There are two controllers that can access any of the devices (as many as 16) on any of the internal paths. ***Any two*** devices in the string may be selected concurrently, one by each controller. The resulting capability is DLS; see "Device Level Selection (DLS)" on page 38.
- When Enhanced Subsystem models are configured as **4-path strings,** there are four controllers that can access any device on the string (as many as 32) on any of the internal paths. ***Any four*** devices on the string can be accessed simultaneously, one by each controller. This capability is called DLSE; see "Device Level Selection Enhanced (DLSE)" on page 39.

The appropriate configuration feature must be used with the Enhanced Subsystem A-units as explained in *IBM 3380 Direct Access Storage Introduction.*

Dynamic Path Selection (DPS)

Note: VM/SP and VM/SP HPO do not support DPS functions, either in native mode or for guest operating systems. VM/XA SF** does **support DPS for both native and guest systems.

Dynamic path selection (DPS) is based on the DASD string having two controllers providing data transfer paths from the 3380 string to the storage directors, with selection of an alternate controller if one controller is unavailable. The storage directors, in turn, attach to two or more channels, thereby achieving multiple paths for transferring commands and data.

For system-initiated communication, a second I/O operation to the string can always be started, provided the two devices are on separate internal paths within the string. If the 3380 controller designated in the I/O address is busy or inoperative, the operating system or channel subsystem making the request is notified and can then select a path to the other controller to access 3380 data. The DPS function allows an alternate path to be established through another storage director and the other controller.

The DPS functions are:

- Simultaneous data transfer. Any two devices can transfer data simultaneously to or from the two controllers (one device to each controller), provided the devices are on separate internal paths within the string.
- Alternate controller selection. Each controller in a 3380 A-unit is capable of accessing any device in the 3380 string. With alternate path selection in a 2-path string, if one controller becomes unavailable or busy, another path is selected through another storage director to the other controller. No manual intervention is required. In a 4-path string, there are four controllers that can access any device in the string, providing even greater availability. An alternate controller can handle I/O operations for any device in the string (although only one device may transfer data at a time on a path).

- Volume reserve (System-Related Reserve). With system-related device reserve, (on 370/Extended Architecture hosts s only), a 3380 device can be reserved for use by a group of paths rather than a single path. Only the paths identified as part of the named group can access the device. System-related device reserve can be used to achieve higher levels of availability and I/O performance. It prevents simultaneous attempts to update critical data (for example, CMS directories).

  Notes for Guest Systems

  Because VM/SP and VM/SP HPO do not support DPS (except PMA guests with dedicated channels on a VM/SP HPO system), they will not accept commands associated with path group identification (such as SET PATH GROUP ID). If used, these commands will cause a unit check condition with command reject in the sense data. Guest operating systems under VM will interpret this to mean that dynamic path selection is not available.

  VM/XA SF (prior to Release 2.0) does not support DPS for guest systems but the guest system will benefit from the DPS performance advantage as the host CP system executes I/O using DPS. VM/XA SF (Release 2.0 and later) accepts commands associated with path group identification that are issued by an MVS/XA guest with dedicated DASD.

- Dynamic path reconnect. When data is not being transferred (such as during a seek operation), the paths to the 3380 device are disconnected. During this time the path is often used to access another device. When the original device requires the path again, it must be reconnected.

  On 370/Extended Architecture hosts, dynamic path reconnect allows the device to reconnect to the first available path within the group, rather than waiting for the use of the original path. This improves I/O performance for 3380 devices by reducing channel path access time, thereby reducing queueing time.

For more information on the functions of dynamic path selection, see *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference* or *IBM 3380 Direct Access Storage Introduction*, and the appropriate *Introduction* manual for your storage control:

3880	Model 2 or 3	*IBM 3880 Storage Control Models 1, 2, 3, and 4 Description Manual*
3880	Model 13	*Introduction to IBM 3880 Storage Control Model 13*
3880	Model 23	*IBM 3880 Storage Control Model 23 Introduction*
3990	Models 1, 2	*IBM 3990 Storage Control Introduction*

Device Level Selection (DLS)

Device level selection (DLS) is a capability of 3380 Models AD4, AE4, AJ4, and AK4 strings that provides two data transfer paths, one from each controller, to each device in the 2-path string (as many as 16 addresses). DLS extends the capabilities of DPS by allowing **any two** devices in the 2-path string to read or write data simultaneously.

When attaching 2-path AJ4 or AK4 strings to a 3990 Storage Control, DLS support mode must be set by the service representative at hardware installation time. For more information on 3990 DLS support mode, see *IBM 3990 Storage Control Introduction*. The Enhanced Subsystem A-units must have the appropriate feature as explained in *IBM 3380 Direct Access Storage Introduction*.

DLS uses two storage directors, in either 3880 or 3990 Storage Controls, that provide two data paths to each device in a 2-path string. When attached to the 3990, each of the storage directors operates as a single-path storage director.

3380 models with DLS offer improved data availability and overall system performance when compared to Standard models. If the selected device is not busy, it may be accessed even if another device on the 2-path string is reading or writing data. When one device on the 2-path string is busy, any of the remaining devices can be selected. This can reduce the amount of time an operating system needs to wait for a path to a device to become available.

Device Level Selection Enhanced (DLSE)

Device level selection enhanced (DLSE) is a capability of 3380 Models AJ4 and AK4 (requiring two interconnected A-units) that provides four data transfer paths, one from each controller, to each device in a 4-path string (as many as 32 addresses). DLSE extends the capabilities of both DPS and DLS by allowing **any four** devices in a 4-path string to read or write data simultaneously.

A 4-path string attaches only to a 3990 Storage Control Model 2, with the storage directors operating as multipath storage directors to provide four data transfer paths. The 3990 operates in DLSE support mode, set by the service representative at hardware installation time. For more information on 3990 DLSE support mode, see *IBM 3990 Storage Control Introduction*. For information on the feature required on the Enhanced Subsystem A-units, see *IBM 3380 Direct Access Storage Introduction*.

3380 models with DLSE offer improved data availability and overall system performance when compared to both Standard and Extended Capability models. If the selected device is not busy, it may be accessed even if any three other devices on the 4-path string are reading or writing data. End-user response can be more consistent during heavy workload periods with DLSE.

With DLSE, any one of the four paths can be quiesced (set offline to the host) without disrupting availability of the devices on the other paths. This allows additional B-units to be added to the 4-path string without disrupting availability of the existing devices, or addition of another 4-path string to the 3990 storage control without disrupting use of an established string. For additional information on the nondisruptive DASD installation capability associated with DLSE, see *IBM 3380 Direct Access Storage Introduction*.

Speed Matching Buffer for 3380

The Speed Matching Buffer for 3380 (SMB) is a feature that can be installed on the storage directors of a 3880 Model 2 or 3 Storage Control. With SMB, you can attach 3380 Model A04 and AA4 strings to processor channels that have data transfer rates of less than 3.0 MB per second. This allows you to use 3380 devices with processors that do not support data streaming.

Since an AA4 string attaches to two storage directors (preferably in different 3880s), ***both*** storage directors to which the AA4 string is attached must have SMB installed. Furthermore, if one storage director in a 3880 Model 3 has SMB installed, the other storage director in that 3880 must have SMB installed.

Note that you cannot attach 3380 Model AD4, AE4, AJ4, or AK4 strings to 3880 units that have the SMB feature. If you have the SMB feature installed on your 3880, you must remove it from both storage directors before installing either the 3380 AD4/AE4 Support feature or the 3380 AJ4/AK4 Support feature on the 3880.

In developing your hardware configuration, plan to isolate 3380 A04/AA4 strings on their own 3880 storage controls if they require the SMB feature.

For more information on SMB, see *IBM 3880 Storage Control Models 1, 2, 3, and 4 Description Manual*. For more information on data streaming, see *IBM 3380 Direct Access Storage Introduction*.

Channel Switches on the Storage Control

Channel switches allow the storage directors of a storage control to share a set of channels. With channel switching, you can establish a number of separate paths to each 3380 string that is attached to a storage control with the appropriate channel switch feature. The channels may be on the same or different processors.

For more information on channel switching and its effects on hardware configuration, see the appropriate manual for your storage control:

3880	Model 2 or 3	*IBM 3880 Storage Control Models 1, 2, 3, and 4 Description Manual*
3880	Model 13	*Introduction to IBM 3880 Storage Control Model 13*
3880	Model 23	*IBM 3880 Storage Control Model 23 Introduction*
3990	Models 1 or 2	*IBM 3990 Storage Control Introduction*
3380	Model CJ2 Direct Channel Attach	*IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference*

Cache Storage

The 3880 Models 13 and 23 contain an area of electronic storage called **cache**. Cache acts as an intelligent, high-speed buffer to DASD, retaining frequently-read data for faster access by the processor.

Having data available in the cache can greatly improve I/O response time. Access time between the cache and the channel is much shorter than that between the DASD and the channel; cache accesses eliminate the seek and rotational time inherent in DASD accesses.

For configuration purposes, 3880 cache storage controls can be attached to as many as two strings of 3380 devices, with as many as 16 devices on each string.

Note: VM/SP ***does not support*** cache storage controls. VM/SP HPO and VM/XA SF support 3880 Model 13 and 23 storage controls as shown in Figure 4 on page 9 and Figure 5 on page 10.

You can read about the types of data suitable for cache storage in "Identifying Candidates for Cache Storage" on page 57. For more information on the use of cache storage, see the appropriate *Introduction* manual for your storage control:

3880	Model 13	*Introduction to IBM 3880 Storage Control Model 13*
3880	Model 23	*IBM 3880 Storage Control Model 23 Introduction*

Configuring for Capacity

Although an elaborate capacity planning study is not required to use 3380s, you should have a high-level understanding of space utilization in your current hardware configuration before installing the 3380. By understanding your current space needs, you can plan for orderly growth and reduce those performance problems that accompany insufficient space.

Figure 20 summarizes the data capacity of the 3350, 3375, and the various 3380 models in terms of blocks and approximate megabytes (MB).

DASD Model	Capacity per	512-byte Blocks	800-byte Blocks	1024-byte Blocks	2048-byte Blocks	4096-byte Blocks
3350	Volume	449 500 blks 230 MB	316 350 blks 253 MB	249 750 blks 256 MB	133 200 blks 273 MB	66 600 blks 273 MB
	2-volume unit	899 100 blks 460 MB	632 700 blks 506 MB	499 500 blks 512 MB	266 400 blks 546 MB	133 200 blks 546 MB
3370 A1, B1, A11, B11	Volume	558 000 blks 285 MB	558 000 blks 223 MB	279 000 blks 285 MB	139 500 blks 285 MB	69 750 blks 285 MB
	2-volume unit	1 116 000 blks 571 MB	1 116 000 blks 446 MB	558 000 blks 571 MB	279 000 blks 571 MB	139 500 blks 571 MB
3375	Volume	460 320 blks 325 MB	345 240 blks 276 MB	287 700 blks 294 MB	161 112 blks 330 MB	92 064 blks 377 MB
	2-volume unit	920 640 blks 470 MB	690 480 blks 552 MB	575 400 blks 589 MB	322 224 blks 660 MB	184 128 blks 754 MB
3380 A04, AA4, B04, AD4, BD4, AJ4, BJ4	Volume	610 650 blks 312 MB	477 900 blks 382 MB	411 525 blks 421 MB	238 950 blks 489 MB	132 750 blks 543 MB
	4-volume unit	2 442 600 blks 1248 MB	1 911 600 blks 1528 MB	1 646 100 blks 1684 MB	955 800 blks 1956 MB	531 000 blks 2172 MB
3380 AE4, BE4	Volume	1 221 300 blks 625 MB	955 800 blks 764 MB	823 050 blks 842 MB	477 900 blks 978 MB	265 500 blks 1087 MB
	4-volume unit	4 885 200 blks 2500 MB	3 823 200 blks 3056 MB	3 292 200 blks 3368 MB	1 911 600 blks 3912 MB	1 062 000 blks 4348 MB
3380 AK4, BK4	Volume	1 831 950 blks 936 MB	1 433 700 blks 1146 MB	1 234 575 blks 1264 MB	716 850 blks 1468 MB	398 250 blks 1631 MB
	4-volume unit	7 327 800 blks 3744 MB	5 734 800 blks 4584 MB	4 938 300 blks 5056 MB	2 867 400 blks 5872 MB	1 593 000 blks 6524 MB
3380 CJ2	Volume	610 650 blks 312 MB	477 900 blks 382 MB	411 525 blks 421 MB	238 950 blks 489 MB	132 750 blks 543 MB
	2-volume unit	1 221 300 blks 625 MB	955 800 blks 764 MB	823 050 blks 842 MB	477 900 blks 978 MB	265 500 blks 1087 MB

Figure 20. Summary of VM Data Capacity for DASD Models

Use the information you gathered in Chapter 3, "Understanding Your Data and Hardware Configuration" to determine the number of 3380 volumes you will need for each group of data. If you are replacing 3350 devices with 3380s and are using a 4096 block size, you can use the following capacity guidelines:

- Slightly less than two full 3350 volumes can be placed on one volume of a 3380 Model A04, AA4, B04, AD4, BD4, AJ4, BJ4, or CJ2.
- Slightly less than four full 3350 volumes can be placed on one volume of a 3380 Model AE4 or BE4.
- Slightly less than six full 3350 volumes can be placed on one volume of a 3380 Model AK4 or BK4.

For planning purposes, you should assume that the 3380 volumes will be filled to roughly 70% to 80% capacity, allowing space for expansion and data movement.

Include spare volumes in your capacity plans. If a device failure occurs, recover the data onto the spare volume.

More information on block size selection can be found on "Selecting a Block Size for CMS Minidisks" on page 76.

Configuring for Performance

In Chapter 3, "Understanding Your Data and Hardware Configuration" you identified groups of data with specific performance requirements based on the nature of the data or the frequency of use. The next step is to design a configuration to favor the volumes containing these groups of data.

To meet your users' requirements for response time and throughput, you must pay close attention to the hardware configuration factors that affect performance. These factors are summarized below.

Capability of processor, channels, and storage controls

In many systems, performance problems can be traced to an imbalance in the capabilities of the various components of the storage subsystem. All components should be sharing the load of I/O appropriately. DASD volumes and paths should be adequate to keep the processor busy during peak load conditions; and no channel, storage control, or device should be so busy that it causes excessive delays. In a balanced subsystem, no single component is the limiting bottleneck.

Number of volumes on a string

To reduce contention for I/O paths, you should limit the number of high-activity volumes on a given string. Most systems do not have the luxury of a separate path for each volume, and availability requirements may imply that volumes should be accessible over many different paths. The DPS (in VM/XA), DLS, and DLSE hardware features can help to improve performance when a number of volumes share access to paths. DPS, DLS, and DLSE are described in "Understanding Hardware Capabilities" on page 35.

For high-activity volumes, the length of the string (that is, the number of volumes attached to the 3380 A-unit string controller) can affect performance. Shorter strings reduce the possibility that all paths will be in use by other volumes when the high-activity volume needs access.

I/O access skew

For a given storage subsystem, the total load is seldom spread uniformly across paths and devices. It is not uncommon to find that 40% to 50% of the devices carry 90% or more of the total I/O load. This nonuniform spread of I/O activity is called I/O access skew.

You should be concerned primarily about the devices in your current configuration toward which I/O activity is most heavily skewed. (In VM these are frequently paging, swapping, and spooling volumes.) These devices are doing more work than other devices in your current configuration. The I/O activity on these devices must be reviewed and balanced in order to maintain system performance.

Configuring for Availability

Availability, for a storage subsystem, is defined as the degree to which a resource is ready when needed to process data. Consequently, availability strongly influences system performance.

Sometimes availability problems in the storage subsystem are easy to spot—for example, when a channel path fails unexpectedly and alternate paths must be used. But you might also have pervasive availability problems without realizing their extent. Availability problems can show up in ways, such as poor response time, that might initially look like performance problems instead.

Lack of availability can severely impact the users of your system. It's important to consider how your hardware configuration affects both system and data availability.

Some techniques you can use to configure for availability are summarized below.

- Establish multiple access paths to 3380 volumes. When a path from the 3380 A-unit string controller to a volume is busy, data should be accessible via another internal path. DPS (in VM/XA), DLS, and DLSE offer significant availability improvements for paths between controllers and devices. (See "Understanding Hardware Capabilities" on page 35 for more information on DPS, DLS, and DLSE.)

  With the 3990 storage control, you can configure 3380 enhanced subsystem models in 4-path strings to improve availability. *IBM 3380 Direct Access Storage Introduction* describes how to set up 4-path configurations.

- Establish alternate paths through alternate storage controls. If a storage control fails, data should be accessible via another storage control. In addition, the 3990 Storage Control 2 improves availability by providing two separate storage clusters, each with two independently-addressable storage directors. Storage controls can be configured as dual-frames. For more information, see the manual describing your storage control.

- Establish alternate paths through alternate processor channels, using channel switches on the storage controls. If a channel fails, data should be accessible via another channel. (See "Channel Switches on the Storage Control" on page 40 for more information on switching.) "Alternate Paths in VM/SP and VM/SP HPO" on page 44 discusses how VM handles alternate paths.

- In a multiple system environment, plan for a backup processor through channel switches, and establish each processor's rights to DASD access during system definition. The devices should be defined to both processors during system definition, then varied offline to the backup processor and online to the primary processor. If the primary processor fails, the alternate processor can then be used as a backup until repairs are made.
- Keep spare volumes available to be used in case of hardware failure. Many systems use their spare volumes to store temporary data until needed for hardware recovery purposes.

The manual *Component Failure Impact Analysis—An Availability Management Technique* describes how to plan and configure for hardware availability. For detailed examples of hardware configurations that illustrate some of these availability techniques, see *IBM 3380 Direct Access Storage Introduction*.

Alternate Paths in VM/SP and VM/SP HPO

The use of alternate paths in VM/SP and VM/SP HPO is different from that of other operating systems, so there are some special considerations for pathing in VM.

On a uniprocessor system, VM attempts to send I/O down the primary path first. If the device, control unit, and channel are busy (queued) or offline, VM tries to send the I/O down the alternate path, if an alternate control unit has been specified for the device. This means that unless your I/O subsystem is very heavily loaded, the alternate paths are rarely used.

On a multiprocessor system, VM assumes that all processors have the same I/O configuration and allows only one alternate path to be specified for each processor. I/O from one processor is sent down the primary path in that processor's channel set; if that path is busy, the processor will attempt to send that I/O request down its alternate path. If both the primary and the alternate path are queued or offline, the second processor will attempt to send the I/O request down *its* primary path.

Because of the way VM handles alternate paths, you have an additional configuration option that might be appropriate for your data processing center. For example, you might:

- Establish a string of 16 volumes connected to processor channels 1 and 2.
- Code the DMKRIO input file so that channel 1 is the primary path to two volumes, and channel 2 is the primary path to the other 14 volumes.

With this configuration, the two volumes on channel 1 could be driven at very high I/O rates without causing channel contention for the other 14 volumes. See Figure 32 on page 65 for an example of how to code the RDEVICE, RCNTLUNIT, AND RCHANNEL macros to accomplish this.

For more information on alternate path use in VM/SP and VM/SP HPO, see the technical bulletin, *Alternate Pathing under VM*.

Alternate Paths in VM/XA

In VM/XA, the operating system does not identify or select alternate I/O paths. Pathing is handled by the microcode of the channel subsystem. This microcode has been designed to use an efficient path allocation algorithm, so that the special considerations for VM/SP and VM/HPO do not apply. For more details, see the appropriate IOCP manual for your processor.

Shared DASD Considerations

Sharing DASD allows multiple systems to access data stored on common volumes. Although sharing improves data availability, it may result in device contention and performance problems.

In general, when sharing minidisks across systems, use VM/ISF. Only one person should have write-access to a given minidisk. More information on VM/ISF is found in the manuals listed in the bibliography.

Performance monitoring must be handled separately for each system that uses the devices.

For a detailed discussion of the considerations in sharing DASD under VM, see the technical bulletin, *DASD Sharing under VM*.

Guest System Considerations

Guest operating systems under VM have special configuration needs and capabilities. For more information, see *VM Running Guest Operating Systems*.

Consider the following general recommendations for hardware configuration in guest operating system environments:

- Dedicate devices and channels to guest systems. Because their heavy I/O activity will degrade performance for interactive users, MVS volumes should always be isolated from those managed by VM. Do **not** mix VM system volumes with guest volumes. See *Using the IBM 3380 Direct Access Storage in an MVS Environment* or *Using the IBM 3380 Direct Access Storage in a VSE Environment* for information on configuring strings of dedicated devices and channels on those guests.
- Remember that MVS guest systems use device reserve/release and string switching functions that are not otherwise available to VM. Configuration possibilities will vary accordingly.

Building the Hardware Configuration

Now that you've reviewed the design considerations for your system service objectives, you're ready to explore some techniques for building more effective hardware configurations.

There are many types of configurations possible, each with its own characteristics. Some typical configurations for achieving performance or availability goals are discussed in this section; more complex types of configurations are illustrated in *IBM 3380 Direct Access Storage Introduction*. The configurations shown here

assume that the 3380 A-units have two string controllers (Models AA4, AD4, AE4, AJ4, and AK4).

By combining or substituting the configurations that meet the specific needs of your system, you can plan your total system hardware configuration. These figures explain how to attach strings, performance will vary by individual environment.

For information on the features listed in the examples, see *IBM 3380 Direct Access Storage Introduction*.

A Basic 3380-3880 Configuration

A basic 3380-3880 configuration is shown in Figure 21.

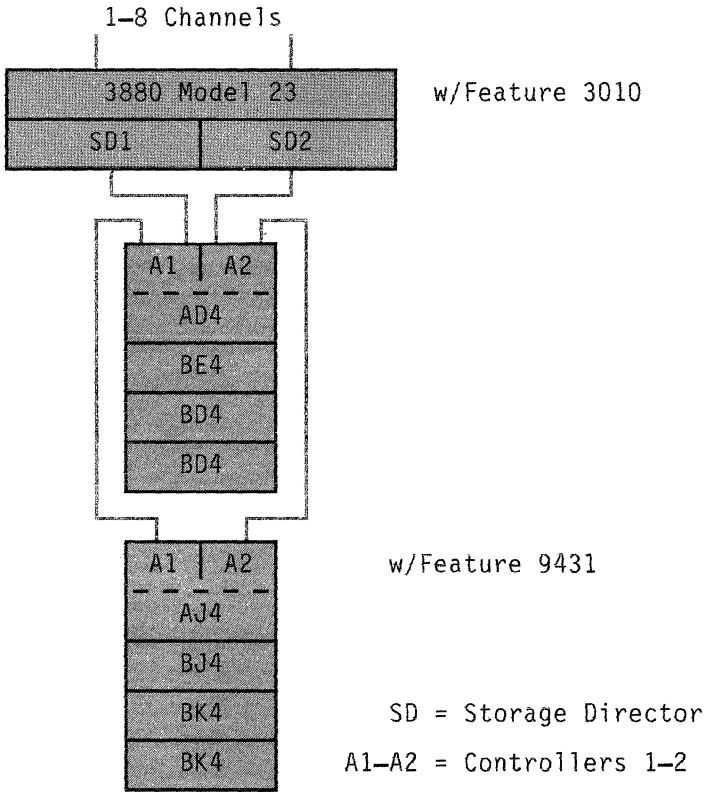

Figure 21. 3380 Model AD4 and AK4 Strings Attached to 3880 Model 23

In this example, two 3380 strings are sequentially connected to both storage directors of a 3880 Model 23. One string contains a 3380 Model AD4 controller followed by a mixture of BD4 and BE4 units; the other string contains a 3380 Model AJ4 controller with BJ4 and BK4 units.

This example is also valid for 3880 Model 3 Storage Controls which would require Feature 3005 in this configuration.

A 3380-3880 Short String Configuration

An example of attaching a short 3380 string to two 3880 storage controls to improve performance and availability is shown in Figure 22.

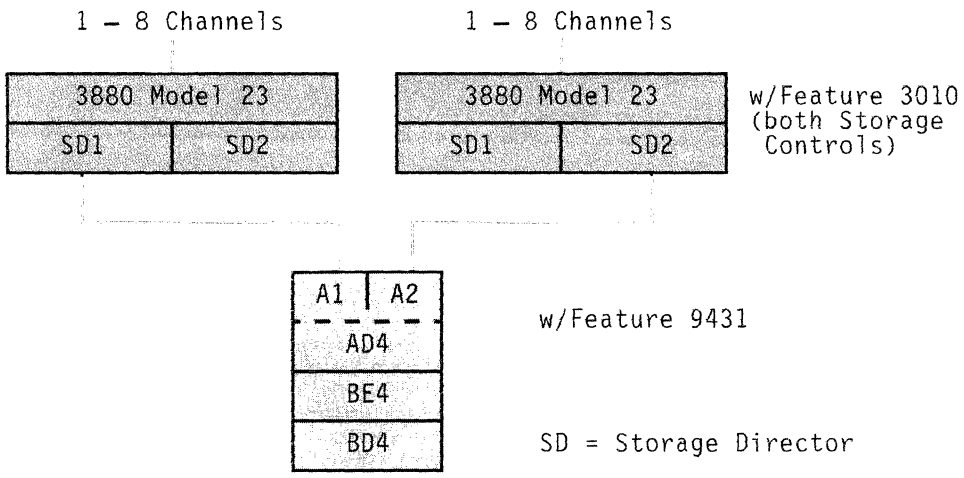

Figure 22. Example of a 3380 Model AD4 String Attached to Two 3880s

In this example, a 3380 string is connected to two storage directors in different 3880 Model 23 storage controls. This configuration improves data availability by allowing access to the 3380 string through either storage control.

The 3380 AD4 string shown here is an example of a short string used to improve performance. A configuration of this type might be appropriate when the performance and availability requirements of the data are critical. For example, minidisks that show high access rates or device utilization might reasonably be contained on these volumes.

This example is also valid for 3880 Model 3 Storage Controls which would require Feature 3005 in this configuration.

Intermixed Strings on Single-path Storage Directors

There are two rules for intermixing strings of different model groups on the same single-path storage director:

- One Model AA4[3] string can be attached with one Model AD4 or AE4 string.
- One Model AA4, AD4, or AE4 string can be attached with one Model AJ4 or AK4 2-path string.

Figure 23 is an example of four 3380 2-path strings, with two sequentially connected to each of the paired storage paths in the same 3990 Model 2.

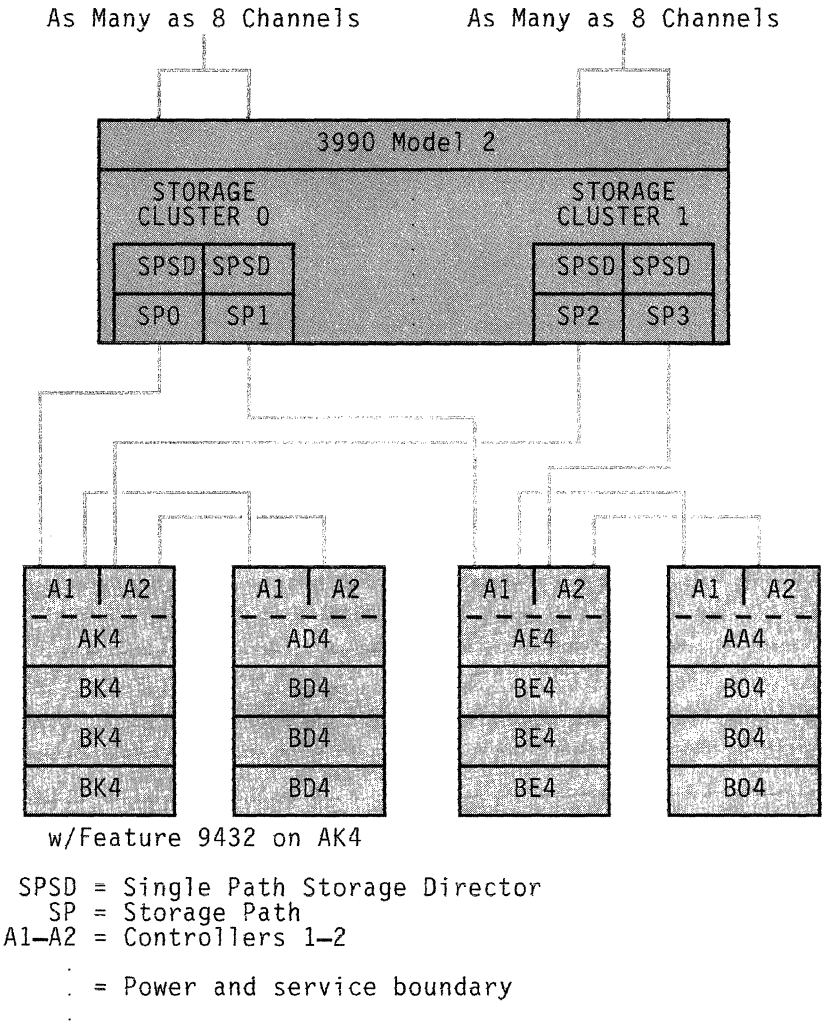

Figure 23. Intermixed 3380 2-Path Strings Attached to a 3990 Model 2

In this example, there are four 3380 strings, two sequentially connected to each of the paired storage paths in the same 3990 Model 2. One string contains standard models, two strings contain extended capability models, and the other string contains enhanced subsystem models.

---

[3] Attachment to a 3990 requires that the Model AA4 have a serial number of 15000 or greater for 60 Hz units or X0300 or greater for 50 Hz units.

A Mixed 2-Path and 4-Path 3380-3990 Configuration

An example of attaching 3380 strings to a 3990 storage control in both 2-path and 4-path configuration is shown in Figure 24.

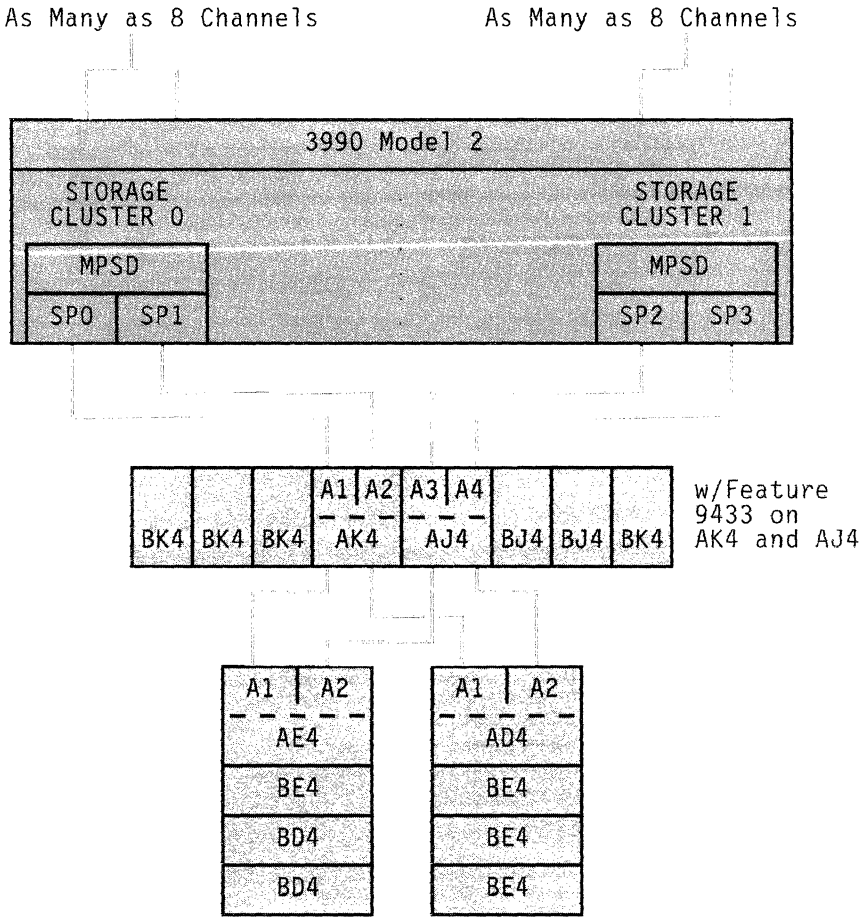

Figure 24. 3380 4-Path and 2-Path Strings Intermixed on a 3990 Model 2

In this example, two 3380 2-path strings and one 3380 4-path string are sequentially connected to two paired storage paths in a 3990 Model 2. In order to attach 3380 strings to 3990 storage controls in 4-path configuration, the two A-units in the 4-path string must be logically and physically connected to each other.

A 3380-3990 4-Path Configuration for High Availability

Figure 25 is an example of two 3380 4-path strings sequentially connected to the paired storage paths in the same 3990 Model 2. This configuration provides high availability for the data on these volumes.

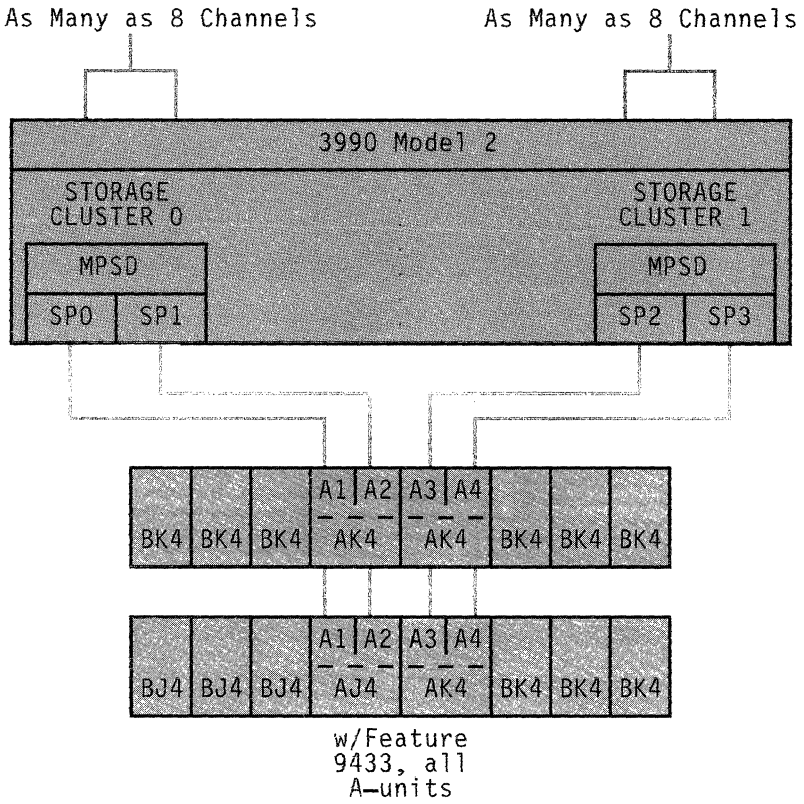

Figure 25. 3380 Model AJ4 and AK4 4-Path Strings Attached to a 3990 Model 2

In order to attach 3380 strings to 3990 storage controls in 4-path configuration, the two A-units in each 4-path string must be logically and physically connected to each other.

A 3380 CJ2 Configuration

An example of attaching a 3380 Model CJ2 string directly to a processor channel is shown in Figure 26.

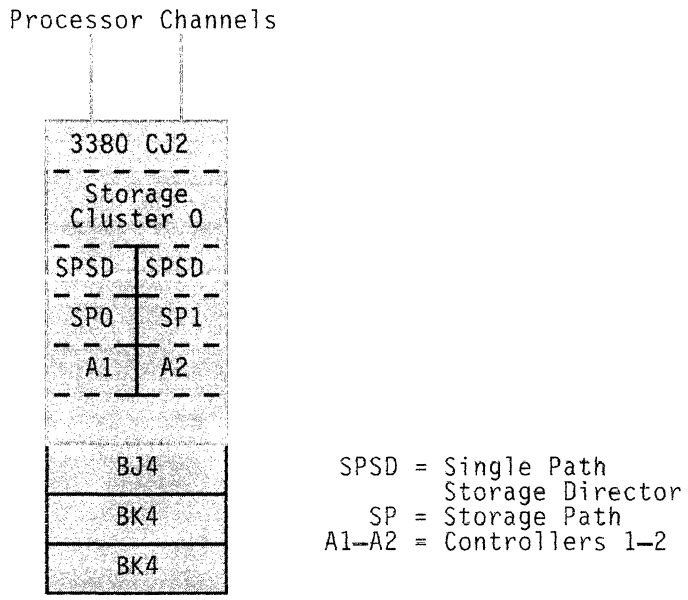

Figure 26. 3380 Model CJ2 String Attached Directly to Processor Channel

In this example, the 3380 Model CJ2 string contains a mixture of J- and K-units. The Model CJ2 contains both a string controller and a storage control, so it can be attached directly to a processor channel.

Documenting the Hardware Configuration

After determining the configuration techniques, document the planned hardware configuration, including the rationale for choosing various configurations. A written hardware configuration plan should include:

- Diagrams of each new or changed DASD string, including storage control and processor attachments.
- A summary of capacity (number of devices/volumes) for each string.
- A description of the elements of the configuration that ensure hardware availability for each string: alternate paths, controllers, storage directors, storage controls, and channels.
- Diagrams showing location and attachments for shared DASD.
- Key contacts for hardware configuration.

Review the plan with all involved system support, operations, and user groups to ensure its viability.

Chapter 5. Planning the Data Configuration

From the analysis of your data configuration that you completed in Chapter 3, you should be able to identify the performance, availability, and capacity requirements of your data. The next step is to translate these requirements into specific volumes and strings of volumes, and to determine exactly where the data should be placed to achieve your service objectives.

This chapter discusses:

- Mapping data to volumes in the hardware configuration
- Planning to move data to new volumes
- Planning for backup and recovery during data movement
- Documenting the data configuration

Mapping Data to Volumes

Based on the hardware configuration you designed in Chapter 4, you should be able to identify appropriate volumes for user minidisks, system data, paging, and temporary space.

For the most part, performance is affected more by *which* volume the data resides on than by *where* on the volume. However, you should still plan to place data with high performance requirements appropriately.

Placing Performance-Critical Data

In VM, there are some minidisks and system areas that have a significant impact on system performance. System minidisks, CP system data, and large data bases may be frequently accessed; these kinds of I/O operations can severely degrade end user response time if the data is not placed effectively.

Before you establish the locations for critical data in your new configuration, check the effectiveness of the previous location of the data. Whenever possible, retain the performance advantages of previously effective tuning efforts.

Some guidelines to follow in placing performance-critical data are listed below.

- Use cache storage where appropriate. See "Identifying Candidates for Cache Storage" on page 57 for more information.
- Place concurrently active minidisks on different paths to avoid contention for the same channel, storage control, and storage director resources. For example, if non-cached, the CMS S- and Y-disks should not be on the same path as the CP directory or CP paging or spooling areas.
- If different paths are not feasible, then try to place concurrently active minidisks on different volumes. If the minidisks must be on the same volume, place them on adjacent cylinders to minimize arm movement.
- Place performance-critical data in the center of the volume, less-active data near the beginning or end of the volume.
- Share minidisks across systems only where necessary.

- Isolate page and swap areas on their own volumes. For details on how to allocate paging, swapping, and spooling space, see the appropriate VM *Planning* manual:

VM/SP	All releases	*VM/SP Planning Guide and Reference*
VM/SP HPO	All releases	*VM/SP HPO Planning Guide and Reference*
VM/XA SF	All releases	*VM/XA SF Virtual Machine Planning*

Many of these performance factors involve balances and compromises that must be evaluated for specific applications; consider the needs of your own users before placing critical data.

When there is more performance-critical data than will fit on the designated volumes, consider the windows within which the various minidisks are used. Combine several types of performance-critical data on the same volume or path if the times during which they are used do not overlap significantly. An example of this is combining batch data with online data, if the online data is used during the prime shift and the batch data is used after the prime shift.

Periodically, reexamine the placement of performance-critical data, due to the changing nature of user and application requirements.

Placing the S- and Y-Disks: A Scenario

You may remember that the sample VMMAP report from Figure 12 on page 27 indicated a problem in minidisk placement on volume VMSY81. Specifically, minidisks 190 (the S-disk) and 19E (the Y-disk) were responsible for a total of 84% of all seeks on the system.

This is an example where I/O service can be improved by appropriate minidisk placement.

Figure 27 shows the original location of minidisks 190 and 19E on volume VMSY81.

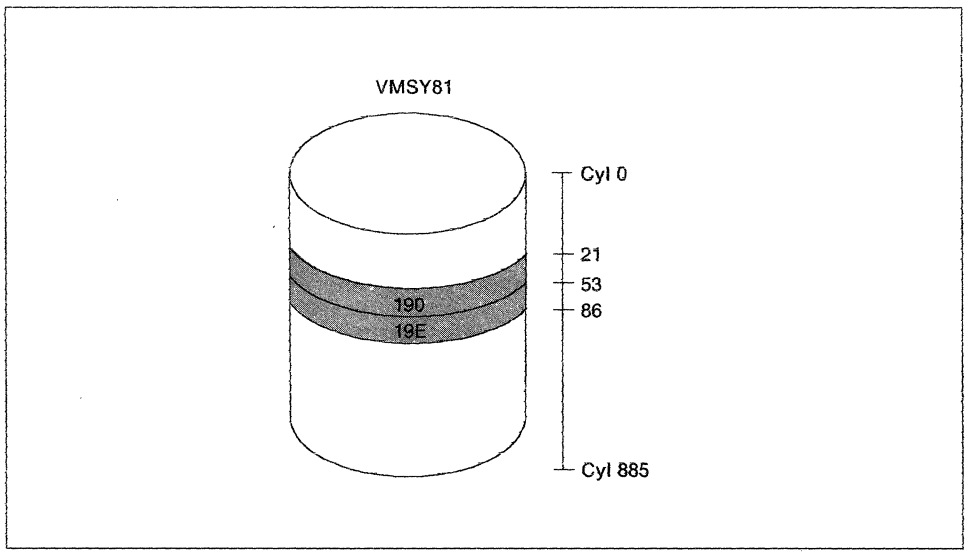

Figure 27. Position of Minidisks 190 and 19E: Before Repositioning

Notice that even though the minidisks are physically adjacent on the volume, there is still an excessive amount of arm movement going on.

There are a number of things that might be done to improve performance for this volume and for these minidisks.

Option 1 is to place the S- and Y-disks behind storage controls with cache to improve performance. "Identifying Candidates for Cache Storage" on page 57 discusses the kinds of data that are suitable for cache.

Option 2 is to split minidisks 190 and 19E onto separate volumes, preferably on separate channels, to reduce device contention. Figure 28 illustrates the position of the minidisks under this option. Each minidisk is placed in the center of the volume, to reduce access arm movement.

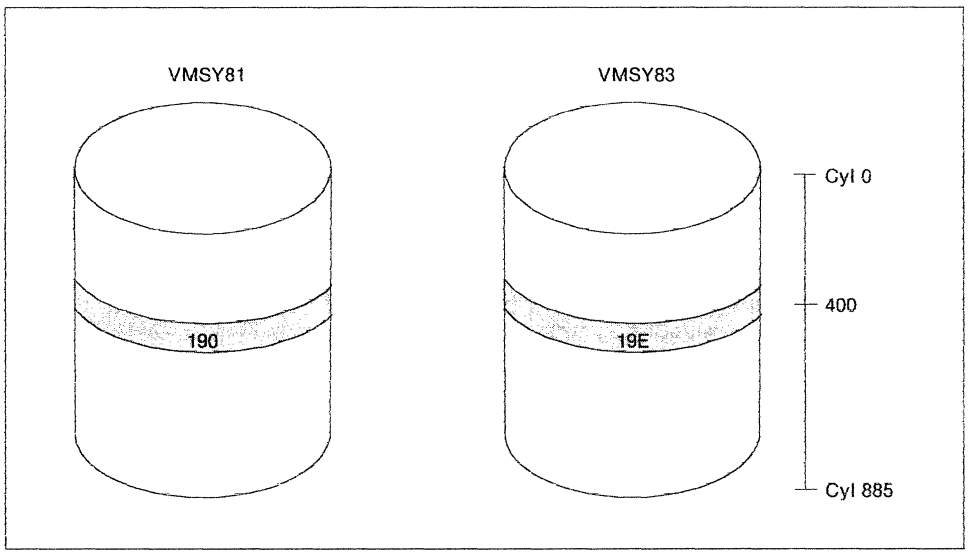

Figure 28. Position of Minidisks 190 and 19E: Option 2

You can use the DDR program to move minidisks among volumes. See Figure 54 on page 88 for sample DDR statements to move minidisks.

If you move the S- and Y-disks, you must update the NAMESYS macro (VM/SP and VM/HPO only) in DMKSNT with the new locations of the minidisks and resave the system. If you use DDR to move the disks, you do not need to resave the system. You must also update the CP directory to reflect the new minidisk addresses. For more information on moving system minidisks, see the appropriate VM *Planning* manual:

VM/SP	All releases	*VM/SP Planning Guide and Reference*
VM/SP HPO	All releases	*VM/SP HPO Planning Guide and Reference*
VM/XA SF	All releases	*VM/XA SF Virtual Machine Planning*

Option 3 is to maintain duplicate S- and Y-disks on a separate string to reduce contention. If contention is a significant problem, you should consider maintaining multiple sets of S- and Y-disks.

Figure 29 shows the S- and Y-disks copied onto a second volume. Again, you should place each minidisk in the center of the volume, to reduce access arm movement.

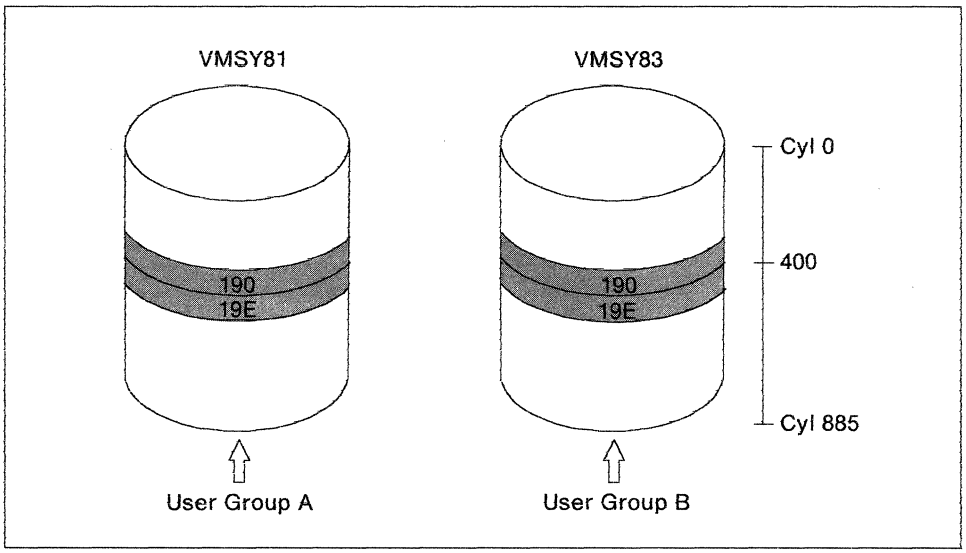

Figure 29. Duplicate Minidisks 190 and 19E: Option 3

You can use the DDR program to copy minidisks among 3380 volumes. See Figure 54 on page 88 for sample DDR statements to copy minidisks. If you are copying from unlike devices to 3380 volumes, you must use the COPYFILE command of CMS. See Figure 55 on page 89 for sample COPYFILE statements used to copy minidisks.

If you use duplicate S- and Y-disks, you must update the NAMESYS macro in DMKSNT with VSYSADR=IGNORE to prevent CP from checking S-disk locations at linkage. If you code VSYSADR=IGNORE, you must also specify VSYSRES and either SYSCYL or SYSBLOK as null (for example, VSYSRES=, and SYSCYL=, or SYSBLOK=,). You must also update the LINK statements for each group of users to point to the proper set of S- and Y-disks. More information on coding the NAMESYS macro can be found in the sections describing how to save/define a saved system in the appropriate VM *Planning* manual:

VM/SP	All releases	*VM/SP Planning Guide and Reference*
VM/SP HPO	All releases	*VM/SP HPO Planning Guide and Reference*
VM/XA SF	All releases	*VM/XA SF Virtual Machine Planning*

Identifying Candidates for Cache Storage

For highly-active volumes, cache storage is be a good way to improve performance and minimize I/O service and response times. Although you should do a complete analysis, some volumes are typically reasonable candidates for attachment to a cache storage control.

In VM, cache is best used for data that is frequently-read but rarely written (such as read-only minidisks). This is primarily because of the way that CMS handles I/O. When a CMS file is updated, the old version of the file stored on DASD is *not* overwritten, rather the data is written to new blocks. After the data has been written, pointers are changed to reflect the location of the newly written data. This is done to ensure that a valid file will exist even if the system or the updating program fails. Because of this, however, ***all CMS I/O writes are considered cache write "misses"*** (that is, the contents of the cache do not match the contents of the record updated by the processor). Therefore, data that is frequently written will not show performance improvements when placed in cache storage.

The following types of data under VM are most appropriate for cache:

- CMS S- and Y-minidisks
- Minidisks containing frequently-read licensed programs (such as PROFS or language compilers)
- Frequently-read (but not frequently-written) system areas (such as the CP directory or saved system area)

Note for Guest Systems

For volumes dedicated to guest systems, see *Using the IBM 3380 Direct Access Storage in an MVS Environment*, where you can find information on the types of data most suitable for cache.

You should also consider cache for volumes that show large average seek lengths for data that is heavily used (illustrated, for example, in the SEEKS ANALYSIS: AVG LENGTH NON-0 field of Figure 11 on page 25). A large seek length indicates that mechanical motion constitutes a significant portion of I/O response time—a portion which can be effectively eliminated when using cache.

For more information on using cache, see the appropriate *Introduction* manual for your cache storage control:

3880	Model 13	*Introduction to IBM 3880 Storage Control Model 13*
3880	Model 23	*IBM 3880 Storage Control Model 23 Introduction*

Planning to Move Data

After you have determined where in the configuration your data should reside, the next step is to plan how and when data will be moved.

Because you undoubtedly have a limited amount of time to move a large amount of data, you will probably not be able to complete the 3380 migration in a single session. You will therefore need to divide your data into groups, so that each group can be managed individually and moved in a single session.

First, plan to move the data that is easiest to move and least critical to your environment. Progress to the more difficult and critical data as you gain migration experience.

When grouping the data, take into account the kind of data you have to move, scheduling constraints, and the potential impact of the group's data movement on other data groups and users. The following sections describe some methods of grouping data for migration purposes.

Moving by Volume

Moving entire volumes at a time can be a good method for most types of user data. Naturally, moving entire volumes leaves the placement of data on the volume intact; so full volume migration should be done only when the volume contains data that requires no reorganization or special placement for performance reasons.

You can combine low-activity volumes onto higher capacity devices if the I/O peaks are relatively low and total I/O activity is uniform. For high-activity volumes, you may want to consider moving each one to a higher capacity volume or to a faster device to improve performance.

For details on how to move volumes under VM, see "Moving Full Volumes From 3380 to 3380" on page 86.

Moving by Minidisk

Moving minidisks is necessary to attain the specific placement required by performance-oriented data, or to redistribute data among volumes. It is more time-consuming to move minidisks on an individual basis, because each case requires individual attention. However, it also ensures that minidisks are placed exactly where you want them.

For details on how to move minidisks under VM, see "Moving Minidisks From 3380 to 3380" on page 88.

Planning for Backup and Recovery

Installing new devices will affect large amounts of data, much of which might be critical to the smooth operation of your data processing center. To minimize the disruption caused by the migration, it's important to plan and implement a strategy for backup and recovery during the migration period.

Your data processing center probably already has data backup and recovery procedures in place. Evaluate these procedures before you move data and extend them to protect data during the device migration period, when it is especially vulnerable.

As a minimum, plan for the following:

- Back up each old volume before its contents are moved.
- Back up each new volume as soon as its contents are complete.
- Keep hardcopy records that will enable you to quickly tell where any file on any moved volume has been located, and where its original backup (from the old volume) and its latest backup are located.
- If possible, keep the backup DASD volumes on the system while data is being moved, to ensure quick data recovery if necessary.

For examples of backing up old and new volumes, see "Backup and Recovery during Migration" on page 84. More information on backup and recovery can be found in *VMBACKUP Management System General Information*.

Documenting the Data Configuration

After deciding on the data movement strategy, document the planned data configuration, including the strategy for data movement and the related backup/recovery procedures. A written data configuration plan should include:

- A description of the groups of data you have selected to move.
- A description of where various groups of data will be placed. This should be in terms of volumes and strings. Where necessary for performance reasons, you should also describe where the specific data areas and minidisks are placed on a volume.
- The plan by which the groups of data will be moved.
- A description of the supplemental backup and recovery procedures that will be used for the data movement period.
- Key contacts for data configuration, data movement, and backup/recovery.

Review the plan with all involved system support, operations, and user groups to ensure its viability.

A written data movement plan has several important benefits:

- It provides a method of verifying the workability of the data movement plan. Potential problems can be handled before, rather than during, the data movement effort.
- It tells the users whose data is being moved what will be happening to their data. Informed users are generally more willing to cooperate with data movement activities than uninformed users.
- It provides a record of the location of data that will be vital if data recovery is necessary.
- It provides a working document that the members of the migration team can use to help coordinate their efforts. With a written plan in place, many team members can act simultaneously.
- It allows management to be aware of the importance of the migration and of the need to provide sufficient financial and personnel resources for data movement.

After your data movement plan has been approved, verify that the necessary blocks of system time have been scheduled. You should now have a reasonably accurate idea of how much time will be needed to move each group of data, and when the time will be needed. After scheduling the necessary system time, you will be ready to begin the actual data movement. Details on how to move data under VM are in Chapter 8, "Moving Data onto 3380 Volumes" on page 83.

Chapter 6. Installing the 3380 under VM

Having completed the planning activity described in Chapters 2 through 5, you're ready to install the 3380 units on your VM system. The DASD installation process includes the following tasks:

- Defining the 3380 units to the system
- Physically installing the 3380 units
- Preparing the 3380 volumes for use

Defining the 3380 to the System

Before you can use the 3380, you must identify its existence and location to VM. It's best to do this before you actually install the 3380 units, so that you can test the configuration definition in advance.

Defining the 3380 to the operating system includes three steps:

1. Assigning I/O addresses or device numbers to the new units
2. Defining the units to VM (by updating the real I/O configuration macros in DMKRIO, for VM/SP and VM/SP HPO, or in HCPRIO for VM/XA)
3. Defining the units to the processor (by running the I/O Configuration Program, where necessary)

The sections that follow describe these steps in more detail.

Assigning I/O Addresses or Device Numbers

VM/SP and VM/SP HPO use the **I/O address** to identify both a 3380 device and the path used to reach it; thus, a given 3380 device can have several different addresses, one for each path by which it can be reached.

VM/XA uses the **device number** to identify a 3380 device only; because the path is selected dynamically by the channel subsystem, there is no specific path identifier within the device number.

You will need to assign I/O addresses or device numbers for all the devices in your new 3380 units and supply these addresses or numbers to the system. You will also need to provide these values to your service representative so the 3880, 3990, and 3380 units can be set to recognize them.

If you plan to install more 3380 units in the future, you may want to assign their addresses or numbers now. You can define these units to VM during the configuration definition process, and at IPL time VM will mark the units that have not been installed as offline.

For detailed information on how to assign I/O addresses or device numbers to the 3380, see either *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference* or *IBM 3380 Direct Access Storage Introduction*.

Defining the Real I/O Configuration to VM

The real I/O configuration file (in VM/SP and VM/SP HPO, called DMKRIO; in VM/XA, called HCPRIO) consists of macros that describe the I/O devices, storage controls, and channels attached to the real processor. Because VM uses this information to schedule I/O and to allocate resources, the macro entries in the real I/O configuration file must accurately represent the real I/O hardware configuration.

If you are adding new devices or changing hardware configurations, you must update the real I/O configuration file to include the new devices.[4]

The following real I/O configuration macros are relevant to 3380 strings:

- RDEVICE, which generates a real device block
- RCTLUNIT, which generates a real storage control block
- RCHANNEL, which generates a real channel block

For VM/XA systems, only the RDEVICE macro is valid.

Note: The following sections are meant to be a summary of the most frequently used parameters. For complete information on parameter combinations and restrictions, see the appropriate VM manual:

VM/SP	All releases	*VM/SP Planning Guide and Reference*
VM/SP HPO	All releases	*VM/SP HPO Planning Guide and Reference*
VM/XA SF	All releases	*VM/XA SF Installation, Administration, and Service*

Coding the RDEVICE Macro

You must code an RDEVICE macro for each I/O device (or group of devices with contiguous addresses or device numbers) in the configuration. The RDEVICE macro for the 3380 Model CJ2 is coded the same as for a 3380 Model AJ4 (even though there are only 2 devices in a 3380 Model CJ2) to ensure proper system handling of all conditions that may arise. The relevant RDEVICE parameters are:

ADDRESS = (*cuu,nn*) or (*ccuu,nn*)

For VM/SP systems, *cuu* is a 3-digit hexadecimal address from 000 to FFF. The address range for 3380s is 100 to FFF, since 3380s may only be attached to block multiplexer channels. The high-order digit is the address of the channel to which the device is attached; the two low-order digits represent the storage control and device address.

For VM/SP HPO systems, *ccuu* is a 4-digit hexadecimal address from 0000 to 1FFF. The address range for 3380s is 0100 to 1FFF, since 3380s may only be attached to block multiplexer channels. The two high-order digits are the address of the channel to which the device is attached; the two low-order digits represent the storage control and device address.

---

[4] Note that once the real I/O configuration file is updated to reflect the new devices, the corresponding device control blocks (RDEVBLOK, RCUBLOK, RCHBLOK) occupy space in real storage, regardless of whether the devices are actually installed. If you have severe real storage limitations, you may want to delay updating the real I/O configuration file until shortly before device installation.

For both systems, *nn* is the number of RDEVBLOKs to be generated (one for each volume in the string). For a full 3380 2-path string, this number should be 16. For two full 3380 2-path strings, code it 32. For a full 3380 4-path string, this number should be 32, and For two full 3380 4-path strings, code it 64.

DEVNO = (*rdevno,nnn*)

For VM/XA systems, *rdevno* is a 4-digit device number from 0000 to FFFF. The DEVNO range for 3380s is 0100 to FFFF, since 3380s may only be attached to block multiplexer channels. The digits are assigned to represent a specific device or volume on a 3380 unit.

nnn is the number of RDEVBLOKs to be generated (one for each volume in the string). For a full 3380 2-path string, this number should be 16. For two full 3380 2-path strings, code it 32. For a full 3380 4-path string, this number should be 32, and For two full 3380 4-path strings, code it 64.

DEVTYPE = *type*

type is the type of device, and should be **DEVTYPE** = **3380**.

ALTCU = *cuu* **or** *ccuu*

cuu or *ccuu* specifies an alternate storage control address to be used if paths through the primary storage control are unavailable. For VM/SP systems, *cuu* is a 3-digit hexadecimal address. For VM/SP HPO systems, *ccuu* is a 4-digit hexadecimal address; for VM/XA systems, *ccuu* is a 4-digit device number. **ALTCU** is an optional parameter; only one **ALTCU** can be specified per RDEVICE macro.

SHARED = **YES|NO**

For VM/XA systems, SHARED indicates whether the device should be defined as shared among systems.

Sample RDEVICE macros are shown in Figure 31 on page 65 through Figure 35 on page 67.

Coding the RCTLUNIT Macro

You must code an RCTLUNIT macro for each real storage control in the configuration. Note that each storage director in a 3880 is separately-addressable and requires its own RCTLUNIT macro. You must code an RCTLUNIT macro for each separately-addressable storage director in a 3990 storage control. The relevant RCTLUNIT parameters are:

ADDRESS = *cuu* **or** *ccuu*

For VM/SP systems, *cuu* is a 3-digit hexadecimal address from 000 to FE0. For 3380s, the range for *cuu* is 100 to FE0. The high-order digit is the address of the channel to which the device is attached; the two low-order digits represent the storage control and device address. The low-order digit in this address must be 0 (for example, ADDRESS = 140).

For VM/SP HPO systems, *ccuu* is a 4-digit hexadecimal address from 0000 to 1FE0. For 3380s, the range for *ccuu* is 0100 to 1FE0. The two high-order digits are the address of the channel to which the device is attached; the two low-order digits represent the storage control and device address. The low-order digit in this address must be 0 (for example, ADDRESS = 0140), even for the 3380 Model CJ2.

CUTYPE = *type*

type is the type of storage control, and should be coded **CUTYPE = 3880** or **CUTYPE = 3990**. For the 3380 Model CJ2, this parameter should be coded as **CUTYPE = 3380**.

FEATURE = *xxx*-**DEVICE**

This parameter must be specified, and is coded **FEATURE = 32-DEVICE** for 3880 or 3990 Model 1 or 2, **FEATURE = 64-DEVICE** for 3990 Model 2. For the 3380 Model CJ2, code this parameter as **FEATURE = 16-DEVICE**.

ALTCH = *(n,n,n)* **or** *(nn,nn,nn)*

(n,n,n) or *(nn,nn,nn)* specifies the alternate channel(s) to be used with the control unit address if the primary channel path is unavailable. For VM/SP systems, *n* represents the one-digit channel address for the alternate channel paths, and for 3380s may be any number from 1 to F. For VM/SP systems, *nn* represents the one- or two-digit channel address for the alternate channel paths, and for 3380s may be any number from 1 to 1F. As many as three alternate paths may be specified for attached processor or uniprocessor systems, only one alternate path may be specified for a multiprocessor system. **ALTCH** is an optional parameter.

Sample RCTLUNIT macros are shown in Figure 31 on page 65 through Figure 35 on page 67.

Coding the RCHANNEL Macro

You must code an RCHANNEL macro for each real channel in the configuration. The relevant RCHANNEL parameters are:

ADDRESS = *address*

For VM/SP systems, *address* is a 1-digit hexadecimal channel address from 0 to F. For 3380s, this range is 1 to F. For VM/SP HPO systems, *address* is a 2-digit hexadecimal channel address from 00 to 1F. For 3380s, this range is 01 to 1F.

CHTYPE = SELECTOR|MULTIPLEXOR|BLKMPXR|FTA

This parameter indicates the type of channel: selector, byte multiplexer, block multiplexer, or file tape adapter. Only block multiplexer, **BLKMPXR,** is valid for the 3380.

Sample RCHANNEL macros are shown in Figure 31 on page 65 through Figure 35 on page 67.

Examples of Updating DMKRIO

The installation examples shown in this book assume the use of either VM/SP 4.0 or VM/SP HPO 4.2.

The configuration shown in Figure 30 includes two 3880 storage controls and four 3380 units. The primary channel path ID is 04 (processor 0, VM channel 4) and the alternate channel path ID is 24 (processor 2, VM channel 4). This provides channel balancing in a symmetrical configuration and applies only to multiprocessor and dual or dyadic processors.

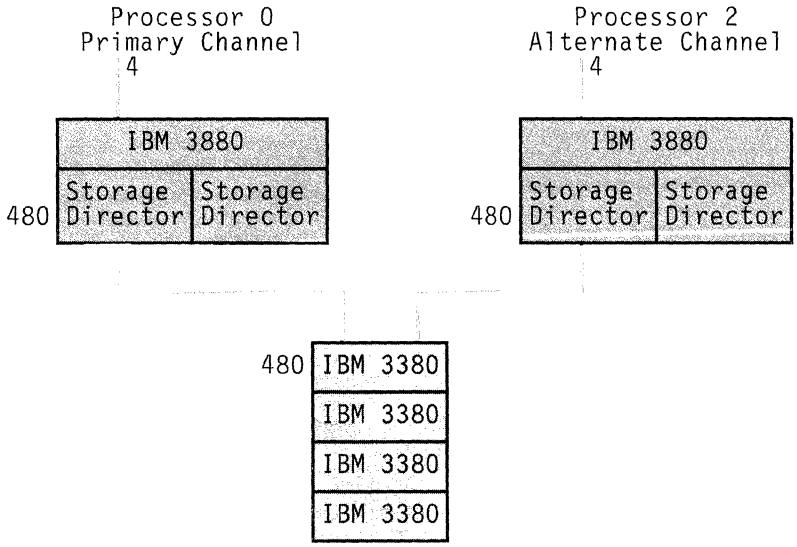

Figure 30. 3380-3880 Sample Configuration

Figure 31 shows how the configuration shown in Figure 30 would be defined to VM/SP in the DMKRIO file.

```
RDEVICE ADDRESS=(480,16),DEVTYPE=3380,ALTCU=480
RCTLUNIT ADDRESS=480,CUTYPE=3880,FEATURE=32-DEVICE
RCHANNEL ADDRESS=4,CHTYPE=BLKMPXR
```

Figure 31. Sample DMKRIO File

Establishing Different Primary Paths to One String

Figure 32 shows how to establish a primary path to two devices on a string from one channel, and a primary path to the other 14 devices from another channel. This enables the processor to drive the two volumes on channel 160 to very high rates without causing channel contention for the other 14 volumes on this string.

```
RDEVICE ADDRESS=(160,2),DEVTYPE=3380,ALTCU=260
RDEVICE ADDRESS=(262,14),DEVTYPE=3380,ALTCU=160
RCTLUNIT ADDRESS=160,CUTYPE=3880,FEATURE=16-DEVICE
RCTLUNIT ADDRESS=260,CUTYPE=3880,FEATURE=16-DEVICE
RCHANNEL ADDRESS=1,CHTYPE=BLKMPXR
RCHANNEL ADDRESS=2,CHTYPE=BLKMPXR
```

Figure 32. Sample DMKRIO File Defining Different Primary Paths to One String. Channel 1 is the primary path to devices 0 and 1, channel 2 is the primary path to devices 2 through F.

Defining a 3380 Four-path String to DMKRIO

Figure 34 shows how to define the configuration shown in Figure 33 to VM/SP in the DMKRIO file.

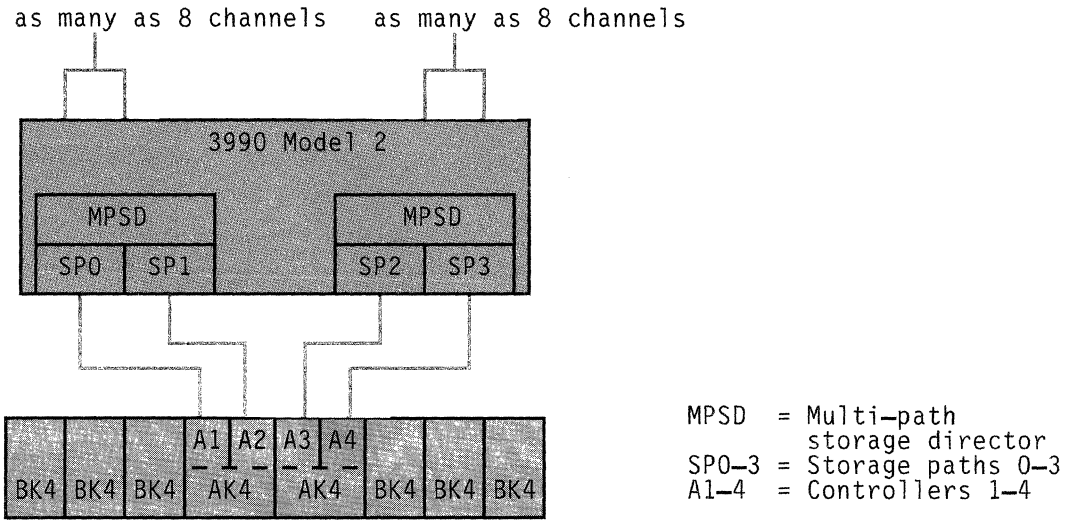

Figure 33. Four-Path Configuration Example

The corresponding IOCP statements to define the 4-path string are shown in Figure 37 on page 68. In a 4-path configuration, the control unit address must begin at 00, 40, 80, or C0.

```
RDEVICE ADDRESS=(140,16),DEVTYPE=3380,ALTCU=240
RDEVICE ADDRESS=(250,16),DEVTYPE=3380,ALTCU=150
RCTLUNIT ADDRESS=140,CUTYPE=3990,FEATURE=32-DEVICE
RCTLUNIT ADDRESS=240,CUTYPE=3990,FEATURE=32-DEVICE
RCHANNEL ADDRESS=1,CHTYPE=BLKMPXR
RCHANNEL ADDRESS=2,CHTYPE=BLKMPXR
```

Figure 34. Sample DMKRIO File Defining a 4-Path String

Defining a 3380 Model CJ2 String to DMKRIO

Figure 35 shows how the configuration shown in Figure 26 on page 51 could be defined to VM/SP in the DMKRIO file.

```
RDEVICE ADDRESS=(3F0,16),DEVTYPE=3380,ALTCU=(4E0)
RCTLUNIT ADDRESS=3F0,CUTYPE=3380,FEATURE=16-DEVICE
RCTLUNIT ADDRESS=4E0,CUTYPE=3380,FEATURE=16-DEVICE
RCHANNEL ADDRESS=3,CHTYPE=BLKMPXR
RCHANNEL ADDRESS=4,CHTYPE=BLKMPXR
```

Figure 35. Sample DMKRIO File for Attaching 3380 Model CJ2 String

For details on how to update DMKRIO or HCPRIO, see the appropriate VM manual:

VM/SP	All releases	*VM/SP Planning Guide and Reference*
VM/SP HPO	All releases	*VM/SP HPO Planning Guide and Reference*
VM/XA SF	All releases	*VM/XA SF Installation, Administration, and Service*

Defining the I/O Configuration to the Processor

For System/370 extended architecture processors (the 308x or 3090 in any mode, or the 4381 in extended architecture mode), you must run the I/O Configuration Program (IOCP) to define the I/O configuration to the processor if you are adding new devices or changing hardware configurations.

Invoke the CMS version of IOCP to generate a new input/output configuration data set. Figure 36 shows how a configuration similar to the one described by Figure 30 on page 65 would be defined to VM in the IOCP source file. For details on running IOCP under CMS, see *IOCP User's Guide and Reference*.

```
CHPID PATH=((04,4,0),(24,4,1),...),TYPE=BL
CNTLUNIT CUNUMBER=043,UNIT=3880,UNITADD=((80,32)),                    -
        SHARED=N,PROTOCL=S,PATH=(04)
CNTLUNIT CUNUMBER=053,UNIT=3880,UNITADD=((80,32)),                    -
        SHARED=N,PROTOCL=S,PATH=(24)
IODEVICE UNIT=3380,ADDRESS=(480,16),CUNUMBR=(043,053)
```

Figure 36. Sample IOCP Source File. Ellipses (...) indicate other parameters that are not shown.

Defining a 3380 Four-path String with IOCP

Figure 37 shows how the configuration defined to DMKRIO by the statements in Figure 34 on page 66 could be defined in the IOCP source file. An additional 32 devices have been reserved for a future nondisruptive DASD install.

```
CHPID PATH=((02,2,0),(03,3,0),(04,4,0),(05,5,0),(06,6,0),...),TYPE=BL
CNTLUNIT CUNUMBR=02C,PATH=(02),SHARED=N,UNIT=3990,                    -
      UNITADD=((40,64)),PROTOCL=S
CNTLUNIT CUNUMBR=03C,PATH=(03),SHARED=N,UNIT=3990,                    -
      UNITADD=((40,64)),PROTOCL=S
CNTLUNIT CUNUMBR=04C,PATH=(04),SHARED=N,UNIT=3990,                    -
      UNITADD=((40,64)),PROTOCL=S
CNTLUNIT CUNUMBR=06C,PATH=(06),SHARED=N,UNIT=3990,                    -
      UNITADD=((40,64)),PROTOCL=S
IODEVICE ADDRESS=(140,64),UNIT=3380,                                  -
      CUNUMBR=(02C,03C,04C,06C)
```

Figure 37. IOCP Generation for a Four-Path 3380-3990 Configuration

Defining a 3380 Model CJ2 String with IOCP

Figure 38 shows how the configuration defined to DMKRIO in Figure 35 on page 67 could be defined in the IOCP source file if you were going to attach that string to an extended architecture processor.

Note: The CNTLUNIT UNIT parameter is 3380 and the IODEVICE ADDRESS parameter has a base address of 0 with a maximum string length of 16.

```
CHPID PATH=((03,3,0),(04,4,0)...),TYPE=BL
CNTLUNIT CUNUMBER=043,UNIT=3380,UNITADD=((F0,16)),                    -
         SHARED=N,PROTOCL=S,PATH=(03)
CNTLUNIT CUNUMBER=044,UNIT=3380,UNITADD=((F0,16)),                    -
         SHARED=N,PROTOCL=S,PATH=(04)
IODEVICE UNIT=3380,ADDRESS=(3F0,16),CUNUMBR=(043,044)
```

Figure 38. Sample IOCP Source File. Ellipses (...) indicate other parameters that are not shown.

Installing the 3380 Units

Any time after the new devices have been defined to the system, you can physically install the 3380 units. Work with your IBM representative to ensure that you have made any physical changes necessary in your computer complex to meet the environmental needs of the 3380.

The physical installation instructions shipped with your 3380 units describe how to arrange power, cabling, and cooling for these devices. For more information on physical device installation, see *IBM Input/Output Equipment: Installation—Physical Planning for System/360, System/370, and 4300 Processors*.

Preparing Volumes for Use

After the 3380 units are installed, you must prepare the volumes for use by VM. Volume preparation includes initialization (defining device addresses and volume labels to VM and preparing the device for use by VM), formatting (preparing the disk surface to receive and store data), and allocation (reserving space on the volume for specific data).

In the VM environment, you can use the following tools to prepare volumes:

- The stand-alone version of Device Support Facilities (ICKDSF), to perform parts of the initialization process, and for initializing minidisks for use by a guest system
- The CP format/allocate program, to format and allocate volumes to be used by CP (for example: paging or spooling)
- The VM/Directory Maintenance program (DIRMAINT), to define minidisks to the VM system
- The CMS command FORMAT, to format CMS minidisks

Avoid initializing and formatting volumes during prime shift, as these tasks may tie up processor channels for extended periods of time.

After being initialized, formatted, and allocated, the volume is ready for use in the VM environment.

Initializing a Volume Using Device Support Facilities

The stand-alone version of Device Support Facilities is used to initialize volumes for use by guest systems. It can also be used as part of the initialization procedure for all devices for post installation verification.

The INIT command at the minimal level is used to write a volume label, a VTOC and other associated information on volumes or minidisks that will be used by guest systems. If the initialization is being done from within a Virtual Machine, then a minimal initialization can be done if the DASD is attached to the Virtual Machine or is a minidisk defined via the CP directory.

The INIT command at the medial level can be used to ensure the best performance of your device. For medial initialization and most other ICKDSF commands, the DASD has to be ATTACHed to the Virtual Machine.

For information on using ICKDSF with 3380s, see *Device Support Facilities: Primer for the User of IBM 3380 Direct Access Storage*. For complete reference information, see *Device Support Facilities User's Guide and Reference*.

Figure 39 shows an EXEC procedure (PREPDISK EXEC) that you can use to invoke Device Support Facilities.

```
/*  PREPDISK - Invoke ICKDSF using commands in a control file.  */
trace 'o'
address command
dsf='IPL DSF *'  /* name of ICKDSF program */
dsfdata='CNTL DSF *' /* name of control file */
'STATE' dsf /* check for ICKDSF program */
 if rc¬=0 then do
               say "ICKDSF program ("dsf") not found."
               exit rc
              end
'STATE' dsfdata /* check for data */
 if rc¬=0 then do
               say "ICKDSF control file ("dsfdata") not found."
               exit rc
              end
'EXECIO 1 CP (VAR CPQ STR Q RDR CLASS I'
if cpq¬='NO RDR FILES' then do
                            say "You have class I reader files."
                            say "PREPDISK cannot run."
                            exit 24
                           end
'CP CLOSE READER' /* send the ICKDSF IPL text to the RDR */
'CP SPOOL PUNCH TO * CLASS I CONT NOHOLD'
'PUNCH' dsf '(NOH' /* send ICKDSF */
'PUNCH' dsfdata '(NOH' /* send the control file */
'CP SPOOL PUNCH NOCONT CLOSE'
'CP SPOOL READER CLASS I NOHOLD'
'CP ORDER READER CLASS I' /* move ICKDSF to front of RDR queue */
say "Push the ENTER key to begin ICKDSF execution."
'CP IPL 00C CLEAR' /* IPL ICKDSF */
```

Figure 39. PREPDISK EXEC—a REXX EXEC to Invoke Device Support Facilities. This EXEC assumes that stand-alone Device Support Facilities has been copied into a CMS file named IPL DSF. You must be operating in EC mode to run this EXEC.

Figure 40 shows the control statement file (CNTL DSF) that contains the Device Support Facilities statements used to initialize the volume.

```
 INIT UNITADDRESS(380) NOVERIFY VOLID(VMSY84)
```

Figure 40. CNTL DSF—Initializing a VM Volume with Device Support Facilities. Note the text begins in column 2.

Execute the PREPDISK procedure by entering the filename of the EXEC:

```
prepdisk
```

After Device Support Facilities has completed loading (this may take several minutes), press the enter key or the request key on the console to continue. The following message will then appear:

```
ICK005E DEFINE INPUT DEVICE, REPLY 'DDDD,CUU' OR 'CONSOLE'
ENTER INPUT/COMMAND:
```

You should enter:

```
2540,00C
```

to indicate that the control statements should be read from your card reader, which is a virtual 2540 device at virtual address 00C. The following message will be issued:

```
ICK006E DEFINE OUTPUT DEVICE, REPLY 'DDDD,CUU' OR 'CONSOLE'
ENTER INPUT/COMMAND:
```

You should enter:

```
console
```

to indicate that the utility output should be sent to your console. The following message will then be issued:

```
ICK003D REPLY U TO ALTER VOLUME 380 CONTENTS, ELSE T
ENTER INPUT/COMMAND:
```

Your reply to message ICK003D should be:

```
u
```

to cause the program to continue.

When the Device Support Facilities program has completed, your virtual machine will be in VM read and you must re-IPL CMS to resume virtual machine execution. If you prefer, you can create an exec to perform the above activities and then request that it execute during off-shift hours.

For information on using ICKDSF with 3380s, see *Device Support Facilities: Primer for the User of IBM 3380 Direct Access Storage*. For a complete description of the control statements used, and a detailed discussion of how to initialize volumes under VM, see *Device Support Facilities User's Guide and Reference*.

Formatting a CP-Owned Volume

CP references to DASD space are always in terms of DASD pages. A DASD page is 4096 bytes of contiguous DASD storage.

CP requires all its system areas (nucleus, error recording, warm start data, checkpoint data, directory, override, saved systems, dump space, paging, and spooling) to be formatted as 4K-byte pages. CP also requires you to allocate specific cylinders on DASD for paging, spooling, dump, temporary space, and for the directory. Any volume containing CP system areas is considered a CP-owned volume.

For CP volumes, you must use the CP format/allocate program to format the specified cylinders into 4K-byte pages. Any remaining space on CP-owned volumes can be used for user minidisks; this space does not have to be formatted into 4K-byte pages, but it must be allocated as PERM.

The steps in preparing a volume for use by CP are:

Step 1. Use the CP format/allocate program to format and allocate the volume.

Step 2. Define the CP volume in the SYSOWN list (located in DMKSYS for VM/SP and VM/HPO, in HCPSYS for VM/XA).

Using the CP Format/Allocate Program

The CP format/allocate program is a service program that:

- Formats volumes for CP use
- Allocates specific cylinders for particular functions
- Writes a volume label

The CP format/allocate program can be loaded from a reader or tape unit into a virtual or a real machine.[5] (If run in a virtual machine, the virtual machine must have write access to the volume being formatted.) The program accepts control statements from the operator's system console or from the IPL device (reader).

If the program finds no control statements at the reader, it issues a prompting message to the console. The proper response causes the prompting message for the next operand to appear until the format, allocate, or label function is completely defined; then the format or allocate is executed.

The CP format/allocate program has been changed to allow you to decide whether or not write verification is performed. The default is no write verification. In most cases, it is unnecessary to perform write verification on previously-formatted 3380 devices. If you are formatting a 3380 that has never been formatted before, you may want write verification to occur and you would reply yes to the message:

```
ENTER 'YES' FOR WRITE VERIFICATION:
```

Figure 41 on page 73 is an example of format/allocate program execution under CP control. All responses are entered after the colon; after a function is complete, the program returns and again issues the message:

```
ENTER FORMAT OR ALLOCATE:
```

---

[5] In VM/XA, you can invoke the CP format/allocate program by using the CP command CPFORMAT.

```
VM/370 FORMAT/ALLOCATE PROGRAM - VM/SP
ENTER FORMAT OR ALLOCATE:
format
FORMAT FUNCTION SELECTED
ENTER DEVICE ADDRESS (CCUU):
391
ENTER DEVICE TYPE:
3380
ENTER START CYLINDER (XXX OR XXXX) OR "LABEL":
0000
ENTER END CYLINDER (XXX OR XXXX):
2654
ENTER DEVICE LABEL:
m35tst
WRITE VERIFICATION NOT PERFORMED UNLESS REQUESTED.
ENTER "YES" FOR WRITE VERIFICATION:
no
FORMAT STARTED
FORMAT DONE
WRITE VERIFICATION WAS NOT PERFORMED
ENTER FORMAT OR ALLOCATE:
allocate
ALLOCATE FUNCTION SELECTED
ENTER DEVICE ADDRESS (CCUU):
391
ENTER DEVICE TYPE:
3380
ENTER DEVICE LABEL:
m35tst
ENTER ALLOCATION DATA FOR VOLUME M35TST
TYPE CYL  CYL
perm 000 000
temp 001 0030
tdsk 031 100
page 101 200
perm 201 2554
end
ALLOCATION RESULTS
PERM 0000 0000
TEMP 0001 0030
TDSK 0031 0100
PAGE 0101 0200
PERM 0201 2554
TEMP 2555 2654
DEVICE 391  VOLUME M35TST ALLOCATION ENDED
```

Figure 41. Formatting and Allocating a Triple Capacity 3380 (under VM/SP)

For more information on using the CP format/allocate program to format CP volumes, see the appropriate VM *Operations* manual:

VM/SP	Up through Release 4	*VM/SP Operator's Guide*
VM/SP	Release 5 or later	*VM/SP CP for System Programming*
VM/SP HPO	All releases	*VM/SP HPO Operator's Guide*
VM/XA SF	All releases	*VM/XA SF Real System Operation*

Defining the CP Volume in the SYSOWN List

After formatting, you must make CP volumes available for use by:

1. Attaching them to the system at IPL or by operator command; and
2. Listing their volume labels in the SYSOWN macro in the DMKSYS module for VM/SP or VM/HPO or the HCPSYS module in VM/XA.

The CP system residence volume and CP volumes containing paging, spooling, directory, and temporary space must be defined in the SYSOWN list stored in DMKSYS or HCPSYS. To add a CP volume to the end of the SYSOWN list, use the SYSOWN macro as shown in Figure 42.

```
SYSOWN VMSRES,VMPG01
```

Figure 42. Example of the SYSOWN Macro

For more information on defining CP volumes in the SYSOWN list, see the appropriate VM *Planning* manual:

VM/SP	All releases	*VM/SP Planning Guide and Reference*
VM/SP HPO	All releases	*VM/SP HPO Planning Guide and Reference*
VM/XA SF	All releases	*VM/XA SF Virtual Machine Planning*

Defining a CMS Minidisk

As shown in Figure 41 on page 73, you allocate space for CMS user minidisks by using the CP format/allocate program. After space has been allocated, you must define the minidisks to the VM directory before they can be used.

CMS minidisks with a block size of 4K can range from 1 to 2655 cylinders in size. However, the CMS file system will not support files with records longer than 2 147 483 647 bytes or files with more than 2 147 483 647 records. If you are allocating a 3380 Model AE4, BE4, AK4, or BK4 volume as a single CMS minidisk, plan your file capacity accordingly.

You can define minidisks in the VM directory using the DIRMAINT command. Figure 43 shows an example of the DIRMAINT command used to define a new minidisk for user DAVEC. The minidisk, at virtual address 193, starts at real cylinder 15 on the 3380 volume called VMMD04, and continues for 10 cylinders.

```
DIRMAINT AMDISK DAVEC 193 3380 15 10 VMMD04
```

Figure 43. Example of the DIRMAINT Command Used to Define a Minidisk

For more information on using the DIRMAINT command to define minidisks, see *VM/Directory Maintenance Installation and System Administrator's Guide*. If you don't have the VM/Directory Maintenance program, see the appropriate VM manual for information on how to update the directory:

VM/SP	All releases	*VM/SP Planning Guide and Reference*
VM/SP HPO	All releases	*VM/SP HPO Planning Guide and Reference*
VM/XA SF	All releases	*VM/XA SF Installation, Administration, and Service*

Formatting a CMS Minidisk

You can format CMS user minidisks with the FORMAT command of CMS. For standard minidisks, the CMS FORMAT command formats the specified tracks into 512-byte, 800-byte, 1024-byte, 2048-byte, or 4096-byte blocks. FORMAT is usually invoked by the owner of the new minidisk.

The CMS FORMAT command can:

- Format a minidisk for use with CMS files
- Count or reset the number of cylinders on a virtual disk
- Write a label on a virtual disk

The relevant parameters for the FORMAT command are:

cuu
: The virtual device address of the minidisk to be formatted, which is defined in the directory. See "Defining a CMS Minidisk" on page 74 for information on this task.

mode
: The filemode letter to be assigned to the minidisk.

(BLKSIZE *n*
: The physical DASD block size of the CMS minidisk, where *n* is 512, 800, 1024, 2048, or 4096 bytes. The recommended block size for DASD is 4096 bytes.

Note: The default block size for the CMS FORMAT command is 4096 bytes in VM/SP and VM/SP HPO Release 5 or later. In VM/XA SF1, or VM/XA SF 2, the default block size is 1024 bytes. You must explicitly specify another block size to override this default. "Selecting a Block Size for CMS Minidisks" on page 76 discusses how to choose a CMS minidisk block size.

An example of the CMS FORMAT command is shown in Figure 44.

```
FORMAT 191 A (BLKSIZE 4096
```

Figure 44. Example of the CMS FORMAT Command

For details on how to code the CMS FORMAT command, see the appropriate *CMS Reference*:

VM/SP	Up through Release 4	*VM/SP CMS Command and Macro Reference*
VM/SP	Release 5 or later	*VM/SP CMS Command Reference*
VM/SP HPO	Up through Release 4	*VM/SP HPO CMS Command and Macro Reference*
VM/SP HPO	Release 5 or later	*VM/SP HPO CMS Command Reference*
VM/XA SF	All releases	*VM/XA SF CMS Command and Macro Reference*

Selecting a Block Size for CMS Minidisks

CMS minidisks and CP data areas are allocated by cylinders. However, because data is transferred to and from DASD in blocks, block size influences performance and is important in determining the best use of space on a 3380 volume.

Small block sizes permit more concurrent operations on a channel, but they reduce the net data transfer rate. Small block sizes require more processor involvement, and use more device space for a given amount of data.

Large block sizes allow a high net data transfer rate and reduce the amount of processor time needed to process a channel program.

The block size used by CP data areas is automatically set at 4096 bytes by the CP format/allocate program. Unlike CP, CMS allows a choice of several block sizes for CMS minidisks. The two main factors that affect the choice of a minidisk block size are the way that CMS stores files, and the capacity of the 3380 for each block size.

More information on block size selection may be found in *Comparison of IBM 3380s and IBM 3350s Used for VM/CMS Minidisks*.

CMS File Storage: The CMS file system uses an index structure to locate each file. The lowest level of index points to the data blocks that contain whole or partial files on the minidisk. Higher levels of the index point to lower levels of the index. As the number of data blocks and files on a minidisk increases, the number of index levels increases.

Each index block is the same size as one data block. Large index blocks can point to a greater number of files and require fewer additional index levels. Large blocks improve overall system performance because a smaller number of index blocks are needed to locate the data blocks for the file.

Similarly, CMS can locate a file that resides entirely within one data block with fewer accesses to the index than a file contained within five data blocks. Therefore, a file stored in large (4096 byte) data blocks can improve system performance by reducing the amount of time needed to read and write the file's data blocks.

For more information on CMS file storage, see the appropriate VM/SP or VM/SP HPO *System Logic and Problem Determination Guide Volume 2 - CMS*.

Capacity of 3380 Devices with Available CMS Block Sizes: Depending on the minidisk block size, a 3380 volume can hold various amounts of data. The capacity (in megabytes) of 3380 volumes at various block sizes (if every block is filled with data), is shown Figure 45. As this chart indicates, a block-size of 4096 bytes allows more data to be placed on a volume. In addition, by increasing the amount of data brought into main storage each time the DASD is accessed, a 4096-byte block size can reduce the number of start I/O requests.

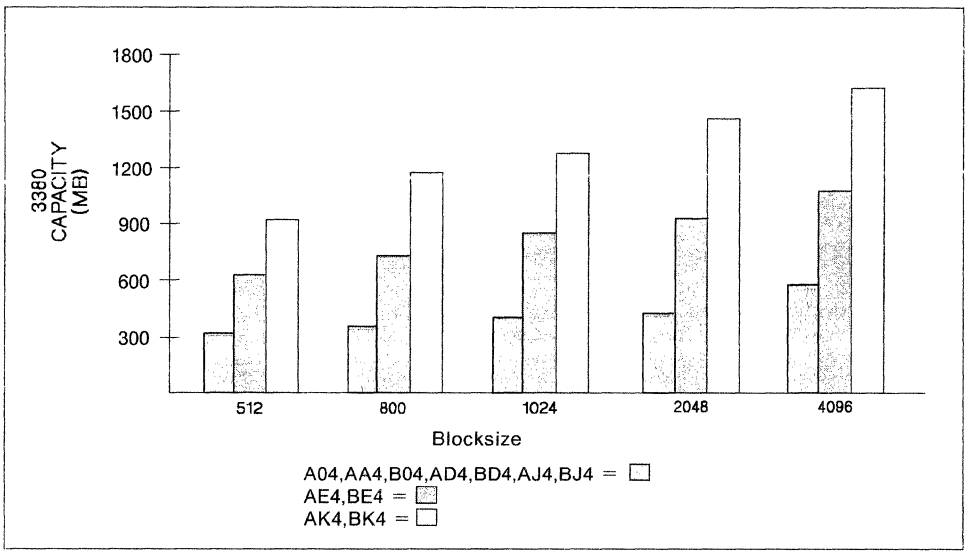

Figure 45. Maximum Volume Data Capacity vs. Block Size

Initializing a Minidisk for MVS, VSE, or VSE/VSAM

You can initialize minidisks for use by MVS or VSE guest systems, or for VSE/VSAM files, using the stand-alone version of Device Support Facilities (ICKDSF). Use the PREPDISK EXEC shown in Figure 39 on page 70 to invoke Device Support Facilities under CMS, and substitute the control statement file (CNTL2 DSF) shown in Figure 46.

```
 INIT UNIT(198) NOVERIFY MIMIC(MINI(3)) DEVTYPE(3380) VOLID(MVSG02)
```

Figure 46. CNTL2 DSF—Initializing a Minidisk in a VM Stand-Alone Environment. Note the text begins in column 2.

This example provides 30 primary and no alternate cylinders on unit 198. The VTOC is written on a default location of cylinder 0, track 1 for a length of one track. The volume is labeled MVSG02, and the label and VTOC are written in OS format. Device 198 has already been linked to the userid. If you want to place IPL text on the device, see *Device Support Facilities User's Guide and Reference*, for specific parameters to include in the INIT statement.

For VSE or VSE/VSAM, modify this example by adding the DOSVTOC parameter. There are restrictions on minidisk size in the VSE environment and you must define VSE minidisks within the size limitations of your VSE system. For example, if the largest DASD device supported in your VSE system is a 3380 Model AE4, then the largest minidisk you can define is 1770 cylinders. For more information, see *Using the IBM 3380 Direct Access Storage in a VSE Environment.*

Chapter 7. Operating the 3380 under VM

This chapter discusses things an operator needs to know about using the 3380 under VM. The intent of this chapter is to give your operations staff enough information to understand how to properly operate the 3380 string, so that you can prepare specific procedures for the operators in your environment.

Most changes to the 3380 operational status are controlled from the system operator userid. System commands are issued from the operator or other privileged userid to control 3380 operations and to obtain the operational status of the 3380 string.

Some 3380 operational status changes are controlled from the operator control panel. A control panel for each 3380 string is located on the end cover of 3380 Model A04, AA4, AD4, and AE4 units; or on the front cover of 3380 Model AJ4, and AK4 units. B-units do not have operator control panels. The 3380 Model CJ2 operator panel is located on the front cover of the unit. *IBM 3380 Direct Access Storage Introduction* and *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference* contain illustrations of the 3380 operator panels and describe the following operator tasks:

Reading the operator control panels
Turning the power on and off
Enabling and disabling controllers and devices

This chapter describes the following operator tasks for the 3380:

Displaying 3380 status
Varying volumes online and offline

Displaying 3380 Status

You can determine or verify the status of a particular volume under VM by using the CP QUERY command. There are several versions of the QUERY command available to display information about DASD. You must be authorized with CP privilege class B to use these **CP commands:**

QUERY DASD ACTIVE
Lists volumes that are online and attached to the system or to a VM userid.

QUERY DASD FREE
Lists volumes that are online but not attached to the system or to a VM userid. Volumes cannot be varied offline (using the VARY command) until they are detached from the system and all users (that is, until they are listed in the FREE list).

QUERY DASD OFFLINE
Lists volumes that are offline (that is, not available for access by any user or by the system).

QUERY DASD PATHS

Displays path status (VM/SP and VM/HPO only)[6] for alternate paths to volumes.

QUERY DASD *volser*

Displays the status of a specific volume with volume serial number *volser*.

QUERY SYSTEM *raddr*

Displays, for the volume with real address *raddr*, the user logon ID, virtual address, and access mode of the virtual disks on the specified volume that are currently being used by logged-on users.

For more information on the CP command QUERY, see the appropriate VM manual:

VM/SP	Up through Release 4	*VM/SP Operator's Guide*
VM/SP	Release 5 or later	*VM/SP CP Command Reference*
VM/SP HPO	All releases	*VM/SP HPO CP Command Reference*
VM/XA SF	All releases	*VM/XA SF CP Command Reference*

There are also two **CMS commands** you may find helpful for displaying device information:

QUERY DISK

Displays information about CMS, VSE, and OS-formatted minidisks that you have accessed.

LISTDS

Displays information about files residing on accessed VSE or OS minidisks.

For more information on the CMS commands QUERY and LISTDS, see the appropriate VM *CMS Reference* manual:

VM/SP	Up through Release 4	*VM/SP CMS Command and Macro Reference*
VM/SP	Release 5 or later	*VM/SP CMS Command Reference*
VM/SP HPO	Up through Release 4	*VM/SP HPO Command and Macro Reference*
VM/SP HPO	Release 5 or later	*VM/SP CMS Command Reference*
VM/XA SF	All releases	*VM/XA SF CMS Command and Macro Reference*

6 Under VM/XA SF, path information can be displayed by using "LOCATE userid vdev" followed by "DISPLAY H," or, if the guest is MVS/XA, using the MVS display matrix command, "D M=DEV(xxx)." For more information, see *VM/XA Systems Facility Planning Guide*.

Varying Volumes Online and Offline

You can use the VARY command to specify that a particular volume is either available (online) or unavailable (offline) for use by a user or by VM. The VARY command varies the volume online or offline, depending on the parameters you specify. Figure 47 shows an example of the VARY command used to vary a 3380 volume online.

```
VARY ONLINE 2C4
```

Figure 47. Example of Varying a 3380 Volume Online. You must be authorized with CP privilege class B to issue the VARY command.

Note: All Power On/Off switches must be in the "On" position and all Enable/Disable switches must be in the "Enable" position before you can vary a volume online to VM. Do not set Power On/Off switches to "Off" or Enable/Disable switches to "Disable" until the volume has been varied offline from each VM system and the offline status has been verified.

When processing several volumes or a range of volumes, VARY command processing continues regardless of whether or not an error is encountered when attempting to VARY any one of the volumes online or offline. An error message is issued for every volume that encounters an error situation.

For all releases of VM, with the exception of VM/XA SF 2, the VARY command will **not** allow you to take a single path to a device offline, but only to take all paths offline.

Notes for Guest Systems

If you are running VSE as a guest system under VM, after placing the Device Enable/Disable switch in the "Enable" position, you should vary the device online to VM then ATTACH the device to the VSE guest userid before making the device available to VSE via DVCUP. Prior to setting the Enable/Disable switch to "Disable," you should first make the device unavailable to VSE with DVCDN, DETACH the device from the VSE guest userid, and then vary it offline to VM.

For details on how to use the VARY command, see the appropriate VM manual:

VM/SP	Up through Release 4	*VM/SP Operator's Guide*
VM/SP	Release 5 or later	*VM/SP CP Command Reference*
VM/SP HPO	Up through Release 4.2	*VM/SP HPO Operator's Guide*
VM/SP HPO	Release 5 or later	*VM/SP HPO CP Command Reference*
VM/XA SF	All releases	*VM/XA SF CP Command Reference*

Chapter 8. Moving Data onto 3380 Volumes

If you are replacing existing DASD with new 3380 units, you will need to move data from old to new volumes. Data movement in VM **can** be easy, but you must take the time to identify the data movement tools available and determine which tools are most appropriate for each type of data.

This chapter describes:

- How to ensure data availability and integrity by backing up volumes before and after data movement;
- How to use VM software tools to move volumes, minidisks, files, and data areas.

Remember that all volumes (both new and old) involved in the data movement process must be defined to the VM system and varied online before any tools can be used to move data.

Tools for Moving Data

There are a number of tools available in VM to move data from one DASD volume to another.

DASD Dump Restore (DDR)
A service program shipped with VM that can be used to dump data from DASD to tape, restore data from tape to DASD, and copy data between like DASD volumes.

CMDISK
A DIRMAINT command that can move minidisks from any device type supported by VM to any other type. In addition, CMDISK can move data from fixed block architecture (FBA) devices to non-FBA devices.

COPYFILE
A CMS command used to copy files (or minidisks) between like or unlike DASD volumes.

MIG3380
A program used to assist in migrating to new 3380 models in the VM/SP 4.0 and VM/SP HPO 4.2 environments (or later). See your IBM marketing representative for applicable SPEs.

SPTAPE
A CP command used to store spool files on tape and to restore them from tape to DASD.

This chapter describes how to use these tools to back up, recover, and move data under VM.

Figure 48 summarizes the tools available and their appropriate uses for various types and amounts of data. The table also indicates where in this chapter the tool is discussed.

Data	From/To	Tool Name	Read Section
Volume	3380/3380	DDR and MIG3380	"Moving Full Volumes From 3380 to 3380" on page 86 "Using the MIG3380 Program to Assist in Volume Migration" on page 87
	non-3380/3380	None; must move by minidisk and file	"Moving Minidisks and Files Between Unlike Devices" on page 89
Minidisk	3380/3380	DDR	"Moving Minidisks From 3380 to 3380" on page 88
	3380/3380	CMDISK	"Moving Minidisks With CMDISK" on page 88
	non-3380/3380	CMDISK	"Moving Minidisks With CMDISK" on page 88
	non-3380/3380	COPYFILE	"Moving Minidisks and Files Between Unlike Devices" on page 89
Data file	3380/3380	COPYFILE	"Moving Individual Files to a 3380" on page 89
	non-3380/3380	COPYFILE	"Moving Minidisks and Files Between Unlike Devices" on page 89
Spool files	3380/tape/3380	SPTAPE	"Moving Spool Files" on page 90
	non-3380/tape/3380	SPTAPE	"Moving Spool Files" on page 90
CP data areas	3380/3380	DDR and MIG3380	"Moving CP Data" on page 90

Figure 48. Summary of Data Movement Tools in VM

Backup and Recovery during Migration

Backup is a means of ensuring data availability and integrity in case of system problems. During 3380 migration, backup is especially important because data is especially vulnerable when it is being moved. You should back up each old volume before moving its contents, and back up each new volume as soon as its contents are complete.

You can use the DDR program to back up data by dumping it from DASD to tape. DDR can also be used to restore data from tape to the DASD device type that the data was dumped from.

DDR can be run interactively under CMS or as a stand-alone program. The examples in this section assume you are using DDR interactively under CMS. Before running DDR, you must attach the tape drives and DASD devices to the virtual machine that is executing the DDR program.

The DDR command "DDR fn ft fm" identifies the file containing the control statements for the DDR program. You need to create this file with its control statements before executing the DDR command. (If no file is found, DDR will look for control statements from the console.)

For example, if the control statements were in a file called CNTL FILE A, you would invoke DDR with the command:

```
DDR CNTL FILE A
```

The control statement file should include the following:

- A SYSPRINT statement, describing the device that will receive the output listing
- An INPUT statement, describing the unit address, device type, and volume serial number of the INPUT volume
- An OUTPUT statement, describing the unit address, device type, and volume serial number of the OUTPUT volume
- A DUMP or RESTORE statement, describing the action to be taken and the cylinders to be moved

Note: Although the examples in this section show backup and recovery for 3380 volumes, you can also back up 3350 volumes by specifying a device type of 3350 instead of 3380 on the control statements. Naturally, 3350 backup volumes must also be restored to 3350 volumes.

For detailed information on how to use DDR, see the appropriate VM manual:

VM/SP	Up through Release 4	*VM/SP Operator's Guide*
VM/SP	Release 5 or later	*VM/SP CP for System Programming*
VM/SP HPO	Up through Release 4.2	*VM/SP HPO Operator's Guide*
VM/SP HPO	Release 5 or later	*VM/SP CP for System Programming*
VM/XA SF	All releases	*VM/XA SF Real System Operation*

Backing up a Volume Using DDR

With DDR, you can back up a volume by dumping its contents to either DASD or tape. Remember to choose a volume of equal or greater data capacity than the volume you are dumping. (For example, don't try to dump a 3380 BE4 volume full of data onto a 3380 BD4.) You will need multiple backup tape reels or cartridges to contain the data dumped from a DASD volume.

Figure 49 summarizes the capacity of various backup devices types by volume, including both tape and DASD.

Device Type	Volume	Data Capacity in 4K Blocks
3350	—	273 MB
3380	Model A04, AA4, B04, AD4, BD4, AJ4, BJ4, CJ2	543 MB
	Model AE4, BE4	1087 MB
	Model AK4, BK4	1631 MB
3420	tape reel	120 MB
3480	tape cartridge	126 MB

Figure 49. Summary of Data Capacity for Backup DASD and Tape Volumes

Figure 50 shows an example of dumping a 3380 volume onto a 3480 cartridge tape.

```
SYSPRINT CONS
INPUT 198 3380 VMSY81
OUTPUT 181 3480 (LEAVE MODE 38K COMPACT
DUMP 000 TO 029
```

Figure 50. Dumping a 3380 Volume to 3480 Tape with DDR

Restoring a Volume Using DDR

If there are problems during data migration, you may need to restore the dumped data from the backup volume or tape. Figure 51 shows an example of restoring a 3380 volume from a DDR backup tape.

```
SYSPRINT CONS
INPUT 181 3480 (LEAVE MODE 38K
OUTPUT 198 3380 VMSY81
RESTORE 000 TO 029
```

Figure 51. Restoring a 3380 Volume from 3480 Tape with DDR

Moving Full Volumes From 3380 to 3380

You can also use DDR to copy an entire volume of data to another DASD volume on a like device (for example, 3380 to 3380). Figure 52 shows an example of copying a 3380 volume to another 3380 volume using DDR.

```
SYSPRINT CONS
INPUT 480 3380 VMSY81
OUTPUT 2C4 3380 VMSY83
COPY ALL
```

Figure 52. Copying a Volume with DDR

After moving the data, update the VM directory to indicate the new location on the new volumes. Remember when moving or copying full volumes, the target volume must be of equal or greater data capacity. Figure 49 on page 85 lists data capacity for common DASD devices.

An XEDIT macro that you can use to update the VM directory to move two standard capacity 3380 volumes to one double capacity 3380 volume is shown in "MERGE Xedit Macro" on page 116.

Because they can rarely be varied offline, system-owned volumes should be moved stand-alone from a real machine. For information on moving CP data areas, see "Moving CP Data" on page 90.

Using the MIG3380 Program to Assist in Volume Migration

VM/SP and VM/SP HPO provide a program called MIG3380 that can help you when moving data from one 3380 volume to another, of equal or greater data capacity. Cylinder 0 of a CP-owned volume contains an allocation map which varies in size (1, 2, or 4k) depending on the capacity of the 3380. MIG3380 saves the CP information from cylinder 0, adjusts the allocation map size on cylinder 0 to the size of the target volume, then rewrites the CP information on cylinder 0 of the target volume. You no longer need to keep track of the source volume's allocation because the target volume need not be re-allocated unless the additional space on a larger target volume is to be used as something other than PERM. Furthermore, if a CP nucleus or directory space exists on the source volume, it will be migrated to the target volume. You do not have to rebuild either the CP nucleus or the directory space.

To use the MIG3380 program:

1. Use DDR to back up the source volume, as described in "Backing up a Volume Using DDR" on page 85.
2. Make sure the target volume is initialized and defined in the SYSOWN list, as described in "Defining the CP Volume in the SYSOWN List" on page 74.
3. Vary the target volume online, as described in "Varying Volumes Online and Offline" on page 81.
4. Use DDR to copy the contents of the source volume to the target volume, as described in "Moving Full Volumes From 3380 to 3380" on page 86.
5. Issue the command `MIG3380 cuu`, where `cuu` is the I/O address of the target volume; for example, `MIG3380 2C4`.
6. If the target 3380 volume has more cylinders than the source volume, use the CP format/allocate program to format the remainder of the target volume and allocate space (if you want it to be allocated as something other than PERM). Not all data areas must be formatted as explained in *VM Planning Guide and Reference*.
7. Use CP format/allocate to re-label the source volume to avoid duplicate volume names.

When this sequence is complete, the target volume is ready for use.

Moving Minidisks With CMDISK

If DIRMAINT is installed, you can use the DIRMAINT CMDISK (change minidisk) command to move individual minidisks. Figure 53 shows how to move a minidisk with the DIRM CMDISK command.

```
DIRM CMDISK userid 191 3380 AUTOV 6 VMSY83
```

Figure 53. Moving a Minidisk with DIRM CMDISK

The following tasks are executed when the command in Figure 53 is issued:

1. Allocate a new minidisk for the user on VMSY83 of size 6 cylinders. This minidisk is given a temporary address and is transferred to the directory of the service machine DATAMOVR (part of DIRMAINT).
2. Transfer the user's 191 minidisk to DATAMOVR at another temporary address.
3. Format the new minidisk using CMS FORMAT.
4. Copy all the files to the new disk using COPYFILE.
5. Transfer the new disk back to the user as his 191.
6. Return the space occupied by the old 191 to the general pool of minidisk space.

Do not use the DIRM CMDISK to move DIRMAINT minidisks. For a procedure you can use to move DIRMAINT minidisks, see *VM/Directory Maintenance Installation and System Administrator's Guide*.

An exec you can use to automate the migration of CMS minidisks is shown in Appendix B, "Migration Aids" on page 115.

Moving Minidisks From 3380 to 3380

Another way to move VM data between like devices is to use the DDR program to copy an entire VM minidisk to a like device. Figure 54 shows an example of copying a minidisk from one 3380 volume to another using DDR.

```
SYSPRINT CONS
INPUT 191 3380 VMSY81
OUTPUT 399 3380 VMSY83
COPY 000 TO 005 REORDER 10
```

Figure 54. Copying a Minidisk with DDR. The six-cylinder minidisk will occupy cylinders 10 through 15 on the output volume VMSY83.

After you have the minidisk safely stored on the new volume, you should update the VM directory to indicate the new location of the minidisk, using the DIRMAINT CMDISK command. See *VM/Directory Maintenance Installation and System Administrator's Guide* for more information.

Moving Individual Files to a 3380

To move individual user files between like or unlike devices, use the CMS command COPYFILE.

Figure 55 shows an example of how COPYFILE can be used to move all the files from a source minidisk (filemode A, linked R/W) to a target minidisk (filemode Q, linked R/W). The TYPE operand specifies that the names of the files copied should be displayed at the console. The OLDDATE operand preserves the original creation/modification date of the copied files.

```
COPYFILE * * A = = Q (TYPE OLDDATE
```

Figure 55. Moving Files to a Minidisk on Like or Unlike Devices

For more information on the COPYFILE command, see the appropriate VM *CMS Reference* manual:

VM/SP	Up through Release 4	*VM/SP CMS Command and Macro Reference*
VM/SP	Release 5 or later	*VM/SP CMS Command Reference*
VM/SP HPO	Up through Release 4	*VM/SP CMS Command and Macro Reference*
VM/SP HPO	Release 5 or later	*VM/SP CMS Command Reference*
VM/XA SF	All releases	*VM/XA SF CMS Command and Macro Reference*

Moving Minidisks and Files Between Unlike Devices

To move data between unlike devices, you can use the CMS command COPYFILE to copy each file to a minidisk on the new device. Figure 55 shows an example of using COPYFILE to copy files between like or unlike devices.

The minidisks on the new volume must have been previously defined in the CP directory, allocated by the CP format/allocate program, and initialized by the CMS FORMAT command. See "Formatting a CMS Minidisk" on page 75 for a description of how to allocate and format minidisks.

You can use the DIRMAINT CMDISK command, as shown in "Moving Minidisks With CMDISK" on page 88, to move minidisks between unlike devices.

An exec you can use to automate the migration of CMS minidisks is shown in Appendix B, "Migration Aids" on page 115.

Moving Spool Files

You can move spool files by storing them on tape and then reloading them onto a different DASD volume, using the SPTAPE command. The restored files retain the same characteristics as the original files, but are assigned a new spoolid to avoid duplicate identification within the spooling system.

An example of dumping spool files to 3480 tape and then retrieving them after a cold start is shown in Figure 56.

```
SPTAPE DUMP 530 RDR ALL MODE 38K RUN PURGE NOHOLD

SPTAPE LOAD 530 RDR ALL
```

Figure 56. Example of Dumping Spool Files Using SPTAPE. You must be authorized with CP privilege class D to use the SPTAPE command.

For more information on the SPTAPE command, see the appropriate VM manual:

VM/SP	Up through Release 4	*VM/SP Operator's Guide*
VM/SP	Release 5 or later	*VM/SP CP Command Reference*
VM/SP HPO	Up through Release 4.2	*VM/SP HPO Operator's Guide*
VM/SP HPO	Release 5 or later	*VM/SP HPO CP Command Reference*
VM/XA SF	All releases	*VM/XA SF CP Command Reference*

Moving CP Data

Most CP system areas cannot be moved but must be rebuilt on the new volumes. The following CP-formatted areas should be rebuilt if they are going to reside on the new volumes:

- Error recording area
- Checkpoint area
- Warm start area
- Paging, spooling and dump space
- Saved systems area
- 3704/3705/3725 control program area
- 3800 printer specification area

Review each of the CP system areas individually before rebuilding the data. If the checkpoint, warm start, spool, or dump areas are moved or rebuilt, a cold start is required.

For more information on rebuilding CP system areas, see the appropriate VM *Planning* manual:

VM/SP	All releases	*VM/SP Planning Guide and Reference*
VM/SP HPO	All releases	*VM/SP HPO Planning Guide and Reference*
VM/XA SF	All releases	*VM/XA SF Virtual Machine Planning*

You can copy the nucleus using the NUCLEUS option on the DDR COPY command, and you can copy the contents of cylinder 0 (including directory, override, and permanent space) using the CPVOL option on the DDR COPY command. For more information on DDR options, see the appropriate VM manual:

VM/SP	Up through Release 4	*VM/SP Operator's Guide*
VM/SP	Release 5 or later	*VM/SP CP for System Programming*
VM/SP HPO	Up through Release 4.2	*VM/SP HPO Operator's Guide*
VM/SP HPO	Release 5 or later	*VM/SP HPO System Facilities for Programming*
VM/XA SF	All releases	*VM/XA SF Real System Operation*

Chapter 9. Monitoring and Maintaining the 3380

To ensure that the 3380 strings are providing predictable and acceptable service, you must learn to manage performance and space for the storage subsystem. Regular performance and space analysis can help you identify potential problems before they occur and plan for orderly storage subsystem growth. In addition, regular media maintenance can help to extend the life of your DASD.

This chapter describes some techniques to improve performance, space utilization, and media maintenance for the storage subsystem.

Managing Performance in the Storage Subsystem

The performance of a VM system depends on many factors: workload, processor speed, channel capacity, paging and swapping frequency, communications network speed, and DASD and storage control configuration. For users, performance is measured primarily by terminal response time, which is affected by all of these factors. DASD I/O accounts for a significant portion of terminal response time in many cases. Because lower response times increase user productivity, it's worth spending time monitoring and improving performance in the storage subsystem.

You should set reasonable response time objectives for your VM system and document them as part of a service level agreement between you and your users. Reliable performance measurements, such as those collected by the VM/Monitor, will tell you whether you are meeting your service objectives.

Collecting Performance Statistics with VM/Monitor

The VM/Monitor is a standard component of VM/SP and VM/SP HPO that collects data on system performance and resource utilization. You can set up various Monitor classes to control the kinds of statistics gathered and the frequency of sampling. These classes include:

PERFORM Samples resource data for the entire system.

USER Samples resource data for individual users.

DASTAP Samples DASD and tape device I/O activity.

SEEKS Records seek activity for specific devices.

You should collect PERFORM, USER, and DASTAP data daily, during prime shift. The default Monitor interval is 60 seconds. PERFORM, USER, and DASTAP cause very little performance overhead and generate a relatively small amount of data.

To set up or change Monitor classes, you can use the SYSMON macro in DMKSYS, or the CP MONITOR command. Figure 57 on page 94 shows an example of how to code the SYSMON macro to establish the PERFORM, USER, and DASTAP classes.

```
SYSMON USERID=OPERATOR,CLASS=M,AUTO=YES,
      ENABLE=(PERFORM,USER,DASTAP),TIME=(09:00,17:00)
```

Figure 57. Sample SYSMON Macro in DMKSYS to Set Monitor Classes

For details on how to use the SYSMON macro, see the appropriate VM *Planning* manual:

VM/SP	All releases	*VM/SP Planning Guide and Reference*
VM/SP HPO	All releases	*VM/SP HPO Planning Guide and Reference*

Figure 58 shows an example of using the CP MONITOR command to add the SEEKS Monitor class for a brief recording interval.

```
MONITOR ENABLE PERFORM USER DASTAP SEEKS INTERVAL 15 SEC -
        SEEKS INCLUDE 283 287 380 382 4C0 TIME FOR 00:03
```

Figure 58. Sample MONITOR Command to Set Monitor Classes. To use the MONITOR command, you must be authorized with CP privilege class A or class E.

For details on how to use the MONITOR command, see the appropriate VM manual:

VM/SP	Up through Release 4	*VM/SP System Programmer's Guide*
VM/SP	Release 5 or later	*VM/SP CP for System Programming*
VM/SP HPO	Up through Release 4.2	*VM/SP HPO System Programmer's Guide*
VM/SP HPO	Release 5 or later	*VM/SP HPO CP for System Programming*

The format of the MONITOR command is described in:

VM/SP	Up through Release 4	*VM/SP Operator's Guide*
VM/SP	Release 5 or later	*VM/SP CP Command Reference*
VM/SP HPO	Up through Release 4.2	*VM/SP HPO Operator's Guide*
VM/SP HPO	Release 5 or later	*VM/SP HPO CP Command Reference*

You should analyze the data collected by the Monitor using the VMMAP or VM/PPF programs. See "Using VMMAP to Review I/O Workload" on page 21 for more information on how to use VMMAP; see "VM Performance Planning Facility (VMPPF)" on page 12 for more information on VM/PPF.

Tuning the Storage Subsystem for Performance

The basic goal of all storage subsystem performance tuning is to improve I/O response time. Using storage controls with cache will improve I/O response time. Other ways to do this are to balance I/O and to reduce I/O activity.

Balancing I/O Load and Reducing Contention

In addition to the availability considerations described in "Configuring for Availability" on page 43, there are several techniques you can use to balance the I/O load and reduce contention in the storage subsystem:

- Balance channel and device activity whenever practical. Some strings can be driven to higher I/O utilization than others, but no one component of the storage subsystem—channel, storage control, or device—should be allowed to limit the capabilities of the other components.
- Spread highly-active system areas and minidisks across volumes. "Placing the S- and Y-Disks: A Scenario" on page 54 describes where to place highly-active data within the configuration.
- Spread temporary disk space across volumes.
- Avoid formatting volumes during prime shift. Formatting ties up channels with steady I/O, limiting the I/O capabilities of other volumes in the string.

Reducing I/O Activity

To reduce the amount of I/O in the storage subsystem, consider the following tuning techniques:

- Use a 4K-byte block size for CMS minidisks. 4K blocks provide more efficient I/O than any other block size.
- Whenever possible, avoid the use of the SYSCLEAR=YES option (which overwrites data with binary zeros). Where security is more important than performance, use SYSCLEAR=YES.
- Use a 4K-byte block for OS-format CMS files (filemode 4) to cause OS simulation to read multiple disk blocks at a time.

Notes for Guest Systems

You can help to optimize the performance of guest systems by taking advantage of facilities that use CP paging in place of guest I/O: facilities such as in-core sorts, Assembler H, and the VM/SP Editor (XEDIT).

Managing Backup and Recovery

Backup is a form of insurance. If things run smoothly, users should never need to return to backup copies of their data. However, if a system component fails unexpectedly, a backup copy may be crucial in restoring full business operations. Regular backup ensures that users will have a current copy of their data available for use even if the original data is somehow damaged or lost.

Because only the owner of the data can determine the backup requirements of a specific minidisk or file, you should use the requirements specified by the users to develop a global backup and recovery strategy for your DASD volumes. Some data is easily re-created and should never need backup; other data is critical to the smooth operation of the system or of your business, and should be backed up frequently.

Generally backup falls into these categories: **minidisk or file backup**, **volume backup**, and **spool file backup**.

Minidisk and file backup are important if a minidisk or file is lost or becomes unusable from logical inconsistencies or media errors. Backup can be done using the VMBACKUP Management System (in VM/SP and VM/SP HPO environments) or the CMS COPYFILE command. "VMBACKUP Management System" on page 12 describes how VMBACKUP can be used to back up changed minidisks and files. "Moving Individual Files to a 3380" on page 89 describes how to use COPYFILE to copy files from one DASD volume to another.

To determine how often minidisks or files should be backed up, consider the following factors:

- The importance of the data
- The rate at which the data changes
- The time and resources needed to re-synchronize or bring the backup copy up to date after recovery
- The relative ease of rebuilding files or minidisks

Volume backup is important if an entire volume fails or where action must be taken against a specific 3380 unit. Because the loss of a single volume may affect many different applications, you should take full volume backups regularly in addition to your incremental minidisk or file backup. For volume backup, you can use the DDR program to dump a full volume of data to DASD or tape. "Backing up a Volume Using DDR" on page 85 describes how to use DDR to back up DASD volumes.

You can also use the CMS TAPE command with the DUMP argument to save the contents of CMS disk files. Once you have saved the files on tape, use the TAPE command with the LOAD argument to restore the files to another device on the system or to transfer them to another system.

For more information on the CMS TAPE command, see the appropriate VM *CMS Reference* manual:

VM/SP	Up through Release 4	*VM/SP CMS Command and Macro Reference*
VM/SP	Release 5 or later	*VM/SP CMS Command Reference*
VM/SP HPO	Up through Release 4	*VM/SP HPO Command and Macro Reference*
VM/SP HPO	Release 5 or later	*VM/SP CMS Command Reference*
VM/XA SF	All releases	*VM/XA SF CMS Command and Macro Reference*

Spool file backup can be done by using SPTAPE to dump the contents of the spool to tape and then restoring the spool file from the tape. This provides you with a copy of the spool file on tape that can be used if the spool is lost. This procedure should be done offshift, on a daily basis. For information on the SPTAPE operator command, see the appropriate manual:

VM/SP	Up through Release 4	*VM/SP Operator's Guide*
VM/SP	Release 5 or later	*VM/SP CP Command Reference*
VM/SP HPO	Up through Release 4.2	*VM/SP HPO Operator's Guide*
VM/SP HPO	Release 5 or later	*VM/SP HPO CP Command Reference*
VM/XA SF	All releases	*VM/XA SF Real System Operation*

Remember to ensure an adequate backup window and enough DASD and tape storage to back up each volume. To determine whether to use DASD or tape to store backup data, consider:

- Cost, both of the media and the personnel needed to operate it
- Speed of recovery necessary
- Amount of data
- Portability of data

Managing DASD Space

In VM, DASD space is a resource that allows little direct management by the end user. End users are assigned a finite number of contiguous cylinders on a specific DASD volume; they cannot "overallocate" their space, and cannot change the amount of space they have without the intervention of the VM administrator. This means that the problems of wasted space, over- or under-allocation, and space fragmentation that characterize other operating systems are rarely significant under VM.

Because space, in the form of minidisks, is permanently assigned to either the system or a user, it is easily identified. When users leave the system, their minidisks can be deleted and the space reused for other minidisks.

You can use the output of the DISKMAP exec to keep track of unallocated space on DASD volumes and to determine when you are running short of DASD space. For an example of DISKMAP output, see Figure 7 on page 18.

Performing Media Maintenance

Media problems with storage devices affect system performance and hardware availability. Your part in media maintenance includes running the System Exception Reports listed in the following section, analyzing them, separating media problems from hardware problems, and handling the media problems. To find out how to obtain these reports, see *Environmental Record Editing and Printing Program User's Guide and Reference.*

Maintaining IBM Storage Subsystem Media explains in detail how each of these reports contributes to the error handling process. It assists you in interpreting the EREP reports and gives you detailed information for using Device Support Facilities (ICKDSF) to handle error situations.

You can fix most media problems with Device Support Facilities. The hardware problems are the responsibility of your service representative. For information on using ICKDSF with 3380s, see *Device Support Facilities: Primer for the User of IBM 3380 Direct Access Storage.* For detailed information on the Device Support Facilities, see *Device Support Facilities User's Guide and Reference.*

System Exception Reports

The Environmental Record Editing and Printing (EREP) program produces a set of System Exception reports, that are based on the contents of the error recording data set (commonly referred to as LOGREC).

- The **System Error Summary (Part 2)** identifies permanent I/O errors, by job and time, associated with data or equipment checks.
- The **Subsystem Exception DASD** report lists accumulated permanent and temporary I/O errors.
- The **DASD Data Transfer Summary** provides details on data checks.

Run the above reports daily as part of normal processing, and designate a member of the data processing center staff to review these reports for actual or potential disk storage problems.

By monitoring your storage devices, you can minimize disruption from problems. When errors occur and the source has been identified, you can alert the responsible party to take action:

- When the source of the error is hardware, call your service representative for service.
- When the source of the error is the volume (media), use Device Support Facilities to handle the specific error condition.

Using the Service Information Messages

With the 3380 Direct Channel Attach Model CJ2, a SIM Alert message is displayed on the operator's console to notify operations personnel that the storage control hardware detected an abnormal condition and has recorded a Service Information Message (SIM) on the Error Recording Data Set (ERDS). It is essential that you run an EREP System Exception Report to get the additional information from the ERDS as quickly as possible. The detailed SIM information includes severity of the abnormal condition and impact of the service action to repair the fault. This information will help you determine when to schedule a service action to repair the fault.

Interpreting the SIM Alert Message

The format for the VM/SP and VM/SP HPO SIM Alert message is shown in Figure 59.

```
DMKDAD403E  0cuu, xxxx,  yyyyyyy, MT=, SER=01aa-dddd, REFCODE=nnnn nnnn nnnn
```

Figure 59. VM/SP and VM/SP HPO SIM Alert Message Format

The format for the VM/XA SF SIM Alert message is shown in Figure 60.

```
HCPERP949E  0cuu, xxxx, yyyyyyy, MT=, SER=01aa-dddd, REFCODE=nnnn nnnn nnnn
```

Figure 60. VM/XA SF SIM Alert Message Format

The fields of the VM SIM Alert messages include:

xxxx is the failing component. SCU specifies a storage control fault.

yyyyyyy is the severity of the failure. The severity can be: ACUTE, SERIOUS, MODERATE, or SERVICE.

MT= is the machine type and model number (7 characters maximum).

SER= is the serial number of the failing unit. The "01" at the beginning of the SER designation indicates that the unit is an IBM-manufactured device.

REFCODE= is a reference code that describes the error and references the list of parts the service representative will need to repair the fault. Thus, the reference code helps the service representative to quickly find the cause of the error and repair the fault.

Figure 61 explains the meaning of the severity field for an SCU failing component.

Severity: SERVICE	Severity: MODERATE	Severity: SERIOUS	Severity: ACUTE
A service-related fault occurred that does not affect storage path operation.	A storage cluster temporary error threshold has been exceeded, but both storage paths are operational.	A permanent error has occurred on one storage path. One storage path remains operational.	A permanent error has occurred on both storage paths in the storage cluster.

Figure 61. Meaning of SIM Alert Severity Field for Failing Component of SCU

Establishing SIM Handling Procedures

It is possible that a SIM alert won't be seen since there is no requirement to respond to the message. The SIM can "roll off" the screen before the operator sees it. However, there are several ways to ensure that you know of all SIM occurrences.

VM/SP and VM/SP HPO contain a Programmable Operator Facility capable of intercepting messages, such as SIMs, and handling them with preprogrammed actions. With this facility you can intercept the SIM, start the execution of the EREP report request with an exec, and route the SIM directly to the storage administrator with a message that the EREP report will be available for examination.

For further information on the Programmable Operator Facility, see the appropriate manual for your system:

VM/SP	Up through Release 4	*VM/SP System Programmer's Guide*
VM/SP	Release 5 or later	*VM/SP CP for System Programming*
VM/SP HPO	Up through Release 4.2	*VM/SP HPO System Programmer's Guide*
VM/SP HPO	Release 5 or later	*VM/SP HPO CP for System Programming*

For instructions on preparing an EREP exec, see *Environmental Record Editing and Printing Program User's Guide and Reference*. For details on the parameters required for selecting the SIM "A3" records, see "Generating EREP Report Requests."

Another way to make sure that you know of all SIM activity is to run and examine EREP reports on a regular basis. To check for SIM records, run either the Detail Edit and Summary report, specifying selection of the "A3" records (as described in "Generating EREP Report Requests") or run the System Exception Reports and review the Informational Messages for SIMs.

Generating EREP Report Requests

You can obtain detail information on SIMs by requesting the Detail Edit and Summary Report. In addition, the System Exception Reports summarize SIM activity in the DASD Informational Messages section.

Detail Edit and Summary Report

The essential parameters for selecting SIM "A3" records are:

```
DEV=(3380)                  Device type of 3380
TYPE=A                      Select A3 records
PRINT=(PT)                  To request Detail Edit and Summary
```

When constructing your EREP report request, use parameters according to your installation's conventions. It is recommended that you ***not*** use DATE and TIME parameters when requesting SIM details. If you do, use a date and time period long enough to ensure all SIM messages relating to the SIM alert are reported. Request details as quickly as possible after the alert message is received at the console.

Sample control statements for generating these reports appear in Figure 62. This example requests SIM Detail Edit and Summary reports from the error log (LOGREC) rather than from a history file. For instructions on constructing report requests, see *Environmental Record Editing and Printing Program User's Guide and Reference*.

```
DEF  STOR  1M                        Specify virtual storage size
EXEC CPEREP STEP1 INPUT A            Contents of file STEP1 INPUT A
                                     are listed here
          TYPE=A                     Select A3
          DEV=(3380)                 Part of A3 selection, device type is 3380
          PRINT=PT                   Requests all the detail print reports
          ACC=N                      Keep the records in LOGREC
```

Figure 62. Sample System Exception Report Request

For each selected "A3" record, SIM sense bytes are formatted and evaluated to provide additional details on the error and its effect on system performance. The 32 SIM sense bytes are printed in hex-character format. A sample SIM record from a Detail Edit and Summary report appears in Figure 63.

```
REPORTING DEVICE:       000844    REPORT: ASYNCHRONOUS                    DAY YEAR
REPORTING DEVICE TYPE:    3380    REPORTING SYSTEM: V370  Rel. n    DATE:  112  87
REPORTING PATH:         20-0844                                         HH MM SS.TH
                                                                  TIME: 23 49 04.76

RECORD TYPE:          DASD SIM
DEVICE DEPENDENT DATA
                      SERVICE INFORMATION MESSAGE
                           SERVICE ALERT  3380-CJ  S/N 12-0010114 REFCODE 26C6-2000-0001
                           PERMANENT ERROR(S) ON 1 OF 2 STORAGE PATHS FOR SSID 0004
                           REPAIR WILL DISABLE STORAGE PATHS 2 AND 1 FOR 0004

HEX DUMP OF RECORD
HEADER    A3831810  00000000  0087112F  23490476  13021170  30810000
    0018  00000000  00000000  00000000  00000000  00000000  00000000  20200844  800F2023
    0038  08000844  D9C1E2F8  F4F40000  00901000  00008FE0  23800020  00000104  22002782
    0058  000426C6  05100212  F1000500
```

Figure 63. Example of a Formatted Detail Edit and Summary Report

If the SIM sense data has irregular information in it, the EREP program does not format the 32 SIM sense bytes into the usual messages for evaluation. Instead, EREP prints the information as hex characters, with byte numbers specified for readability, as shown in Figure 64.

```
REPORTING DEVICE:       0001C4    REPORT: ASYNCHRONOUS                        DAY YEAR
REPORTING DEVICE TYPE:    3380    REPORTING SYSTEM: V370 REL. n        DATE:  117  87
REPORTING PATH:         11-01C4                                            HH MM SS.TH
                                                                     TIME: 17 01 40.95

RECORD TYPE:          DASD SIM
DEVICE DEPENDENT DATA
 DEVICE DEPENDENT DATA NOT FORMATTED
SYSTEM INFORMATION DATA
BYTE 00 01 02 03 04 05 06 07
     D9 C1 E2 F1 F8 F4 00 00

SUBSYSTEM INFORMATION DATA
BYTE 00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15
     00 90 10 00 00 00 8F E0 23 9C 00 10 00 00 0E 04
BYTE 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
     22 00 27 7F F5 00 31 21 05 10 42 11 F4 F4 00 00

HEX DUMP OF RECORD
HEADER    A3831810  00000000  0087117F  17014095  13221170  30810000
    0018  00000000  00000000  00000000  00000000  00000000  00000000  201101C4  800F2021
    0038  080001C4  D9C1E2F1  F8F40000  00901000  00008FE0  239C0010  00000E04  2200277F
    0058  F5003121  05104211  F4F40000
```

Figure 64. Example of an Unformatted Detail Edit and Summary Report

System Exception Reports

It is best to run the System Exception Reports as a normal part of your operation. Summary information of SIM activity appears in the DASD Informational Messages report, as shown in Figure 65.

The COUNT field on this report indicates the number of times a specific SIM is reported. If duplicate SIMs are received, the message is listed once and the COUNT field indicates how many times it occurred during the time period specified by the report request.

For instructions on requesting the System Exception Reports, see *Environmental Record Editing and Printing Program User's Guide and Reference.*

```
DASD INFORMATIONAL MESSAGES                          REPORT DATE 126 87
                                                     PERIOD FROM 112 87
                                                            TO   117 87
PHYSICAL  SYMPTOM
ID         CODE  COUNT            MESSAGE
********************************************************************************
0010111.X  0013     1  SERVICE INFORMATION MESSAGE
                         SERVICE ALERT  3380-CJ  S/N 12-0010111 REFCODE 2600-2000-0013
                         PERMANENT ERROR(S) ON 1 OF 4 STORAGE PATHS FOR SSID F400 AND F500
                         REPAIR WILL DISABLE STORAGE PATHS 2 AND 1 FOR F400 AND F500
0010111.X  0015     1  SERVICE INFORMATION MESSAGE
                         SERVICE ALERT  3380-CJ  S/N 12-0010111 REFCODE 3140-3000-0015
                         PERMANENT ERROR(S) ON 1 OF 4 STORAGE PATHS FOR SSID F500 AND F400
                         REPAIR WILL DISABLE STORAGE PATHS 2 AND 3 FOR F500 AND F400
0010111.X  0017     1  SERVICE INFORMATION MESSAGE
                         SERVICE ALERT  3380-CJ  S/N 12-0010111 REFCODE 3122-2000-0017
                         TEMPORARY ERROR(S) ON 1 OF 4 STORAGE PATHS FOR SSID F400 AND F500
                         REPAIR WILL DISABLE STORAGE PATHS 2 AND 1 FOR F400 AND F500
0010114.X  0001     1  SERVICE INFORMATION MESSAGE
                         SERVICE ALERT  3380-CJ  S/N 12-0010114 REFCODE 26C6-2000-0001
                         PERMANENT ERROR(S) ON 1 OF 2 STORAGE PATHS FOR SSID 0004
                         REPAIR WILL DISABLE STORAGE PATHS 2 AND 1 FOR 0004
```

Figure 65. Example of SIMs in System Exception Informational Messages

Appendix A. Sample DASD Installation and Migration Project Plan

This appendix presents a sample project plan designed for a typical VM DASD installation and migration scenario. The project plan itself documents the following:

- The tasks that need to be completed, including interdependences among tasks
- The implementation schedule
- Actual completion dates for tasks
- Responsible individuals

You can produce adequate planning documents and schedules with simple tools. You can prepare planning information similar to that on the following pages using an editor that has sort-on-column capability. Then a variety of reports can be generated; for example, reports for each person who has responsibility for tasks, lists in planned-completion-date sequence, lists of behind-schedule tasks, and so forth. Useful project plans can also be manually generated and maintained. No matter how you create them, your plans should be written down and understood by all involved.

The sample project plan on the following pages groups tasks by area of activity. Whether you are replacing all of your old DASD with new DASD, or adding a new string, document and review your plans to make your transition to the new hardware as smooth as possible.

SAMPLE DASD INSTALLATION AND MIGRATION PROJECT PLAN

FOCUS AREA ==▶ INSTALLATION PLANNING

TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
		START	END	START	END	
Identify requirements for floor space, power, air conditioning, cables						
Identify prerequisite features						
Verify equipment delivery schedules						
Identify licensed program requirements: order, install and test programs						
Identify and install software tools						
Review vendor and application software support						
Identify and apply software fixes						
Identify publication requirements						
Create project team						
Assign responsibility for project plan						
Schedule system time for volume preparation and data movement						
Plan for personnel training						
Document new operating and programming procedures						
Notes:						

FOCUS AREA ==▶ UNDERSTANDING CURRENT DATA AND HARDWARE CONFIGURATION

TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
		START	END	START	END	
Identify data groups						
Identify performance requirements for data						
Identify availability requirements for data						
Identify space requirements for data						
Identify security requirements for data						
Identify device dependencies						
Measure effectiveness of current hardware configuration						
Notes:						

SAMPLE DASD INSTALLATION AND MIGRATION PROJECT PLAN

FOCUS AREA ==▶ PLANNING HARDWARE CONFIGURATION

TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
		START	END	START	END	
Understand hardware functions and capabilities						
Understand capacity considerations						
Understand performance considerations						
Understand availability considerations						
Understand shared DASD considerations						
Understand guest system considerations						
Plan configuration						
Plan I/O addresses						
Document planned hardware configuration						
Notes:						

FOCUS AREA ==▶ PLANNING DATA CONFIGURATION

TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
		START	END	START	END	
Identify performance-critical data						
Identify candidates for cache volumes						
Plan data movement strategy by volume and/or minidisk						
Plan for backup and recovery during data movement						
Document data configuration and data movement strategy						
Schedule system time for data movement						
Notes:						

SAMPLE DASD INSTALLATION AND MIGRATION PROJECT PLAN

FOCUS AREA ==▶ INSTALLING DASD

TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
		START	END	START	END	
Assign I/O addresses/device numbers to new devices						
Define devices to VM in DMKRIO or HCPRIO generation file						
Define devices to processor with IOCP generation						
Physically install new units						
Initialize volumes with ICKDSF						
Format and allocate CP volumes						
Define CP volumes in SYSOWN list						
Define CMS minidisks						
Format CMS minidisks						
Initialize minidisks for guest systems						
NOTES:						

FOCUS AREA ==▶ PLANNING OPERATIONS CHANGES

TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
		START	END	START	END	
Document changes in operator control panels						
Document power on/off procedures for devices and storage controls						
Document enable/disable procedures for devices and storage controls						
Document procedures to display device status						
Document procedure to vary volumes online/offline						
Document SIM alert procedures						
Document nondisruptive DASD install procedures						
Train operators in new procedures (service and/or remote support)						
Notes:						

SAMPLE DASD INSTALLATION AND MIGRATION PROJECT PLAN

FOCUS AREA ==▶ MOVING DATA

TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
		START	END	START	END	
Identify tools for moving data						
Back up volumes before moving data						
Move full volumes						
Move minidisks and files						
Move spool files						
Move CP data areas and rebuild as necessary						
Notes:						

FOCUS AREA ==▶ MONITORING AND MAINTAINING DEVICES

TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
		START	END	START	END	
Monitor system for SIM alerts						
Set up VM/Monitor to collect performance statistics						
Tune storage subsystem for performance						
Use DISKMAP to track space utilization						
Develop backup and recovery strategy						
Run EREP system exception reports to track DASD errors						
Use ICKDSF to fix media errors						
Notes:						

Appendix B. Migration Aids

The REXX Exec shown in Figure 66 moves CMS minidisks from one volume to another, using DIRMAINT to perform the actual data movement and directory updates. To use this exec:

1. Issue the 'USEDEXT' command to DIRMAINT to generate a list of minidisks to be moved.
2. Edit the returned file to delete any non-CMS minidisks (since the exec will only move CMS minidisks) or minidisks which you do not wish to move.
3. Issue the command 'MOVE5080 fromvol tovol' to start the exec.

The MOVE5080 exec looks for a file named 'fromvol USEDEXT' and moves all the minidisks found in the file. The exec should finish fairly quickly; however, DIRMAINT will continue to format the new minidisks and copy all the files to them.

```
/* MOVE5080 EXEC
   Generates DIRMAINT commands to move CMS minidisks from one pack
   to another.
      FORMAT:      MOVE5080 fromvol tovol
   'fromvol' is the volser of the 3350 to be moved, and 'tovol' is
   the volser of the 3380 to which is to be moved.
   The user must issue the 'DIRMAINT USEDEXT fromvol' command before
   starting this exec to build a list of minidisks on fromvol.    */

arg frompack topack .
if topack = '' then do
  Say 'No target volser entered'
  exit 8
  end
'state' frompack 'usedext a'
if rc ¬= 0 then do
  Say 'Issue DIRM USEDEXT' frompack
  exit rc
  end
'execio * diskr' frompack 'USEDEXT A'
pull line /* trash header record */
Do while queued() > 0
  pull line
  parse var line 1 userid . 10 vaddr . 35 startcyl . 44 size .
  if startcyl = 0 then
    Say 'Minidisk' userid vaddr 'starts at cylinder 0, will not be moved'
    Else  'DIRM CMD' userid vaddr '3380 AUTOV' size topack
  end
```

Figure 66. Minidisk Migration Exec

Note: This exec is presented as an example of what you can do to automate migration procedures. It was written and used for migrating from 3350s to double capacity 3380s. For very different cylinder sizes, adjust the target minidisk size.

Do not use this exec to move a DIRMAINT minidisk. For a procedure you can use to move DIRMAINT minidisks, see *VM/Directory Maintenance Installation and System Administrator's Guide*.

MERGE Xedit Macro

The XEDIT macro shown in Figure 67 will update the text version of the VM directory to move user minidisks to a double capacity 3380. The system programmer will have to install the updated directory with the 'DIRECT' command, or if a directory management program such as DIRMAINT is used, it will have to be reinitialized using the new directory. Full-pack minidisks and other nonstandard directory entries may also have to be updated manually after using the macro.

```
/* Modify VM directory to move two 3380 volumes to one 3380-E volume.
   FORMAT:    MERGE  volid1 volid2 volid3
'volid1' is the volser of the pack to be moved to the front half
(cylinders 0 to 884) of the target 3380-E (volid3)
and 'volid2' is the volser of the pack to be move to
the back half (cylinders 885 to 1769) of the target volume (volid3).
DDR will perform actual data movement (not included in this exec) */

Arg src1 src2 target rest
if target = '' | rest ¬= '' then do
  Say 'Incorrect number of parameters'
  Exit 8
  end
mdcount. = 0
'TOP'
'set wrap off'
'set autosave off'
DO forever
  'LOCATE /' src1 '/ | /' src2 '/'
  IF RC ¬= 0 THEN LEAVE
  STACK 1
  Pull ltype vadd dtype STRTCYL length volser rest
  if ltype = 'MDISK' & (volser = src1 | volser = src2) then do
    if volser = src2 then strtcyl = strtcyl + 885
    'replace' ltype vadd dtype right(strtcyl,4,'') length target rest
    mdcount.volser = mdcount.volser + 1
  end
END
Say mdcount.src1 'minidisks moved from' src1 'to' target
Say mdcount.src2 'minidisks moved from' src2 'to' target
Exit 0
```

Figure 67. VM Directory Update Macro

Note: If DIRMAINT is one of the users being migrated, you cannot use this macro. For a procedure you can use to move DIRMAINT minidisks, see *VM/Directory Maintenance Installation and System Administrator's Guide*.

Glossary

This glossary contains disk storage subsystem terms that are used in the various manuals in the Storage Subsystem Library. Each of the terms included here is not necessarily used in *this specific* manual. To help explain some of the terms related to configuration of storage subsystems, several illustrations are included at the end of this glossary. The definitions of certain terms include references to these illustrations.

A

A-unit. The direct access storage unit that contains the controller functions to attach to the storage control. An A-unit controls the B-units that are attached to it and is often referred to as a head of string.

access mechanism. See actuator.

actuator. A set of access arms and their attached read/write heads, which move as an independent component within a head and disk assembly (HDA). For example, the 3380 Model AK4 has two HDAs, each containing two actuators. See also device and volume.

alternate track. On a direct access storage device, a track designated to contain data in place of a defective primary track.

B

B-unit. A direct access storage unit that attaches to the subsystem through an A-unit. A B-unit has no controller functions.

C

C-unit. A direct channel attach 3380 direct access storage unit that contains both the storage control functions and the DASD controller functions. A 3380 C-unit functions as a head of string and controls the B-units that are attached to it.

cache. A random access electronic storage in selected storage controls used to retain frequently used data for faster access by the channel. For example, 3880 Model 23 and 3990 Model 3 contain cache.

cache fast write. A form of fast write where the data is written directly to cache storage without using nonvolatile storage and is available for later destaging. This 3990 Model 3 Storage Control function should be used for data of a temporary nature, or data which is readily recreated, such as the sort work files created by the appropriate release of DFSORT.

channel interface (CHI). The circuitry of a storage control that attaches storage paths to a host channel.

check-1 error. In the storage control and DASD, an error that does not allow the use of normal machine functions to report details of the error condition.

check-1 error. In the storage control and DASD, an error that can be reported using the normal machine functions.

connection check alert. The electronic signal used by the 3380 to indicate a check-1 error condition to the storage control. See check-1 error.

control interface (CTI). The hardware connection between the storage control function and the DASD controller function.

controller. The hardware component of a DASD head of string unit that provides the path control and data transfer functions. For example, there are two controllers in a 3380 Model AE4 or AK4.

count-key-data (CKD). A DASD data recording format employing self-defining record formats in which each record is represented by a count area, that identifies the record and specifies its format, an optional key area that may be used to identify the data area contents, and a data area that contains the user data for the record. CKD is also used to refer to a set of channel commands that are accepted by a device that employs the CKD recording format. See extended count-key-data architecture.

D

DASD. Direct access storage device; for example, a 3380.

DASD fast write. A form of fast write to cache storage where the data is written concurrently to cache storage and nonvolatile storage and automatically scheduled for destaging to the DASD. Both copies are retained in the storage control until the data is completely written

to the DASD, providing data integrity equivalent to writing directly to the DASD. DASD fast write is available with a 3990 Model 3 Storage Control with nonvolatile storage.

DASD subsystem. One or more DASD strings and the storage control(s) to which the DASD are attached.

demotion. The process of removing the image of one or more records from cache storage. A set of one or more DASD records is demoted either by being selected for replacement (overlay) by another set of DASD records or by being marked invalid. Compare to promotion.

destage. The asynchronous write of new or updated data from cache storage or nonvolatile storage to DASD. This is used only for the fast write and dual copy functions of 3990 Model 3. See also fast write and write hit.

device. A uniquely addressable part of a DASD unit that consists of a set of access arms, the associated disk surfaces, and the electronic circuitry required to locate, read, and write data. See also volume.

device address. Three or four hexadecimal digits that uniquely define a physical I/O device on a channel path in System/370 mode. The one or two leftmost digits are the address of the channel to which the device is attached. The two rightmost digits represent the unit address.

device ID. An 8-bit identifier that uniquely identifies a physical I/O device.

device level selection (DLS). A DASD function available with 3380 Models AD4, BD4, AE4, BE4, AJ4, BJ4, AK4, BK4, and CJ2. With DLS, each of the two controllers in the DASD string has a path to all devices in the string (as many as 14 addresses for a CJ2 or 16 addresses for other string types), and any two devices in the 2-path DASD string can read or write data simultaneously. See DLS support mode, and see Figure 68 on page 123 and Figure 69 on page 124.

device level selection enhanced (DLSE). A DASD function available with 3380 Models AJ4, BJ4, AK4, and BK4. With DLSE, each of the four controllers in the 4-path DASD string (as a result of interconnecting two A-units), has a path to all devices in the string (as many as 32 addresses), and any four devices in the 4-path DASD string can read or write data simultaneously. See DLSE support mode, and see Figure 70 on page 125 and Figure 71 on page 126.

device number. Four hexadecimal digits that logically identify an I/O device in a System 370/Extended Architecture system.

device release. A command that terminates the reservation of the device from the channel issuing the command or from all channels on the interface path group.

device reserve. A command that reserves the device for the channel issuing the command, or for all channels in the same interface path group.

device support facilities program (ICKDSF). A program used to initialize DASD at installation and provide media maintenance.

device support tracks. Reserved tracks of a DASD volume that store defect skipping information. This information is used by the subsystem (for example, at IML time) and by host utility programs such as ICKDSF.

diagnostic tracks. Tracks used by the diagnostic programs for testing the read/write function.

diskette drive. A direct access storage device that uses diskettes as the storage medium. A 3880 uses a read-only diskette drive for microcode storage; a 3990 and a 3380 Model CJ2 use a read/write diskette drive for microcode storage and storage control error logs error logs.

DLS support mode. A mode of operation in a 3990 Storage Control that supports 3380 2-path strings, including 3380 AA4 strings and 3380 AD4, AE4, AJ4, and AK4 2-path strings. DLS support mode must be specified by the IBM service representative at installation for the 3990. See single-path storage director, and see Figure 69 on page 124.

DLSE support mode. A mode of operation in a 3990 Model 2 or 3 Storage Control that supports 3380 AJ4 and AK4 4-path strings. DLSE support mode must be specified by the IBM service representative at installation time for the 3990. See multipath storage director, and see Figure 70 on page 125 and Figure 71 on page 126.

dual copy. A high availability function made possible by nonvolatile storage in a 3990 Model 3. Dual copy maintains two identical copies of designated DASD volumes in the logical 3990 Model 3 subsystem, and automatically updates both copies every time a write operation is issued to the dual-copy logical volume.

dual-copy logical volume. A logical volume comprised of two physical devices with all data recorded twice, once on each device. A 3990 Model 3 Storage Control with nonvolatile storage automatically ensures that both devices are updated with each write operation to the dual-copy volume. Also called a duplex pair.

dual-frame configuration. Consists of two like storage controls physically interconnected. Pairs of 3880 Model 13 and 23 and 3990 Model 2 and 3 Storage Controls can be dual-framed. In a dual-frame configuration, each storage director in a logical DASD subsystem is in a

different storage control. When a 3990 Storage Control is in DLS support mode, each DASD string has one path to a single-path storage director in each of the 3990 Storage Controls. When a 3990 Storage Control is in DLSE support mode, each DASD string has two paths to a multipath storage director in each of the 3990 Storage Controls.

duplex mode. Two devices in a 3990 Model 3 subsystem are in duplex mode when they have been made into a dual-copy logical volume.

duplex pair. See dual-copy logical volume.

dynamic path reconnect. A function of dynamic path selection (DPS) that allows disconnected DASD operations to reconnect over any available channel path rather than being limited to the one on which the I/O operation was started. It is available only on System 370/Extended Architecture systems. For example, when a 3990 Storage Control (in DLSE support mode) having four host channels is connected to a 3380 Model AJ4 or AK4 4-path string, any device can reconnect on any one of four completely independent data paths, providing improved performance and availability.

dynamic path selection (DPS). DASD subsystem functions available with all 3380 heads of string except Model A04. These functions include:

- Two controllers providing data paths from the 3380 strings to the storage directors
- Simultaneous transfer of data over two paths to two devices, providing the two devices are on separate internal paths within the string
- Sharing DASD volumes by using System-Related Reserve and Release
- Providing dynamic path reconnect to the first available path (with System 370/Extended Architecture hosts only)

E

error burst. A sequence of bit errors counted as one unit, or burst.

error correcting code (ECC). A code designed to detect and correct error bursts by the use of check bytes.

extended count-key-data (ECKD) architecture. A set of channel commands that use the CKD track format. This architecture employs the Define Extent and Locate Record commands to describe the nature and scope of a data transfer operation to the storage control to optimize the data transfer operation. The 3990 Storage Control supports the ECKD architecture.

F

fast write. In a 3990 Model 3 Storage Control, a write operation at cache speed that does not require immediate transfer of data to a DASD. The data is written directly to cache storage and/or nonvolatile storage and is available for later destaging. Fast write reduces the time an application must wait for the I/O operation to complete. See also DASD fast write, cache fast write, and destage.

fence. To separate one or more paths or elements from the remainder of the logical DASD subsystem. The separation is by logical boundaries rather than power boundaries. This separation allows isolation of failing components so that they do not affect customer operation.

G

gigabyte (GB). 10^9 bytes.

H

head and disk assembly (HDA). A field replaceable unit in a direct access storage device containing the disks and actuators. A 3380 Model AK4 has two HDAs.

head of string. The unit in a DASD string that contains controller functions. For example, a 3380 Model AE4, AK4, or CJ2.

home address (HA). The first field on a CKD track that identifies the track and defines its operational status. The home address is written after the index point on each track.

I

ICKDSF. See Device Support Facilities program.

IDCAMS. A Data Facility Product program for MVS that is also referred to as access method services.

identifier (ID). A sequence of bits or characters that identifies a program, device, controller or system.

index point. The reference point on a disk surface that determines the start of a track.

initial microcode load (IML). The act of loading microcode.

I/O device. An addressable input/output unit, such as a direct access storage device, magnetic tape device, or printer.

invalidation. The process of removing records from cache storage because of a change in status of a subsystem facility or function, or because of an error while processing the cache image of the set of records. When such a cache image is invalidated, the corresponding records cannot be accessed in cache and the assigned cache space is available for allocation.

L

least recently used algorithm (LRU). The algorithm used to identify and make available the cache space that contains the least recently used data.

logical DASD subsystem. Two storage directors attached to the same DASD strings together with those DASD strings. See Figure 68 on page 123, Figure 70 on page 125, and Figure 71 on page 126.

M

maintenance analysis procedure (MAP). A step-by-step procedure for tracing a symptom to the cause of a failure.

megabyte (Mb). 10^6 bytes.

multipath storage director. A storage director in a 3990 Storage Control operating in DLSE support mode. Each multipath storage director in a storage control is associated with two storage paths. All storage paths in a multipath storage director respond to the same range of control unit addresses on a channel. See Figure 70 on page 125 and Figure 71 on page 126.

N

nondisruptive install. Provides for the physical installation of additional Enhanced Subsystem B-units to an existing 4-path DASD string or an additional 4-path DASD string, concurrently with customer operations, providing access to existing data when DASD unit installation activity is occurring. Nondisruptive install is available when only 4-path Enhanced Subsystem DASD are attached to a 3990 Model 2 or 3 Storage Control.

nonvolatile storage (NVS). Additional random access electronic storage with a backup battery power source, available with a 3990 Model 3 Storage Control, used to retain data during a power failure. Nonvolatile storage, accessible from all storage directors, stores data during DASD fast write and dual-copy operations.

O

orientation. A control state within a storage path that indicates the type of area (home address, count, key, or data field) that has just passed under the read/write head of the device.

P

physical ID. A unique designation to identify specific components in a data processing complex. Pinned data exists only when using fast write functions.

predictable write. A fast write operation that formats, in cache storage only, the entire user area of the track and creates a track image. This full-track image is available for later destaging to a DASD.

primary device. One device of a dual-copy volume. All channel commands to the dual-copy logical volume are directed to the primary device. The data on the primary device is duplicated on the secondary device. See also secondary device.

primary track. On a direct access storage device, the original track on which data is stored. See also alternate track.

promotion. The process of moving a track image from a DASD to cache.

R

read hit. When data requested by the read operation are in the cache.

read miss. When data requested by the read operation are not in the cache.

rotational position sensing (RPS). A function that permits a DASD to reconnect to a block multiplexer channel when a specified sector has been reached. This allows the channel to service other devices on the channel during positional delay.

S

secondary device. One of the devices in a dual-copy logical volume that contains a duplicate of the data on the primary device. Unlike the primary device, a limited subset of channel commands may be directed to the secondary device. See also primary device.

service information message (SIM). A message, generated by the host processor upon receipt of sense information from a 3990 or a 3380 Model CJ2, that contains notification of a need for repair or customer action. The SIM identifies the affected area of the storage control and the effect of the expected service action. A host Error Recovery Procedure (ERP) causes a SIM Alert to be sent to the operator console.

SIM Alert. An operator console message that alerts the operator that an action requiring attention has occurred. The service information message (SIM) can be obtained from the EREP exception report.

simplex mode. A volume is in simplex mode if it is not part of a dual-copy logical volume. Terminating a dual-copy logical volume returns the two devices to simplex mode. In this case, there is no longer any capability for either automatic updates of the secondary or for logging changes.

single-frame configuration. In a single-frame configuration, the storage directors of a logical DASD subsystem are located inside one storage control.

single-path storage director. A storage director in a 3990 or 3380 Model CJ2 operating in DLS support mode. Each single-path storage director in the storage cluster is associated with one storage path. A storage path on a single-path storage director responds to a unique control unit address on the channel. A single-path storage director in a 3990 is like a storage director in a 3880. See Figure 69 on page 124.

storage cluster. In the 3990 Storage Control and 3380 Model CJ2, a power and service region containing two independent transfer paths and either one multipath storage director or two single-path storage directors. It is designed so that should a failure or maintenance action occur, it will be independent of the other storage cluster in a 3990 Model 2 or 3 Storage Control. The 3990 Model 1 and the 3380 Model CJ2 each have a single storage cluster; the 3990 Model 2 and 3 each have two storage clusters. See also storage director, single-path storage director, and multipath storage director.

storage control. The component in a DASD subsystem that connects the DASD to the host channels. It performs channel commands and controls the DASD devices. For example, the 3990 Model 2 and 3 are storage controls.

storage director. In a 3990 storage control, a logical entity consisting of one or more physical storage paths in the same storage cluster. In a 3880, a storage director is equivalent to a storage path. See also storage path, single-path storage director, and multipath storage director.

storage director ID. For 3880 Storage Control configurations, an 8-bit designation that uniquely identifies the storage director regardless of its selection address. It identifies to the service representative, by means of EREP, a failing subsystem component (storage director) without having to translate a selection address (which may have little relation to a physical address) to a physical component. The storage director ID is the number shown on the operator panels of 3880s and the attached DASD units.

storage facility. See 4-path string.

storage path. The hardware within the 3990 Storage Control that transfers data between the DASD and a channel. See also storage director.

storage subsystem. One or more storage controls and their attached storage devices.

string. A series of connected DASD units sharing one or more controllers (or heads of string). For example, a 3380 Model AE4 with the attached B-units is one string.

string address. The 1-bit address used by the storage control to direct commands to the correct DASD string on the CTL-I.

string ID. An 8-bit identifier that uniquely identifies the physical string regardless of the selection address. It identifies to the service representative, by means of EREP, a failing subsystem component (controller, device) without having to translate a selection address (which may have little relation to a physical address) to a physical component. The string ID is the number shown on the operator panel.

substring. In a 4-path Enhanced Subsystem DASD configuration, one of the two A-units and the physically adjacent B-units (as many as three B-units).

subsystem identifier (SSID). In a 3990 Storage Control configuration, a number that identifies the physical components of a logical DASD subsystem. This number is set by the service representative at time of installation, and is included in the vital product data in the support facility. This number is identified on the 3380 Enhanced Subsystem models and 3990 operator panels.

subsystem storage. A term used for cache storage in a 3880 Model 13 or 23. See cache storage.

support facility (SF). A component of each 3990 and 3380 Model C2 storage cluster that provides initial microcode load, error logging, maintenance panel, MAPs, and microdiagnostic functions for that cluster.

U

unit address. The last two hexadecimal digits of a DAS device address. This identifies the storage control and DAS string, controller, and device to the channel subsystem. Often used interchangeably with channel unit address and device address in System/370 mode.

V

volume. The DASD space accessible by a single actuator. A 3380 Model AK4 contains four volumes, each with 1.89 gigabytes of space.

W

write hit. When data requested by the write operation are in the cache.

write miss. When data requested by the write operation are not in the cache.

Numeric

2-path string. A series of physically connected DASD units in which the head of string unit provides two data transfer paths that can operate simultaneously.

4-path string. A series of physically connected DASD units in which the heads of string provide four data transfer paths that can operate simultaneously. A 4-path string requires two 3380 Enhanced Subsystem model A-units.

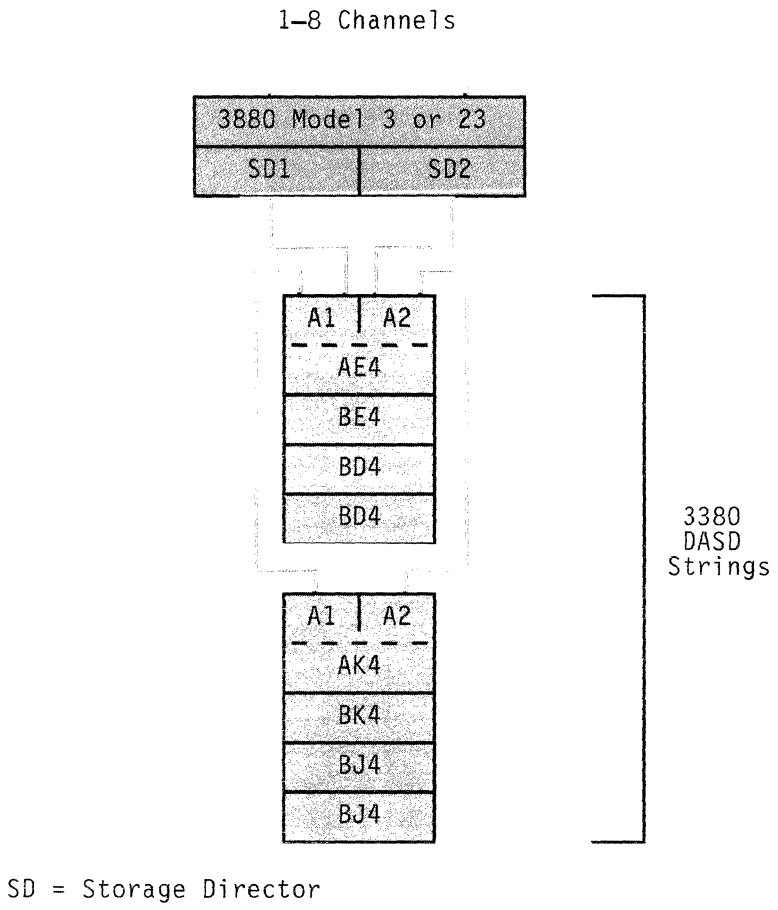

Figure 68. Example of 3380 Model AE4 and AK4 2-Path Strings Attached to a 3880. Two 3380 strings sequentially connected to both storage directors of a 3880 Model 3 or 23 with appropriate MESs installed. In this example, upper string contains a 3380 Model AE4 controller and a mixture of BD4 and BE4 units. The lower string contains a Model AK4 controller, and a mixture of BJ4 and BK4 units. In a 3880, a storage director performs the same functions as a storage path.

This example shows one logical DASD subsystem.

Definitions of the following glossary terms refer to this illustration:

- DLS
- Logical DASD subsystem

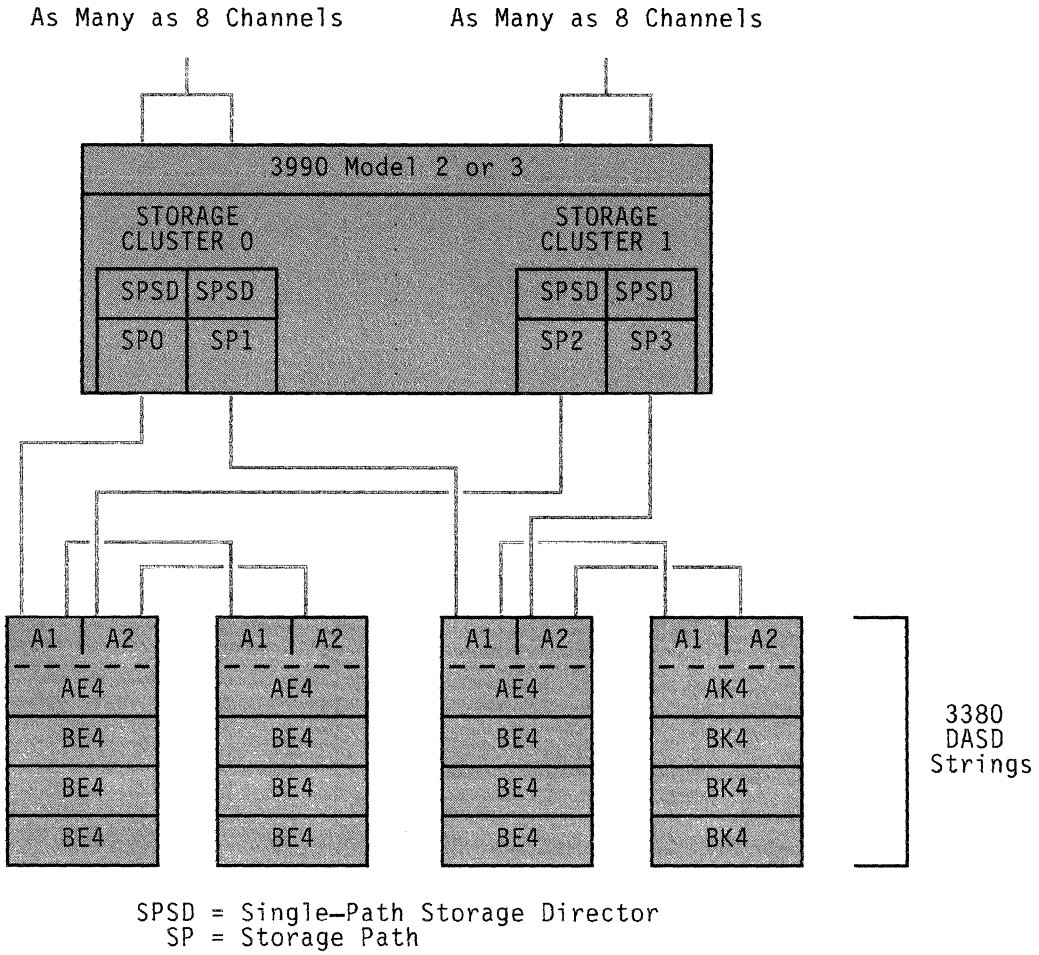

Figure 69. Example of 3380 2-Path Strings Attached to a 3990 Model 2 or 3 in DLS Support Mode. This example shows two logical DASD subsystems. Storage Paths 0 and 2 and the attached DASD comprise one subsystem, with a unique subsystem ID, while Storage Paths 1 and 3 and the attached DASD comprise a second subsystem with a unique subsystem ID. Compare this 2-path Enhanced Subsystem string with 4-path Enhanced Subsystem strings shown in Figure 70.

Definitions of the following glossary terms refer to this illustration:

- DLS
- DLS support mode
- Single-path storage director

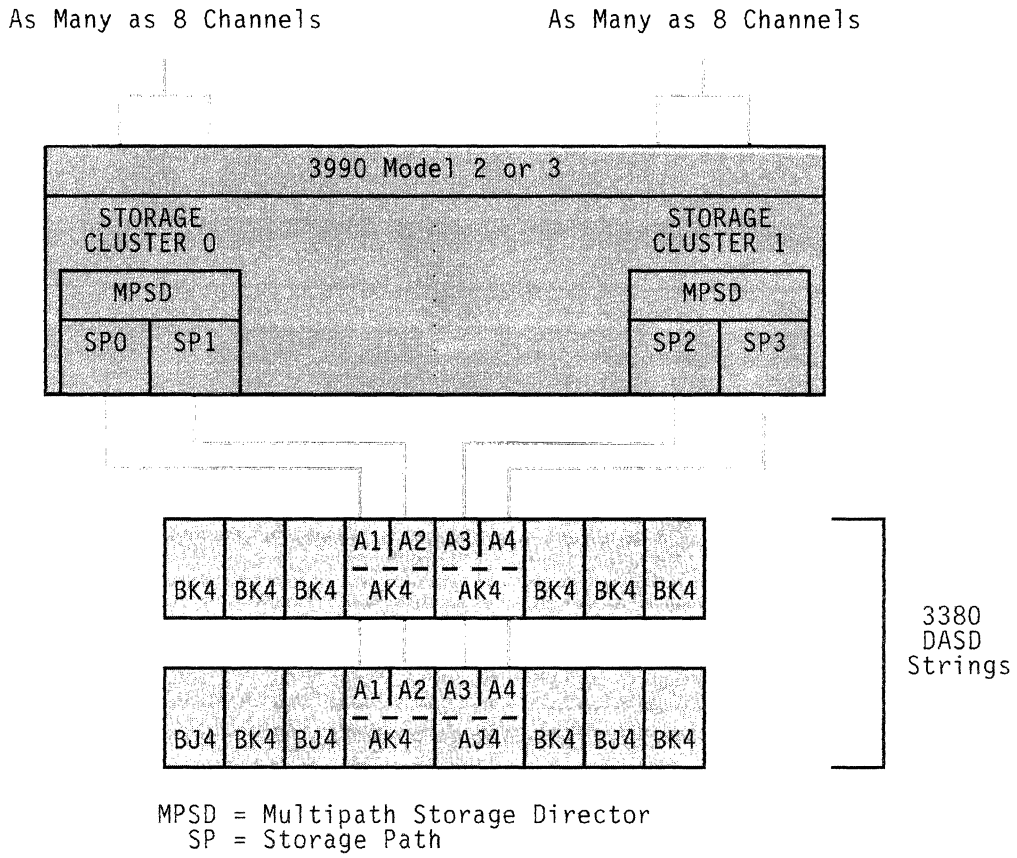

Figure 70. Example of 3380 Enhanced Subsystem Model 4-Path Strings Attached to a 3990 in DLSE Support Mode. Two 3380 4-path strings sequentially connected to the multipath storage directors in the same 3990 Model 2 or 3.

This example shows one logical DASD subsystem.

Definitions of the following glossary terms refer to this illustration:

- DLSE
- DLSE support mode
- Logical DASD subsystem
- Multipath storage director

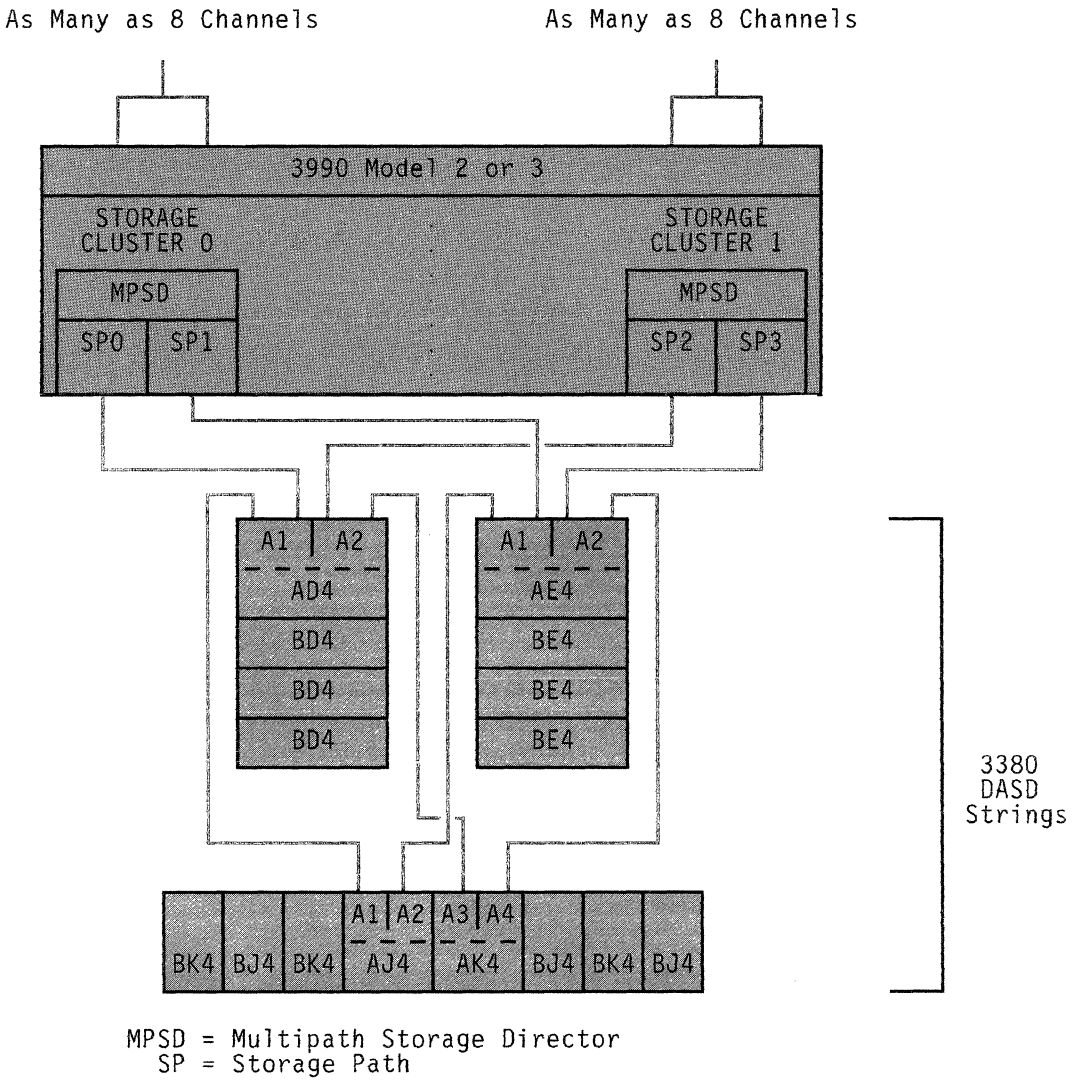

Figure 71. Example of 3380 2-Path and 4-Path Strings Attached to a 3990 in DLSE Support Mode. Two 3380 2-path strings and one 4-path string, sequentially connected to the multipath storage directors in the same 3990 Model 2 or 3.

This example shows one logical DASD subsystem.

Definitions of the following glossary terms refer to this illustration:

- DLSE
- DLSE support mode
- Logical DASD subsystem
- Multipath storage director

Bibliography

The manuals listed in the table below contain more detailed information on the subjects discussed in this book. For each manual referenced, the table shows the short and expanded title with the manual's order number, and a short description of its contents.

For information on how to order these manuals, contact your local IBM branch office.

Short Title	Full Title	Order Number	Contents
Storage Subsystem Library	**3380 Manuals**		
IBM 3380 Direct Access Storage Introduction	*IBM 3380 Direct Access Storage Introduction*	GC26-4491	Overview of 3380 functions
IBM 3380 Direct Access Storage Reference Summary	*IBM 3380 Direct Access Storage Reference Summary*	GX26-1678	Summary of 3380 device characteristics
Maintaining IBM Storage Subsystem Media	*Maintaining IBM Storage Subsystem Media*	GC26-4495	Description of DASD media maintenance and error handling
Storage Subsystem Library Master Index	*Storage Subsystem Library Master Index*	GC26-4496	Index to information in 3380 and 3990 manuals
Using the IBM 3380 Direct Access Storage in an MVS Environment	*Using the IBM 3380 Direct Access Storage in an MVS Environment*	GC26-4492	Discussion of 3380 use in MVS/XA and MVS/370
Using the IBM 3380 Direct Access Storage in a VSE Environment	*Using the IBM 3380 Direct Access Storage in a VSE Environment*	GC26-4494	Discussion of 3380 use in VSE
Storage Subsystem Library	**3990 Storage Control Manuals**		
IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide	*IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*	GA32-0100	Detailed information on installation and use of the 3990 storage control
IBM 3990 Storage Control Introduction	*IBM 3990 Storage Control Introduction*	GA32-0098	Overview of 3990 storage control functions
IBM 3990 Storage Control: Reference	*IBM 3990 Storage Control: Reference*	GA32-0099	Reference information on 3990 functions including descriptions of DASD channel command set
3880 Storage Control Information			
IBM 3880 Storage Control Models 1, 2, 3, and 4 Description Manual	*IBM 3880 Storage Control Models 1, 2, 3, and 4 Description Manual*	GA26-1661	Overview of 3880 Models 1, 2, 3, and 4 functions
Introduction to IBM 3880 Storage Control Model 13	*Introduction to IBM 3880 Storage Control Model 13*	GA32-0062	Overview of 3880 Model 13 functions
IBM 3880 Storage Control Model 23 Introduction	*IBM 3880 Storage Control Model 23 Introduction*	GA32-0082	Overview of 3880 Model 23 functions
VM Reference Information			
Alternate Pathing under VM	*Alternate Pathing under VM*	GG22-9381	Description of how VM/SP uses alternate paths in the storage subsystem
CMS Reference	*VM/SP CMS Command and Macro Reference* (Release 4)	SC19-6209	Discussion of CMS commands, including FORMAT and COPYFILE for the VM/SP and VM/SP HPO environments
	VM/SP CMS Command Reference (Release 5)		
	VM/XA SF CMS Command and Macro Reference	GC19-6231	Discussion of CMS commands, including FORMAT and COPYFILE
Comparison of IBM 3380s and IBM 3350s Used for VM/CMS Minidisks	*Comparison of IBM 3380s and IBM 3350s Used for VM/CMS Minidisks.*	GG22-9347	Evaluates 3380 disk capacity and performance in VM/SP HPO
Component Failure Impact Analysis—An Availability Management Technique	*Component Failure Impact Analysis—An Availability Management Technique*	GC20-1865	Planning for hardware availability through configuration
CP for System Programming	*VM/SP CP for System Programming* (Release 5)	SC24-5285	Discussion of system programming tasks and commands, including CP INDICATE, SYSOWN, MONITOR

Short Title	Full Title	Order Number	Contents
	VM/SP HPO CP for System Programming (Release 5)	SC19-6224	Discussion of system programming tasks and commands, including CP INDICATE, SYSOWN, MONITOR
CP Reference	*VM/SP CP Command Reference*	SC19-6211	Discussion of CP commands for both general and non-general users
	VM/SP HPO CP Command Reference	SC19-6227	Discussion of CP commands for both general and non-general users
	VM/XA SF CP Command and Diagnosis Reference	GC19-6215	Discussion of CP commands for both general and non-general users
DASD Sharing under VM	*DASD Sharing under VM*	GG22-9380	Description of how DASD can be shared among guest systems
VM/SP AJ4,BJ4,AK4,BK4 User's Guide	*VM/SP IBM 3380 AJ4/BJ4/AK4/BK4 User's Guide*	GC24-5371	Contains general description of VM/SP software support (within CP and CMS) for the Enhanced Subsystem 3380s.
3380 VM Performance Analysis	*IBM 3380 Extended Capability DASD VM Benchmark Performance Analysis*	GG66-0262	Helps plan migration from 3380 standard models to 3380 extended capability models, providing migration execs
Storage Hardware Maintenance			
EREP User Guide and Reference	*Environmental Record Editing and Printing Program User's Guide and Reference*	GC28-1378	Description of EREP functions and commands for DASD media reporting
ICKDSF Primer	*Device Support Facilities: Primer for the User of IBM 3380 Direct Access Storage*	GC26-4498	Describes how to use ICKDSF with the 3380
ICKDSF User's Guide and Reference	*Device Support Facilities User's Guide and Reference*	GC35-0033	Description of ICKDSF functions and commands for DASD initialization and maintenance
Physical Planning and Installation Information			
IBM Input/Output Equipment: Installation – Physical Planning for System/360, System/370, and 4300 Processors	*IBM Input/Output Equipment: Installation – Physical Planning for System/360, System/370, and 4300 Processors*	GC22-7064	Description of physical planning for I/O hardware
9370 Information System Installation Manual – Physical Planning	*9370 Information System Installation Manual – Physical Planning*	GA24-4031	Contains physical planning information for the 9370 family of processors
Installation	*VM/SP Installation Guide*	SC24-5237	Discussion of VM/SP installation tools, including the DISKMAP exec
	VM/SP HPO Installation Guide	SC38-0107	Discussion of VM/SP HPO installation tools, including the DISKMAP exec
	VM/XA SF Installation, Administration, and Service	GC19-6217	Discussion of VM/XA SF installation tools, including the DISKMAP exec
IOCP User's Guide and Reference	*Input/Output Configuration Program User's Guide and Reference*	GC28-1027	Shows how to define the I/O configuration data required by the processor complex to control I/O requests, describing the MVS version, the VM version, and the standalone version of IOCP
System Operations and Planning Information			
Operations	*VM/SP Operator's Guide*	SC19-6202	Discussion of VM/SP operator commands and facilities, including DDR, MONITOR, VARY, and DMKFMT
	VM/SP HPO Operator's Guide (Release 4.2)	ST00-1898	Discussion of VM/SP HPO operator commands and facilities, including DDR, MONITOR, VARY, and DMKFMT
	VM/SP HPO CP Command Reference (Release 5)	SC19-6227	Discussion of VM/SP HPO operator commands and facilities.
	VM/XA SF Real System Operation	GC23-0139	Discussion of VM/XA SF operator commands and facilities, including DDR, MONITOR, VARY, and DMKFMT
	VM/XA SF Virtual Machine Operation	GC23-0138	Discussion of VM/XA SF operator commands and facilities

Short Title	Full Title	Order Number	Contents
Planning	*VM/SP Planning Guide and Reference*	SC19-6201	Discussion of VM/SP hardware and software planning, system design, and system definition
	VM/SP HPO Planning Guide and Reference	SC19-6223	Discussion of VM/SP HPO hardware and software planning, system design, and system definition
	VM/XA SF Virtual Machine Planning	SC23-0167	Discussion of VM/XA SF hardware and software planning, system design, and system definition
	VM/XA Systems Facility Planning Guide	GG24-1709	Provides advanced installation and planning information for VM/XA SF, and includes planning information on MVS/370, MVS/XA, VSE and VM guests under VM/XA SF
Data Security Information			
RACF General Information	*Resource Access Control Facility General Information*	GC28-0722	Overview of RACF/VM functions for data security
Performance and Monitor Information			
VM/RTM Program Description/Operations Manual	*VM Real Time Monitor Program Description/Operations Manual*	SH20-2337	Description of RTM functions, commands, and reports for performance monitoring
RTM/SF Program Description/Operations Manual	*VM/XA Realtime Monitor/Systems Facility Program Description/Operations Manual*	SH26-7000	Description of RTM/SF functions, commands, and reports for performance monitoring
System Facilities for Programming	*Virtual Machine System Facilities for Programming* (Release 5)	ST24-5288	Discussion of system programming tasks and commands, including DDR for both the VM/SP and VM/SP HPO environments
System Logic and Problem Determination Guide Volume 2—CMS	*VM/SP System Logic and Problem Determination Guide Volume 2—CMS*	LY20-0893	Discussion of CMS file storage logic in the VM/SP and VM/SP HPO environments
System Programmer's Guide	*VM/SP System Programmer's Guide* (Release 4)	SC19-6203	Discussion of system programming tasks and commands, including CP INDICATE, SYSOWN, MONITOR
	VM/SP HPO System Programmer's Guide (Release 4.2)	ST00-1897	Discussion of system programming tasks and commands, including CP INDICATE, SYSOWN, MONITOR
VM Running Guest Operating Systems	*VM Running Guest Operating Systems*	GC19-6212	Discussion of guest operating systems under VM
	VM/XA SF Virtual Machine Planning	SC23-0167	Discussion of VM/XA SF hardware and software planning, system design, and system definition
	VM/XA SF Virtual Machine Operation	GC23-0138	Discussion of VM/XA SF operator commands and facilities, including VARY
VMBACKUP Management System General Information	*VMBACKUP Management System General Information*	GH20-6248	Overview of VMBACKUP and VMArchive functions for data backup and archival
VM/Directory Maintenance General Information	*VM/Directory Maintenance General Information*	GC20-1836	Overview of DIRMAINT functions for directory maintenance
VM/Directory Maintenance Installation and System Administrator's Guide	*VM/Directory Maintenance Installation and System Administrator's Guide*	SC20-1840	Description of DIRMAINT administrator commands
VM/ISF General Information	VM/Intersystems Facility General Information	GC23-0397	Provides an overview of using VM/ISF for sharing minidisks
VM/ISF Planning and Installation	VM/Intersystems Facility Planning and Installation	SC23-0399	Provides guidance and reference information for those planning system resources' usage for VM/ISF and for installers of VM/ISF
VM/ISF Operation and Use	VM/Intersystems Facility Operation and Use	SC23-0400	Describes operation of VM/ISF functions, and contains CP and CMS command syntax and messages and codes for CP, CMS, and VM/Pass-Through Facility
VMMAP General Information	*VM Monitor Analysis Program General Information*	GC34-2164	Overview of VMMAP functions for performance monitoring
VMMAP User's Guide and Reference	*Virtual Machine Monitor Analysis Program User's Guide and Reference*	SC34-2166	Description of VMMAP commands and reports

Short Title	Full Title	Order Number	Contents
VM/PPF General Information	*VM Performance Planning Facility General Information*	GC34-2126	Overview of VM/PPF functions for performance analysis and modeling
VM/Integrated System Information			
VM/IS Planning For Your System	*VM/Integrated System Planning For Your System*	SC24-5337	Provides pre-installation planning instructions for VM/IS and should be read before *VM/IS Installing Your System*
VM/IS Installing Your System	*VM/Integrated System Installing Your System*	SC24-5341	Provides step-by-step instructions for installing VM/IS; use this book in conjunction with *VM/IS Planning For Your System*
VM/IS Managing Your System	*VM/Integrated System Managing Your System*	SC24-5338	Provides operation and administration instructions for VM/Integrated System
VM/IS Reporting System Problems	*VM/Integrated System Reporting System Problems*	SC24-5339	Provides problem reporting instructions for VM/IS and is based on VM/Interactive Productivity Facility ' s Problem Control Facility
VM/IS Learning to Use Your System: Error and Information Messages	*VM/Integrated System Learning to Use Your System: Error and Information Messages*	SC24-5351	Describes error and information messages produced by VM/IS; includes cross-reference table to help locate messages produced by other products and functions included in VM/IS

Index

W

Numerics

Reader's Comment Form

Using the IBM 3380 Direct Access Storage
in a VM Environment
GC26-4493-0

This manual is part of a library that serves as a reference source for system analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Your comments will be sent to the author's department for whatever review and action, if any, are deemed appropriate.

Note: Do not use this form to request IBM publications. If you do, your order will be delayed because publications are not stocked at the address printed on the reverse side. Instead, you should direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.

If you have applied any technical newsletters (TNLs) to this book, please list them here: _____

Chapter/Section _____

_____ Page No. _____

Comments:

Note: Staples can cause problems with automated mail sorting equipment. Please use pressure sensitive or other gummed tape to seal this form.

If you want a reply, please complete the following information.

Name _____ Phone No. (____)_____

Company _____

Address _____

_____

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

Reader's Comment Form

Fold and tape | Please do not staple | Fold and tape

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
P.O. Box 50020
Programming Publishing
San Jose, California 95150

Fold and tape | Please do not staple | Fold and tape

IBM®

File Number
S370-07, 4300-07

GC26-4493-00

Printed in U.S.A.