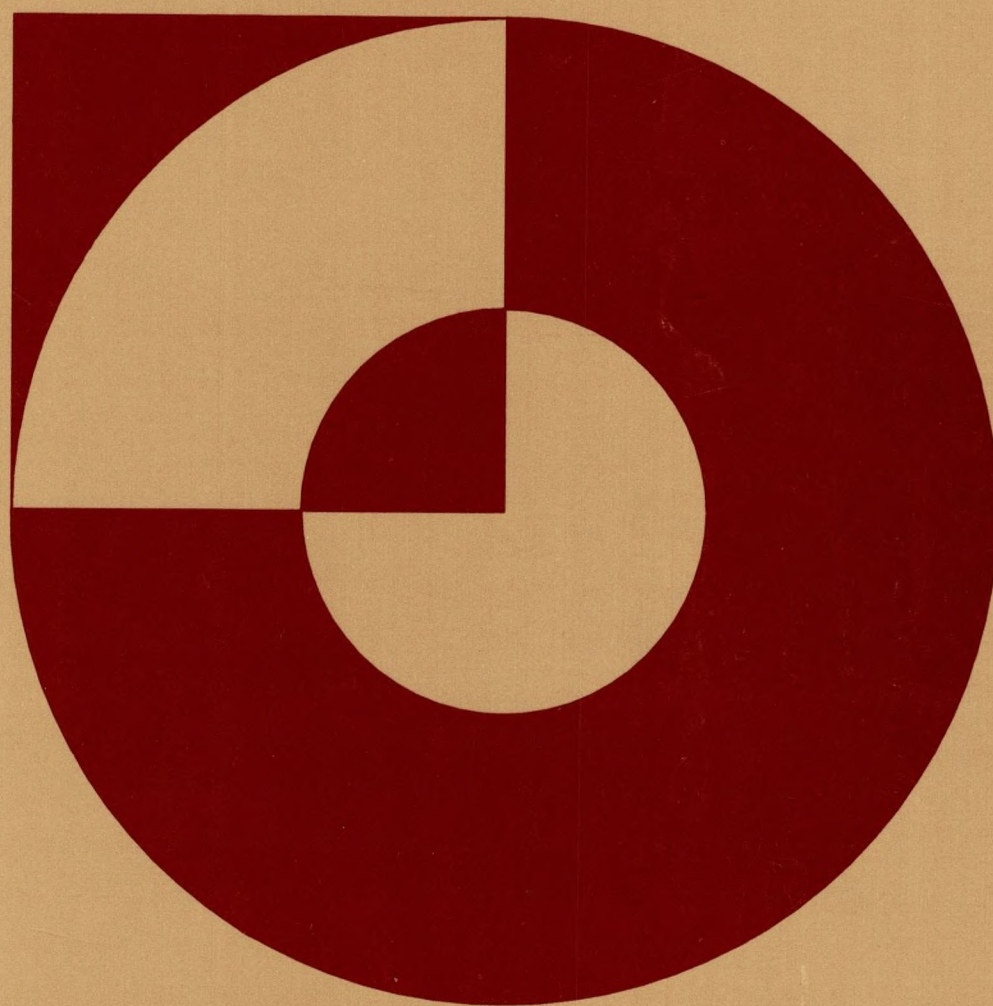




Storage Subsystem Library

GA32-0098-1

IBM 3990 Storage Control Introduction





Storage Subsystem Library

GA32-0098-1

IBM 3990 Storage Control Introduction

Second Edition (June 1990)

This is the second edition of the *IBM 3990 Storage Control Introduction*, a member of the Storage Subsystem Library. This edition replaces and makes obsolete the previous edition, GA32-0098-0. The changes for this edition are summarized under "Summary of Changes". Consult the latest edition of the applicable IBM system bibliography for current information on this product.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address given below. A form for readers' comments appears at the back of this publication. If the form has been removed, address your comments to IBM Corporation, General Products Division, Department 61C, Tucson, Arizona, U.S.A. 85744.

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1987, 1990. All rights reserved.

Note to US Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Federal Communications Commission (FCC) Statement

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

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Canadian Department of Communications compliance statement: This equipment does not exceed Class A limits per radio noise emissions for digital apparatus, set out in the Radio Interference Regulation of the Canadian Department of Communications. Operation in a residential area may cause unacceptable interference to radio and TV reception requiring the owner or operator to take whatever steps are necessary to correct the interference.

Avis de conformité aux normes du ministère des Communications du Canada: Cet équipement ne dépasse pas les limites de Classe A d'émission de bruits radioélectriques pour les appareils numériques, telles que prescrites par le Règlement sur le brouillage radioélectrique établi par le ministère des Communications du Canada. L'exploitation faite en milieu résidentiel peut entraîner le brouillage des réceptions radio et télé, ce qui obligerait le propriétaire ou l'opérateur à prendre les dispositions nécessaires pour en éliminer les causes.

The United Kingdom Telecommunications Act 1984: This apparatus is approved under the approval No. NS/G/23/J/100003 for the indirect connections to the public telecommunications systems in the United Kingdom.

この装置は、第一種情報装置（商工業地域において使用されるべき情報装置）で商工業地域での電波障害防止を目的とした情報処理装置等電波障害自主規制協議会（VCCI）基準に適合しております。

従って、住宅地域またはその隣接した地域で使用する、ラジオ、テレビジョン受信機等に受信障害を与えることがあります。

取扱説明書に従って正しい取り扱いをして下さい。

100A

Contents

Special Notices	1
Trademarks	1
Preface	3
About This Book	3
Terminology	4
The Storage Subsystem Library	4
Storage Control Publications	4
3390 Publications	5
3380 Publications	5
Shared SSL Publications	6
Storage Subsystem Library Ordering Information	8
3390 and 3990 Publications	8
3380 and 3990 Publications	9
Storage Subsystem Library Binders	10
Related Publications	10
Summary of Changes	11
Second Edition, June 1990	11
Chapter 1. 3990 Overview	13
3990 Standards of Data Availability and Performance	14
3990 Standards of Reliability, Availability, and Serviceability	14
3990 Performance Features	15
Faster and More Efficient Microprocessing	15
Four Storage Paths	16
Cache (3990 Model 3 Only)	16
Fast Write Capabilities (3990 Model 3 Only)	17
Attachment to 4.5 or 3.0MB Per Second Channels	17
3390 Data Transfer Rate	18
3990 Reliability, Availability, and Serviceability Features	18
Dual Line Cord	18
Dual Copy	18
Two Independent Storage Clusters	18
Support Facility	19
Service Information Messages (SIMs)	20
Ease of Migration	21
Nondisruptive Installation or Removal Capability	21
Nondisruptive Installation Considerations	22
Intermixed Strings of DASD	23
Environmental Features	24
Chapter 2. IBM 3990 Family	25
3990 Model 1	25
3990 Model 2	25
3990 Model 3	26
Summary of 3990 Models 1, 2, and 3 Improvements	27
Advantages of 3990 Storage Controls	30
Fast Write	34
Dual Copy	36
Branching Transfers, Dual Data Transfer, and Logical Device Allegiance in the 3990 Model 3	38

Nonvolatile Storage	38
3990 Storage Control Components	38
Channels	40
Storage Directors	40
Storage Paths	40
Shared Control Array	40
Support Facility	41
Cache	42
Channel Commands	44
Hardware and Data Availability Considerations	44
Processor Attachment	46
DASD Attachment	47
 Chapter 3. 3990 Storage Control Models and Configurations	 49
3990 Model 1	49
3990 Models 2 and 3	51
3990 Dual-Frame Configuration	54
 Chapter 4. Programming Support	 57
MVS Support for the 3990	57
Minimum MVS Program Levels for 3990 Basic Operations	57
MVS/ESA and MVS/XA Software Capabilities for 3990 Model 3 Extended Functions	58
DFP Releases that Support the 3990 Models 1, 2 and 3	60
VM Support for the 3990	60
Minimum VM Program Levels for 3990 Model 1 and 2 Support	60
VM Software Capabilities for 3990 Model 3 Basic Caching	60
VSE Support for the 3990	61
Minimum VSE Program Levels for 3990 Models 1 and 2 Support	61
Minimum VSE Program Levels for 3990 Model 3 Support	61
TPF Support for the 3990	61
 Chapter 5. Using Dual Copy, DASD Fast Write, and Cache Fast Write in MVS	 63
Environments	63
Using Dual Copy in MVS/ESA and MVS/XA Environments	63
Using DFSMS to Manage Placement of Data Sets on Duplex Pairs	64
3990 Model 3 Fast Dual Copy Performance Considerations	64
Fast Dual Copy Operation	65
Fast Dual Copy Performance Measurements	65
TSO Fast Dual Copy	67
IMS Fast Dual Copy	68
Summary	69
DASD Fast Write in MVS/ESA and MVS/XA Environments	75
MVS DASD Fast Write Data Selection Guidelines	75
DASD Fast Write Performance Measurements	76
Example 1 — DASD Fast Write (TSO)	78
Example 2 — DASD Fast Write (TSO)	79
Example 3 — DASD Fast Write (TSO)	80
Example 4 — DASD Fast Write (Online/TSO)	81
Example 5 — DASD Fast Write (IMS)	82
Guidelines for Using Cache Fast Write with DFSORT in MVS Environments	83
MVS Cache and Cache Fast Write Data Set Selection Guidelines	84
 Chapter 6. Optional Features	 85
Features Available for 3990 Storage Controls	85
Four Channel Switch, Additional	85

Remote/Local Capability	85
Dual Line Cord	85
3390 Attachment	85
Features Available for 3380 Direct Access Storage Devices	87
3990 AJ4/AK4 2-Path String Attachment	87
3990 AJ4/AK4 4-Path String Attachment	87
Engineering Changes Required for Other 3380's	87
3390 Features	87
 Acronyms and Abbreviations	 89
 Glossary	 91
 Bibliography	 99
 Index	 113

Figures

1.	The Storage Subsystem Library	7
2.	A 3990 Model 3 Storage Control with 3390 DASD	13
3.	Examples of Data Transfer Capabilities of DLS and DLSE Subsystems	16
4.	A 3990 Model 3 Storage Control Attached to 3390 and 4-path 3380 DASD	24
5.	A 3880 Model 3 and a 3990 Model 1	27
6.	Two Cross-Configured 3880 Model 3s and a Single 3990 Model 2	28
7.	Two 3880 Model 23s in Dual-Frame and a Single 3990 Model 3	29
8.	3990 Model 2 or Model 3 in DLSE Mode Attached to 3380 Models AJ4 or AK4 4-Path Strings	32
9.	Input/Output Operation Example, 4-Path DASD (DLSE)	33
10.	DASD Fast Write Operation	34
11.	Cache Fast Write Operation	35
12.	Dual Copy Operation	37
13.	Power and Service Boundaries of the 3990 Model 3	39
14.	Modular Power Subsystem	46
15.	3990 Model 1 in DLS Mode Attached to 2-Path DASD Strings	50
16.	3990 Model 3 in DLSE Mode Connected to 3380 Models AJ4 or AK4 4-Path Strings	51
17.	3990 Model 2 in DLS Mode Attached to 2-Path DASD Strings	52
18.	Dual-Frame 3990 Model 2 in DLSE Mode Attached to 4-Path DASD Strings	54

Special Notices

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 or other IBM product in this publication is not intended to state or imply that only IBM's program or other product may be used. Any functionally equivalent program which does not infringe any of IBM's intellectual property rights may be used instead of the IBM product. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Commercial Relations, IBM Corporation, Purchase, NY 10577.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States and/or other countries and may be used throughout this publication:

DB2

DFSMS

ECKD

Enterprise Systems Architecture/370

ESA/370

IBM

MVS

MVS/DFP

MVS/ESA

MVS/SP

MVS/XA

System/360

System/370

VM/XA

3090

Preface

This book is part of the Storage Subsystem Library (SSL)—a set of books that provides information about the hardware components of IBM disk storage subsystems. The SSL includes both direct access storage (DAS) and storage control publications; this book is part of the SSL subset concerned primarily with the 3990 storage control.

This book is intended to help you plan for and use the IBM 3990 Storage Control in MVS, VM, and VSE environments. It describes the various models of the 3990 product and the functions provided by each model. This book is not intended to describe programming interfaces in MVS, VM, or VSE environments.

To use this book effectively, you may also want to read the *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*, the *IBM 3990 Storage Control Reference*, *IBM 3990 Operations Study Guide*, and *IBM 3990 Operations and Recovery Reference*. See the *Maintaining IBM Storage Subsystem Media*, for information on overall DAS maintenance. The books describing the direct access storage devices attached to the 3990 should be available for reference. You should also have a complete set of publications for your operating system available for reference.

About This Book

This book contains:

“Chapter 1. 3990 Overview” provides an overview of the 3990 family of Storage Controls, including an overview of the data availability and performance improvements of each 3990 model, and the enhancements to the reliability, availability, and serviceability characteristics.

“Chapter 2. IBM 3990 Family” provides a detailed introduction to the 3990 Storage Control Models, the advantages of the family of storage controls, and their hardware components.

“Chapter 3. 3990 Storage Control Models and Configurations” describes the 3990 Storage Control Models and their configurations in various environments.

“Chapter 4. Programming Support” describes an overview of the software support in the MVS, MVS/XA, MVS/ESA, VM, VSE, and TPF environments.

“Chapter 5. Using Dual Copy, DASD Fast Write, and Cache Fast Write in MVS Environments” describes how to use the 3990 Model 3 functions DASD fast write, dual copy, and cache fast write in MVS environments.

“Chapter 6. Optional Features” describes the various channel switch features, the dual line cord feature, and DASD features required to attach to a 3990 Storage Control.

“Acronyms and Abbreviations” lists the acronyms used in the Storage Subsystem Library.

“Glossary” lists the terms and abbreviations used in the Storage Subsystem Library.

“Bibliography” lists books that have additional information on the various subjects discussed in this book.

Terminology

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

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 various IBM operating environments. The library is designed to provide hardware and software related information for DAS and storage controls.

Figure 1 on page 7 shows the relationships among the 3990, 3390, and 3380 SSL books in terms of high-level tasks described in each book.

Storage Control Publications

The 3990 subset of the SSL includes:

- *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, the book provides descriptions of the configuration attachment options, optional features, performance characteristics, and software support of the 3990 Storage Control.

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

Provides a functional description of the 3990 Storage Control. The book describes the planning, program installation, and storage management tasks used in typical environments. Configuration examples as well as sample job control language (JCL) for controlling the various functions of the 3990 Storage Control are provided.

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

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

- *Cache Device Administration*, GC35-0101

Specifies the access method services tools for administering a cache device under MVS. The book supports these storage controls: 3990 Model 3, 3880 Model 23, 3880 Model 21, 3880 Model 13, and 3880 Model 11.

- *IBM 3990 Operations Study Guide*, GA32-0131

A study guide for operators of 3990 storage subsystems. Provides general information on system control program commands and messages, and guidelines for basic problem determination.

- *IBM 3990 Operations and Recovery Reference*, GA32-0133

A reference for operational and recovery procedures for 3990 storage subsystems. Provides general guidelines for 3990 problem determination, operations with 3990 extended functions, and also provides recommended recovery actions for the 3990.

- *Introduction to Nonsynchronous Direct Access Storage Subsystems*, GC26-4519

Provides specific information for programmers responsible for writing DASD channel programs that operate in a nonsynchronous environment. This book defines synchronous and nonsynchronous operations, explains ECKD data transfer commands and provides examples of using ECKD commands to build nonsynchronous channel programs.

3390 Publications

The 3390 subset of the SSL includes:

- *IBM 3390 Direct Access Storage Introduction*, GC26-4573

Provides a complete description of each 3390 model, including characteristics, features, and capabilities. In addition, the configuration and attachment options are described along with other information that will help you design a subsystem to meet your needs.

- *Using IBM 3390 Direct Access Storage in an MVS Environment*, GC26-4574

Provides specific guidance for using 3390s in an MVS/ESA, MVS/XA or MVS/370 operating environment. The book provides detailed instruction for planning the addition of new 3390 devices, installing devices, moving data to new devices, and performing ongoing storage subsystem management.

- *Using IBM 3390 Direct Access Storage in a VM Environment*, GC26-4575

Provides specific guidance for using 3390s in a VM/SP HPO, or VM/XA SP operating environment. The book provides detailed instruction for planning the addition of new 3390s, installing devices, moving data to new devices, and performing ongoing storage subsystem management. In addition, storage considerations related to guest systems are addressed.

- *IBM 3390 Direct Access Storage Reference Summary*, GX26-4577

Provides a summary of 3390 capacity, performance, and operating characteristics.

3380 Publications

The 3380 subset of the SSL includes:

- *IBM 3380 Direct Access Storage Introduction*, GC26-4491

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 book does *not* cover 3380 Model CJ2.

- *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference*, GC26-4497

Provides a complete description of the 3380 direct channel attach Model CJ2 characteristics, features, capabilities, and string configuration options.

- *Using the IBM 3380 Direct Access Storage in an MVS Environment*, GC26-4492

Provides specific guidance for using the 3380 in MVS/ESA, MVS/XA, or MVS/370 operating environments. The book 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 some ongoing activities to maintain a reliable storage subsystem.

- *Using the IBM 3380 Direct Access Storage in a VM Environment*, GC26-4493
Provides specific guidance for using the 3380 in a VM/SP, VM/SP HPO, or VM/XA SP operating environment. The book 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, hardware considerations related to guest systems are addressed.
- *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. The book provides instruction for planning the addition of new 3380 devices, installing devices, moving data to new devices, and performing ongoing storage subsystem management.
- *IBM 3380 Direct Access Storage Reference Summary*, GX26-1678
Provides a summary of 3380 capacity, performance, and operating characteristics.

Shared SSL Publications

The following publications contain information relevant to the entire Storage Subsystem Library.

- *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.
- *Storage Subsystem Library Master Bibliography, Index, and Glossary*, GC26-4496
Provides a central source for information related to storage subsystem topics. Books for IBM 3390 DAS, IBM 3380 DAS, and 3990 Storage Controls are indexed in this publication. An overview of the material in the Storage Subsystem Library is provided with this index.

Storage Subsystem Library

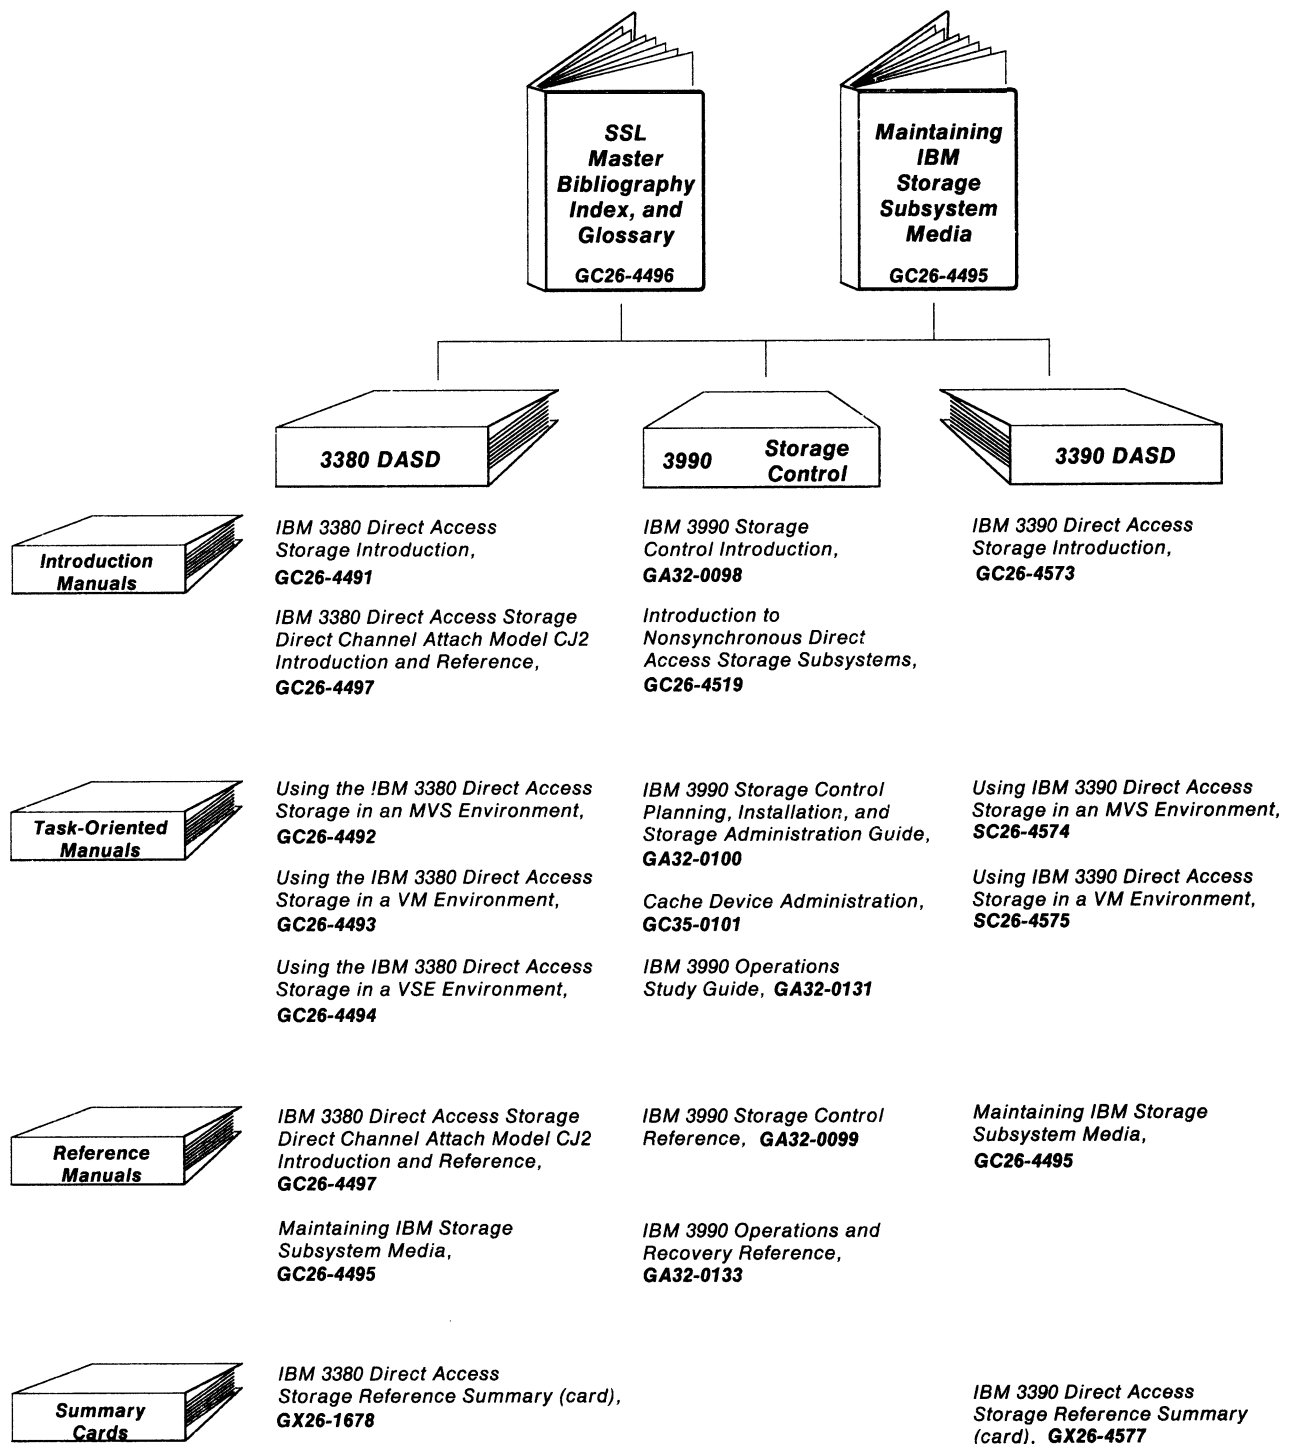


Figure 1. The Storage Subsystem Library

Storage Subsystem Library Ordering Information

The entire SSL or parts of it tailored to your hardware and software environment can be ordered with bill of form numbers.

3390 and 3990 Publications

You can obtain a copy of **every manual** in the 3390 and 3990 subsets of the SSL with one order number, **SBOF-3124**. Select one of the following bill of form numbers to obtain information tailored to your hardware and software environment. To obtain an individual manual, use its order number.

Table 1. 3390 and 3990 Subsets				
Title	MVS SBOF- 3121	VM SBOF- 3122	3990 GBOF- 0366	SBOF- 3124
<i>IBM 3390 Direct Access Storage Introduction, GC26-4573</i>	X	X		X
<i>Using IBM 3390 Direct Access Storage in an MVS Environment, SC26-4574</i>	X			X
<i>Using IBM 3390 Direct Access Storage in a VM Environment, SC26-4575</i>		X		X
<i>Maintaining IBM Storage Subsystem Media, GC26-4495</i>	X	X		X
<i>IBM 3390 Direct Access Storage Reference Summary, GX26-4577</i>	X	X		X
<i>IBM 3990 Storage Control Introduction, GA32-0098</i>			X	X
<i>IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide, GA32-0100</i>			X	X
<i>IBM 3990 Storage Control Reference, GA32-0099</i>			X	X
<i>Cache Device Administration, GC35-0101</i>			X	X
<i>IBM 3990 Operations Study Guide, GA32-0131</i>			X	X
<i>IBM 3990 Operations and Recovery Reference, GA32-0133</i>			X	X
<i>Introduction to Nonsynchronous Direct Access Storage Subsystems, GC26-4519</i>			X	X
<i>Storage Subsystem Library Master Bibliography, Index, and Glossary, GC26-4496</i>	X	X	X	X
<i>Binder and 3390 inserts, SX26-3777</i>	X	X		X
<i>Binder and 3990 inserts, GX26-3768</i>			X	X

3380 and 3990 Publications

You can obtain a copy of **every manual** in the 3380 and 3990 subsets of the SSL using one General Bill of Forms (GBOF) number, **GBOF-1762**. Select one of the following bill of form numbers to obtain information tailored to your hardware and software environment. To obtain an individual manual, use its order number.

Table 2. 3380 and 3990 Subsets							
Title	MVS GBOF- 1756	VM GBOF- 1757	VSE GBOF- 1758	CJ2/MVS GBOF- 1759	CJ2/VM GBOF- 1760	CJ2/VSE GBOF- 1761	3990 GBOF- 0366
<i>IBM 3380 Direct Access Storage Introduction, GC26-4491</i>	X	X	X				
<i>IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference, GC26-4497</i>				X	X	X	
<i>Using IBM 3380 Direct Access Storage in an MVS Environment, GC26-4492</i>	X			X			
<i>Using IBM 3380 Direct Access Storage in a VM Environment, GC26-4493</i>		X			X		
<i>Using IBM 3380 Direct Access Storage in a VSE Environment, GC26-4494</i>			X			X	
<i>Maintaining IBM Storage Subsystem Media, GC26-4495</i>	X	X	X	X	X	X	
<i>IBM 3380 Direct Access Storage Reference Summary, GX26-1678</i>	X	X	X	X	X	X	
<i>IBM 3990 Storage Control Introduction, GA32-0098</i>							X
<i>IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide, GA32-0100</i>							X
<i>IBM 3990 Storage Control Reference, GA32-0099</i>							X
<i>Cache Device Administration, GC35-0101</i>							X
<i>IBM 3990 Operations Study Guide, GA32-0131</i>							X
<i>IBM 3990 Operations and Recovery Reference, GA32-0133</i>							X
<i>Introduction to Nonsynchronous Direct Access Storage Subsystems, GC26-4519</i>							X
<i>Storage Subsystem Library Master Bibliography, Index, and Glossary, GC26-4496</i>	X	X	X	X	X	X	X
<i>Binder and 3380 inserts, GX26-3767</i>	X	X	X	X	X	X	
<i>Binder and 3990 inserts, GX26-3768</i>							X

Storage Subsystem Library Binders

Binder kits are available to help organize your library. Kits consist of a binder with identifying cover and spine inserts for 3390, 3380, or 3990 manuals, and are included when you order the following numbers:

- SBOF-3121 and SBOF-3122 include a binder with 3390 inserts.
- GBOF-1756 through GBOF-1761 include a binder with 3380 inserts.
- GBOF-0366 includes a binder with 3990 inserts.
- SBOF-3124 includes binders and inserts for both 3990 **and** 3390.
- GBOF-1762 includes binders and inserts for both 3990 **and** 3380.

Binder kits may also be ordered separately.

- Order number GX26-3777 contains a binder and 3390 inserts.
- Order number GX26-3767 contains a binder and 3380 inserts.
- Order number GX26-3768 contains a binder and 3990 inserts.

Related Publications

"Bibliography" on page 99 contains a list of publications that are referred to in this book or that may provide additional, related information. The bibliography includes a short description of each publication.

IBM 3990 Transaction Processing Facility Support RPQs describes the 3990 RPQs available for support of the Transaction Processing Facility operating system. Initialization procedures, channel commands, and sense data specific to these RPQs are provided.

Summary of Changes

Second Edition, June 1990

Technical Changes

The following feature numbers have been changed to engineering changes:

- 9003 - 3990 AA4 Installation/Maintenance Update Feature has been changed to engineering change EC444944 ECA130.
- 9052 - 3990 AD4/AE4 Installation/Maintenance Update Feature has been changed to engineering change EC465359 ECA131.

Information has been added on the attachment to 3390 DASD, the concurrent media maintenance as provided by ICKDSF, the dual line cord feature of the 3990, and the media maintenance on duplex pairs as provided by ICKDSF Release 12.

Other Changes

This book has been revised with minor technical updates and an acronym list has been added to this book.

The books, *Cache Device Administration*, GC35-0101, and *Introduction to Nonsynchronous Direct Access Storage Subsystems*, GC26-4519, have been added to the Storage Subsystem Library.

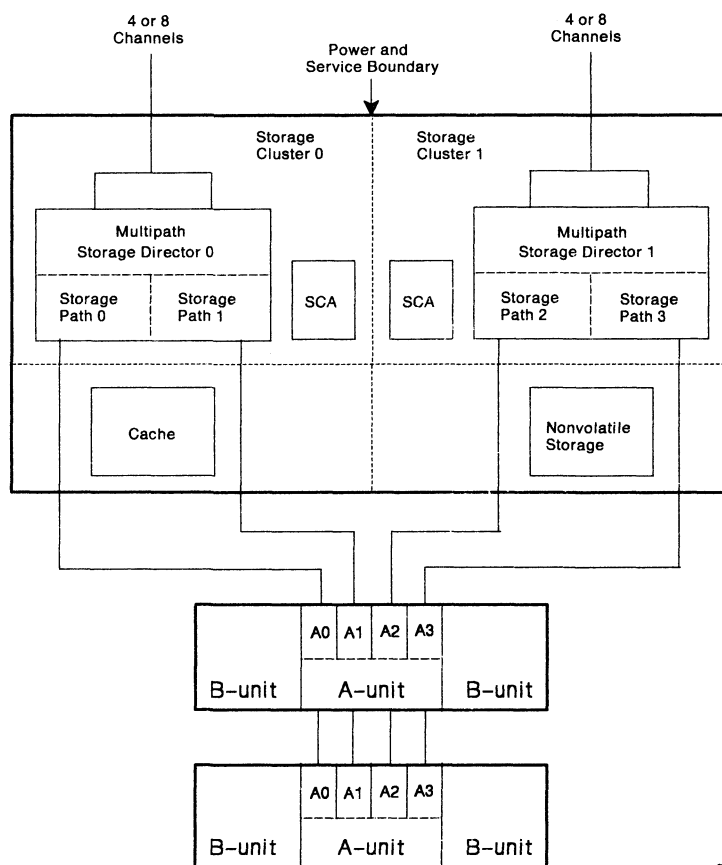
Chapter 1. 3990 Overview

This chapter is an **overview** of the IBM 3990 family of storage controls. Today's businesses expect exceptional performance from their storage subsystems. The 3990 family of storage controls is designed to meet and exceed the data availability, performance, and reliability requirements that businesses demand. More detailed descriptions are provided in "Chapter 2. IBM 3990 Family" on page 25.

The 3990 Storage Control is available in a full range of configurations and models:

- **3990 Model 1**, designed for intermediate systems with a requirement for two DASD paths.
- **3990 Model 2**, designed for intermediate-to-large systems requiring more than two DASD paths.
- **3990 Model 3**, also designed for intermediate-to-large systems requiring more than two DASD paths *and* also offering four cache sizes from 32 to 256MB and 4MB of nonvolatile storage (NVS).

Figure 2 shows a 3990 Model 3 Storage Control with 3390 DASD.



Legend: SCA = Shared Control Array

SL220N38

Figure 2. A 3990 Model 3 Storage Control with 3390 DASD

3990 Standards of Data Availability and Performance

All 3990 models address data availability and performance requirements by using:

- Internal processing that is faster than that in previous IBM Storage Controls.
- Faster access to data on the DASD is achieved because control information that formerly was kept in the DASD A-unit is now kept in the 3990.

In addition, the 3990 Models 2 and 3 provide:

- Four storage paths between the 3990 Storage Control and the DASD; twice as many paths as available with the 3880 Models 3 or 23.
- Four independent I/O operations can be active concurrently within the same 4-path string when attached to the IBM 3390 A-units or 3380 Models AJ4 or AK4.
- Two independent storage clusters per frame provide separate power and service regions.
- The capability to install the dual line cord feature. With this feature, a single-frame 3990 provides all the data availability and reliability that was previously available only in a 3990 dual-frame configuration.
- 3390 attachment at the 4.2MB per second data transfer rate.

In addition to those improvements given above, the 3990 Model 3 provides:

- Larger cache sizes that result in finding the requested data in cache more often, thus reducing subsystem response time.
- Use of IBM's one million-bit chip for cache.
- Cache slot segmentation resulting in more efficient cache space utilization.
- Improved internal cache algorithms resulting in greater efficiency.
- Dual data transfer, branching data transfers, and logical device allegiance between the 3990 and the channel, and the 3990 and the device.
- Two fast write capabilities, **DASD fast write** and **cache fast write**, each allowing cache speed write operations.
- A **dual copy** capability that greatly increases the availability and accessibility of data. The dual copy capability writes all selected data to two different devices, independent of host processing. Information is automatically available on the secondary device if access to the primary device is not possible.
- Separate power and service regions for cache and nonvolatile storage.

3990 Standards of Reliability, Availability, and Serviceability

All models of the 3990 family address reliability, availability, and serviceability requirements by using new components and design enhancements including:

- Each storage cluster, within its own power and service region, contains a separate microprocessor-controlled support facility. These independent support facilities provide effective availability and serviceability functions.
- Service Information Messages (SIMs) for improved problem determination and impact assessments.
- Internal design enhancements and newly designed, reliable components to assure a high level of reliable hardware and efficient operation.
- Improved fault isolation for more efficient problem determination.
- Vital product data (VPD) storage for subsystem configuration information.
- A reset allegiance and a reset event notification capability. (See "Minimum MVS Program Levels for 3990 Basic Operations" on page 57 for more information on these two capabilities.)

- Concurrent media maintenance with ICKDSF Release 11 and above allowing media errors to be fixed without taking the device offline.

In addition, the 3990 Models 2 and 3 provide:

- A *nondisruptive DASD installation* capability that may allow additional:
 - 4-path 3380 B-units, or
 - a second 4-path 3380 string, or
 - a second 3390 string or
 - 3390 B-units, or B-unit upgrades
 to be installed in a 3990 subsystem without disrupting host access to the existing DASD, assuming the new devices are properly defined to the host and the storage control.
- Two independent storage clusters per frame providing separate power and service boundaries.
- Better reliability, availability, and serviceability characteristics than similarly configured 3880s.
- The capability to install the dual line cord feature. With this feature, a single-frame 3990 provides all the data availability and reliability that was previously available only in a 3990 dual-frame configuration.

In addition to those improvements given above, the 3990 Model 3 provides:

- With ICKDSF Release 12 and the appropriate level of microcode in the 3990 Model 3, media maintenance may be performed on either device in a duplex pair without resetting the pair to simplex state.
- A 4MB nonvolatile storage (NVS) that supports both the DASD fast write and the dual copy capabilities. Dual copy provides for the automatic creation, update, and cut-over to a second copy of critical data. The NVS uses a battery to protect data for up to 48 hours.

Note: For maximum data availability, it is highly recommended that all 3990 Model 3 subsystems using dual copy or DASD fast write have either the dual line cord feature or be dual-frame configurations.

The 3990 Storage Control answers the needs of both today and the future's data processing requirements. Each model of the 3990 Storage Control directly addresses the data availability, performance, and reliability issues of storage management. Even greater performance is available from those 3990 models with cache and the nonvolatile storage.

3990 Performance Features

The 3990 provides a number of performance features. The features described below are significant performance improvements over the previous family of IBM storage controls.

Faster and More Efficient Microprocessing

All the 3990 Storage Controls use faster internal processing techniques than previously available and more efficient microprocessing that permits faster data access.

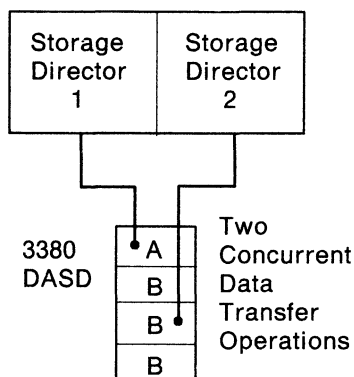
Four Storage Paths

As a new standard of data availability and overall performance, twice as many storage paths are provided in the 3990 Models 2 or 3 than in the 3880 Model 3 or 23 Storage Controls. When attached to 4-path DASD, the 3990 can operate in a high performance and high data availability mode: Device Level Selection Enhanced (DLSE). When the 3990 Model 2 or 3 is attached to 4-path DASD, the accessibility of data is even greater than the accessibility of data when attached to 3380 AD4/AE4 devices operating in Device Level Selection (DLS) mode. Attaching to 3390 A-units or the 4-path 3380 AJ4/AK4 units allows the 3990 to operate in DLSE mode giving twice as many paths to each device. (See "Device Level Selection Enhanced Mode of Operation" on page 30 for more information.)

With the Device Level Selection Enhanced mode of data transfer, simultaneous data transfer is possible over all four storage paths within the same 4-path string. A DLSE storage subsystem, even with one of its four paths disabled, performs better than two DLS subsystems with all four paths available. (See "Device Level Selection Mode of Operation" on page 30 for more information.)

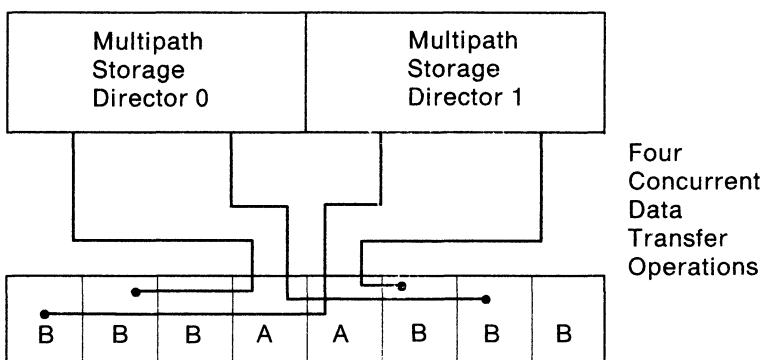
Figure 3 shows the data transfer capabilities of both the DLS subsystem and the DLSE subsystem.

**A Single 3880
Model 3 Storage Control**



A DLS Subsystem

**A Single 3990
Model 2 Storage Control**



3380 4-Path DASD

A DLSE Subsystem

SILVCO7M8

Figure 3. Examples of Data Transfer Capabilities of DLS and DLSE Subsystems

Cache (3990 Model 3 Only)

3990 Model 3 cache performance enhancements include:

- Cache sizes of 32, 64, 128, or 256MB.
- Up to eight concurrent operations inside the cache provide more internal efficiency. This efficiency means less contention for components within the subsystem and a higher data throughput.
- Logical device allegiance, branching transfers, and dual data transfers permit the 3990 Model 3 to perform more tasks simultaneously, resulting in greater performance. See "Branching Transfers, Dual Data Transfer, and Logical Device Allegiance in the 3990 Model 3" on page 38 for more information on these capabilities.

- Sequential cache management algorithms have been refined to provide faster cache access to sequential data.
- Cache is managed in 16K segments providing more efficient use of the cache.

Fast Write Capabilities (3990 Model 3 Only)

Two significant performance benefits to 3990 Model 3 users are:

- **DASD fast write:**
 - Permits write operations to be performed at cache speeds.
 - Provides data integrity equivalent to DASD writes.
 - Eliminates the requirement to write the data to the DASD immediately.
 - Maintains a copy of the data in the NVS until it is destaged (written from the cache to the DASD).
- **Cache fast write:**
 - Recommended for small- to intermediate-size sort work files, cache fast write permits the selected write operations to be performed at cache speeds without requiring writes to DASD.

DASD Fast Write and Cache Fast Write Advantages

Both fast write capabilities permit write operations to be performed at cache speed, eliminating the need to write data to the DASD immediately. These capabilities provide faster response times, particularly for DASD volumes with high write activity resulting in improved overall subsystem performance. (See "Fast Write" on page 34 for more information on the DASD fast write and cache fast write capabilities.)

Attachment to 4.5 or 3.0MB Per Second Channels

The 3990 Model 3 can transfer cache data at channel speed. The actions required to run the 3990 at the 4.5MB per second channel speed are: run the I/O Configuration Program (IOCP) specifying 4.5MB per second as the channel speed and set the channel speed in the 3990 Vital Product Data for the appropriate channel.

On the same 3990 Model 3, some channel interfaces can operate at 4.5MB per second while other channel interfaces can operate at 3.0MB per second, depending on the processor type and the attached DASD. All channels accessing 3990 volumes must operate at 4.5MB per second.

3990 storage controls operating at 4.5MB per second must be installed ahead of non-3990 control units, regardless of their operating speed. A 3990 operating at 3MB per second may be installed ahead of a 3990 operating at 4.5MB per second. The 3990 storage controls that operate only at 3.0MB per second are not required to be installed ahead of all other storage controls, but it is recommended.

3990s may be installed using blue cables for channel lengths up to 400 feet. For channel lengths up to 300 feet, blue and gray cables may be mixed. For distances between 300 and 400 feet, gray cables may be used after submission of RPQ 8B0115.

3390 Data Transfer Rate

3390s attached to either the 3990 Model 2 or 3 transfer data at 4.2MB per second. This data transfer speed is a significant improvement over the 3.0MB per second data transfer rate of the IBM 3380.

3990 Reliability, Availability, and Serviceability Features

The 3990 provides a number of reliability, availability, and serviceability features. New components and design enhancements improve the IBM 3990's reliability, availability, and serviceability (RAS) characteristics over those of the previous family of IBM Storage Controls. These features are briefly described below.

Dual Line Cord

The dual line cord feature provides protection from failure equivalent to a dual-frame configuration in a single machine. See "Dual Line Cord" on page 85 for details about features numbers.

Dual Copy

A major improvement in data reliability and availability is provided by **dual copy**. Available on the 3990 Model 3, dual copy provides a duplicate copy of a volume on another device in the subsystem. Dual copy is defined for each volume by the installation. For a dual copy volume, each time the primary device is updated, the 3990 Model 3 automatically updates the secondary device. The data transfer operations are automatically switched to the secondary device if a failure occurs on the primary device.

Dual copy provides important advantages:

- Protects a critical application from a single device failure.
- Provides for near-continuous operations by providing a duplicate copy if loss of access occurs to one of the dual copy devices.
- Ensures that the second copy is identical to the primary copy.

See "Dual Copy" on page 36 for more information on the dual copy capability.

The DASD fast write capability and the dual copy capability can be combined resulting in a **fast dual copy**. See "Chapter 5. Using Dual Copy, DASD Fast Write, and Cache Fast Write in MVS Environments" on page 63 for additional information.

Two Independent Storage Clusters

Single-frame 3990 Models 2 and 3 provide a *pair of independent storage clusters* (usually shortened to clusters). A clear benefit of each of the clusters is that they are independent components. Each cluster provides a separate power and service region and two separate paths to the DASD. Loss of power to one cluster does not prevent access to data because processing continues through the other storage cluster. This processing is possible because all DASD are connected to both storage clusters. A service representative can take one cluster for maintenance while cache and DASD access continues through the remaining cluster's two paths.

In the 3990 Model 3, cache and nonvolatile storage are shared by the storage paths, but are logically and physically separate from the storage clusters.

Support Facility

Each storage cluster has its own support facility. A major reliability, availability, and serviceability (RAS) enhancement, the support facility and the independent service regions permit concurrent maintenance. The support facility also provides a remote maintenance support capability. Among other tasks, the support facility monitors subsystem activity, generates the service information messages (SIMs), communicates with the other support facility (in 3990 Models 2 and 3), runs maintenance analysis procedures (MAPs) and diagnostics, and logs error conditions on diskette storage.

Vital Product Data (VPD) Storage

Each storage cluster contains a vital product data (VPD) storage module. The subsystem configuration information saved in VPD storage includes: 3990 features, the subsystem mode of operation, subsystem identifiers, subsystem configuration, control unit addresses for each channel, channel speed, and the number of addressable devices that can be attached to the logical DASD subsystem. The service representative keys the VPD into the 3990 during installation. With varying levels of concurrency, the service representative can change the VPD at any time, but the changes are not activated until the next 3990 IML. The contents of VPD storage are duplicated on the writeable diskette for each storage cluster.

Concurrent Maintenance

Concurrent maintenance greatly improves the availability of DASD data. Because of the power and service boundaries built into the 3990, one storage cluster can continue to access cache, NVS, and DASD while maintenance is taking place on the other storage cluster.

A service action can be performed on cache while direct access to DASD is provided through the storage clusters.

A service action can be performed on nonvolatile storage while caching operations and direct access to DASD continues through the storage clusters.

The concurrent maintenance and remote maintenance support capabilities in combination with the service information messages (SIMs) reduce problem determination time and effort, the total maintenance time, and the impact of the service on the installation's operations.

Concurrent Media Maintenance

The Device Support Facilities (ICKDSF) provide concurrent media maintenance for 3380 and 3390 devices. With concurrent media maintenance, Device Support Facilities can INSPECT tracks on a volume while allowing normal access to all the data on the volume. You do not have to quiesce or vary the volume offline but can continue normal job stream processing for most media maintenance activity. Concurrent media maintenance is automatically invoked when you specify the PRESERVE parameter in the INSPECT command. For more information on the INSPECT command, see *Device Support Facilities User's Guide and Reference*. For more information on concurrent media maintenance, see *Maintaining IBM Storage Subsystem Media*.

To have concurrent media maintenance for either 3380s or 3390s, you need the 3390 Attachment Feature 6120 on a 3990 Model 2 or 3 and ICKDSF Release 11 or higher.

Media Maintenance to Duplex Pairs

When the 3990 Model 3 has the most current level of microcode for extendfunction installed, ICKDSF Release 12 provides the capability for performing media maintenance on a duplex pair without first needing to reset it to simplex. An inspect operation may be performed to the primary device of the pair while normal dual copy operations continue. If such operations must be performed on the secondary device, the duplex pair must first be suspended. After the maintenance is complete, normal dual copy operation is resumed through use of the IDCAMS RESETTODUPLEX command. Because changed cylinders were logged while the pair was suspended, only changed cylinders need to be copied from the primary device to the secondary device. (Before Release 12, the pair had to be reset to simplex for maintenance and resumption of dual copy operation required a full volume copy to the secondary device.)

Remote Maintenance Support

The remote maintenance support capability permits a support representative in a remote field support center to establish communication with either storage cluster in a 3990 through an external modem. Once the communication link is established, the remote support representative can analyze the error data and send maintenance information to the service representative on site.

An access code, valid for one hour, is required to authorize a remote connection. All remote maintenance support sessions are protected by a 3990-generated access code. Having the 3990 generate the access code ensures the integrity of the remote session and prevents unauthorized data links to the 3990.

Writeable Diskette

The diskettes contain 3990 microcode, microcode patches, error log, and the soft-copy maintenance analysis procedures (MAPs).

During either local or remote maintenance support activities, microcode patches can be transmitted to the 3990 support facility and stored on the 3990's diskettes. Microcode patches written on the diskettes are not lost across IMLs. Microcode patches can be downloaded to the diskettes and saved for later use without disrupting host access to data.

The installation retains control of when microcode patches are implemented. Those patches written on the diskettes are not implemented until the installation asks a local service representative to do so.

Service Information Messages (SIMs)

Service information message-format sense data (SIMs) and SIM alert messages provide a much improved method of notifying the user of any failure. SIMs are reported by the 3990 storage control and 3390 DASD. For more information on 3990 SIMs, see *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*. For more information on 3390 SIMs, see *Maintaining IBM Storage Subsystem Media*.

SIM Alert Messages

SIM alert messages inform operators of a hardware failure. The SIM alert messages are displayed by a host error recovery program (ERP). The SIM alert messages are intended to prompt the operator to initiate the action necessary for an installation to get a problem fixed.

SIM Messages

The SIM identifies the impact of the failure, and the effect of the repair action through an Environmental Record Editing and Printing (EREP) detail edit report. The installation can use this information to schedule the repair action, based on the severity of the failure, and the resource required to repair it. These messages provide information about the failure so that the service representative can bring the most likely repair parts, significantly reducing repair time and the possibility of return calls.

Because of the service information message and the SIM alert message, any failure can, in most cases, be fixed quickly through replacement parts identified by the SIM. Also, the installation can make a more informed decision about scheduling any repair action. See *IBM 3990 Operations and Recovery Reference* for more information about problem determination.

Ease of Migration

All 3990 features are field-installable and field-upgradable. Any 3990 model can be upgraded to a higher model number. Larger cache size increments can be added to a 3990 Model 3 in the customer's installation. (See "Chapter 6. Optional Features" on page 85 for information on the features available on a 3990 Storage Control.)

Nondisruptive Installation or Removal Capability

The 3990 Storage Control in DLSE (4-path) mode allows 4-path DASD units to be installed or removed while retaining access to data on the other units in the subsystem.

This capability applies to:

- The installation or removal of a 4-path 3380 second string, where the first string is either a 3380 4-path or a 3390 string
- The installation or removal of a 3390 as a second string, where the 3990 already has 3390 Attachment Feature 6120 installed and the first string is a 3380 4-path string or a 3390 string with a channel speed of 4.5MB
- The installation or removal of end-of-string B-units to or from an existing 3380 4-path string or a 3390 string
- Conversion of a 3390 Model 1 to a Model 2
- The 2-HDA 3390 capacity upgrade, where data on the other HDAs in that unit can be accessed as well.

The 3990 subsystem nondisruptive installation or removal capability allows the additional devices to be installed and brought online (or removed) without:

- Disrupting any application
- Removing power from the 3990
- Stopping the operating system
- Taking any *channel paths* offline.

During the nondisruptive DASD installation process, only one of the four internal paths to the DASD string must be quiesced at any given time.

Note: To take advantage of the nondisruptive DASD installation capability, the devices planned for future installation must have assigned addresses that are predefined to the operating system, the storage control, and the processor channel subsystem in supported processor environments. Also, the 3990 Storage Control must have the 3390 Attachment Feature 6120 installed prior to the 3390 DASD installation. 3390s must be attached to channels with 4.5MB per second data transfer rate and these channels must be identified in the 3990 VPD.

Nondisruptive Installation Considerations

When planning for your 3390 installation, you should carefully consider the factors that can keep a particular action from being disruptive.

The following actions are disruptive, and should be completed prior to the installation of the first 3390 on a 3990 subsystem.

- The 3390 Attachment Feature 6120 must be installed on the 3990 involved.
- For MVS environments, an IOGEN is required to change unit names. All unit name changes require an IOGEN, and because 3390 is a new device type, the unit name must be changed.
- For VM environments, the real I/O configuration file and the I/O configuration program (IOCP) must be changed, specifying the new device type.
- A change to 3990 vital product data (VPD) is required. The 3990 VPD must be changed when:
 - A 3990 running in DLS mode must have the VPD changed to DLSE mode.
 - A 3990 Model 2 running in DLSE mode, with the channel speed set for 3.0MB per second data transfer rate, must have the VPD changed to 4.5MB per second channel speed to support the 3390 4.2MB per second data transfer rate.
 - A second string is being added to a 3990 whose VPD defines only one string. The VPD must be changed to reflect the new subsystem configuration.
 - Changing from intermix DLSE to full DLSE configuration.

Table 3 on page 23 shows which 3390 installations are nondisruptive during the actions listed. A *Yes* in the column for any action means that access to other data is unaffected (nondisruptive). In contrast, a *No* in the column means that access to other data is not possible during the action indicated. For example, during a 3390 Model 1 to Model 2 conversion, data stored on other units in the same string and on the other string on the storage control can be accessed, but data on the head disk assemblies (HDA) in the affected unit cannot be accessed.

Table 3. 3390 and 4-Path 3380 Nondisruptive Installations — Quick Reference

Action...	Can access data on the other string	Can access data on other units in the same string	Can access data on HDAs in the affected unit
Installation or removal of a second string (3390 A-unit), where the first string is a 4-path string, and the 3990 already has Feature 6120 installed	Yes	N/A ¹	N/A ¹
Installation or removal of a second 3380 4-path string (3380 AJ4/AK4), where the first string is a 3380 4-path or 3390 string	Yes	N/A ¹	N/A ¹
Installation or removal of 3390 B-units	Yes	Yes	N/A ¹
Installation or removal of end-of-string 4-path 3380 B-units	Yes	Yes	N/A ¹
3390 Model 1 to Model 2 Conversion			
Model A14 to Model A24	Yes	Yes	No
Model A18 to Model A28	Yes	Yes	No
Model B14 to Model B24	Yes	Yes	No
Model B18 to Model B28	Yes	Yes	No
Model B1C to Model B2C	Yes	Yes	No
3390 2-HDA Capacity Upgrade			
Model A14 to Model A18	Yes	Yes	Yes
Model A24 to Model A28	Yes	Yes	Yes
Model B14 to Model B18	Yes	Yes	Yes
Model B18 to Model B1C	Yes	Yes	Yes
Model B24 to Model B28	Yes	Yes	Yes
Model B28 to Model B2C	Yes	Yes	Yes
Note: ¹ N/A (not applicable).			

Intermixed Strings of DASD

A 3990 Model 2 or 3 can attach:

- One or two 3390 strings
- One or two 3380 AJ4 or AK4 4-path strings
- One 3390 string and one 3380 AJ4 or AK4 4-path string
- One 3390 string and one or two 2-path 3380 AJ4, AK4, AD4, or AE4 strings
- One 3380 AJ4 or AK4 4-path string and one or two 3380 2-path AD4, AE4, or AA4 strings
- One to four 2-path 3380 AA4, AD4, AE4, AJ4, or AK4 strings.

Note: 3390 and 3380 Model AA4 strings **cannot attach** to the same 3990 Storage Control.

Figure 4 on page 24 shows a 3990 Model 3 storage control attached to 3390 and 4-path 3380 DASD. For figures showing the supported 3990 configurations, see *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*.

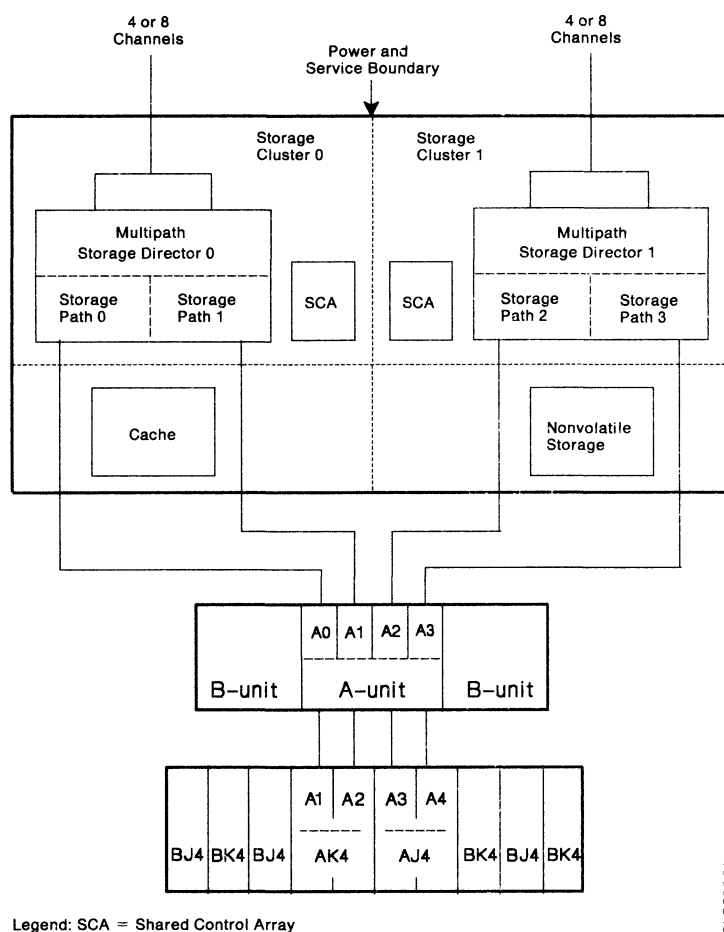


Figure 4. A 3990 Model 3 Storage Control Attached to 3390 and 4-path 3380 DASD

Environmental Features

Improved technology and packaging reduce the power and cooling requirements for a 3990 Storage Control from those required for the 3880 Storage Control on a per path basis comparison.

For information on power and cooling requirements, see the *IBM Input/Output Equipment Installation – Physical Planning for System/360, System/370, and 4300 Processors*.

Also, because four paths to DASD are built into each single-frame 3990 Model 2 or 3, twice as many paths to DASD are provided over the number of paths available in the previous family of storage controls. For example, *one 3990 Model 2* provides the same number of paths (four) to the DASD as there are in *two 3880 Model 3s or 23s*. The 3990 Model 2 or 3 provides four paths to DASD in the same floor space as two paths from a 3880 to DASD. Because side clearance is no longer required, the total required floor space is reduced by more than half.

Because of these improved environmental features, twice as many devices can be attached to the 3990 Model 2 or 3 as can be attached to the 3880 Model 3 or 23.

Chapter 2. IBM 3990 Family

The structure and features of the 3990 family are a logical extension and significant improvement over the structure and capabilities of the IBM 3880 family of storage controls. The IBM 3990 Storage Control family attaches all models of the IBM 3390 and 3380 Direct Access Storage Devices except 3380 Models A04 and CJ2.

The 3990 is available in three models:

- 3990 Model 1 with two separate storage paths to DASD
- 3990 Model 2 with four separate storage paths to DASD
- 3990 Model 3 with four separate storage paths to DASD and 32, 64, 128, or 256MB of cache, and 4MB of nonvolatile storage.

A single 3990 Model 2 or a single 3990 Model 3 can attach up to 16 different channels.

For a listing of the software support for the 3990, see "Chapter 4. Programming Support" on page 57.

3990 Model 1

A single 3990 Model 1 provides the same function as a single-frame 3880 Model 3 plus the performance and reliability characteristics of the 3990 family. It is intended for midrange systems that require two storage paths to the DASD.

See Figure 5 on page 27 for a comparison of a 3880 Model 3 and a 3990 Model 1. See "3990 Model 1" on page 49 for a more complete description of a 3990 Model 1.

The 3990 Model 1 attaches all models of the IBM 3380 Direct Access Storage Devices except 3380 Models A04 and CJ2. The 3990 Model 1 cannot attach the IBM 3390.

3990 Model 2

A single 3990 Model 2 may provide the same function as two cross-configured 3880 Model 3s plus the performance and reliability characteristics of the 3990. It is intended for intermediate-to-large systems that require up to four storage paths to the DASD.

The single 3990 Model 2 is a fundamental enhancement of two cross-configured 3880 Model 3s. Because of its improved data availability, performance, and the increased reliability, availability, and serviceability characteristics, the 3990 Model 2 will satisfy many DASD configuration requirements.

See Figure 6 on page 28 for a comparison of two 3880 Model 3s and a single 3990 Model 2. See "3990 Models 2 and 3" on page 51 for a more complete description of a 3990 Model 2.

3990 Model 3

The 3990 Model 3, with 32 to 256MB of cache and 4MB of NVS, provides all of the performance and reliability characteristics of the 3990.

The single 3990 Model 3 is a fundamental enhancement of a dual-frame 3880 Model 23. In one 3990 Model 3, with the added benefits of cache fast write, DASD fast write, and dual copy:

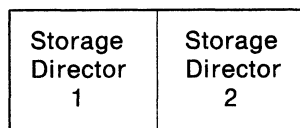
- The cost per path to DASD is substantially less than a pair of 3880 Model 23s in a dual-frame configuration.
- The performance is much better.
- The DASD availability is significantly improved.

See Figure 7 on page 29 for a comparison of a pair of 3880 Model 23s in a dual-frame configuration and a single 3990 Model 3. See "3990 Models 2 and 3" on page 51 for a more complete description of a 3990 Model 3.

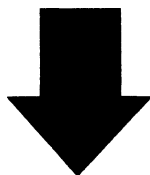
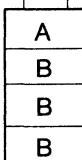
Summary of 3990 Models 1, 2, and 3 Improvements

The following figures summarize the 3990 Model 1, 2, and 3 improvements over the 3880 Storage Controls.

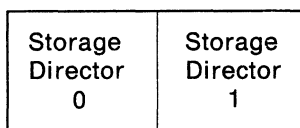
A Single 3880 Model 3 Storage Control



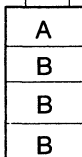
3380
DASD



A Single 3990 Model 1 Storage Control



3380
DASD



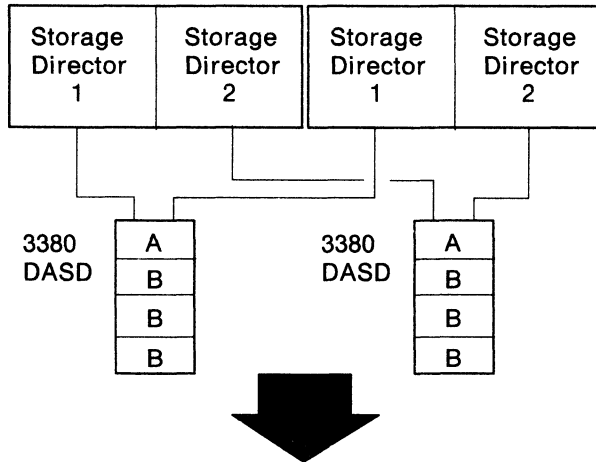
3990 Model 1 Improvements

- Faster Microprocessing
- Service Information Messages (SIMs)
- Online SIM Alert Messages
- Faster access to DASD by moving DPS array into the Storage Control
- Field Upgradable to 3990 Model 2
- Remote Maintenance
- Support Facility
- Writeable Diskette
- Reduced Power and Cooling Requirements

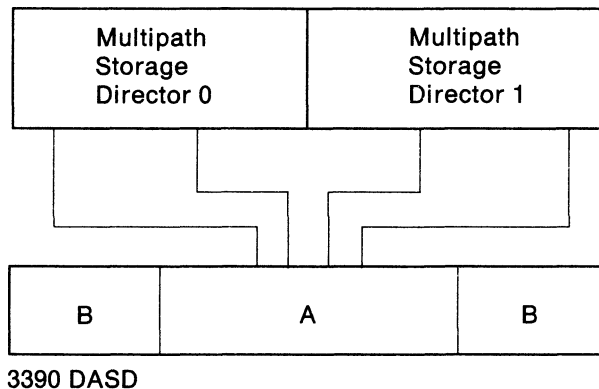
Figure 5. A 3880 Model 3 and a 3990 Model 1

slvc01m1

Two Cross-Configured 3880 Model 3 Storage Controls



A Single 3990 Model 2 Storage Control in DLSE Mode



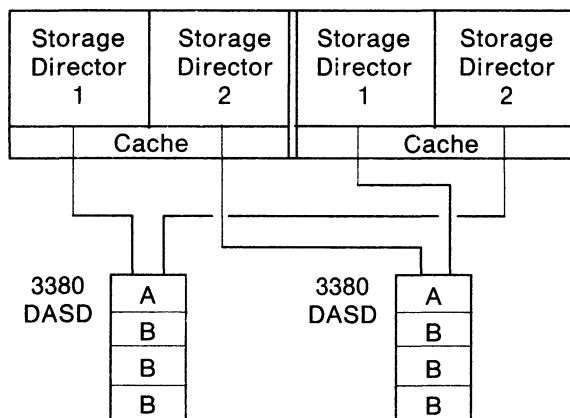
3990 Model 2 Improvements

- 3390 Attachment
- Dual line cord
- Nondisruptive DASD Installation/Removal
- DLS or DLSE Mode (Two or Four Data Transfer Paths to DASD)
- Multipath Storage Directors
- Concurrent Maintenance
- Independent Support Facilities
- Two Independent Storage Clusters
- Uses Less than Half the Floor Space of Two Cross-Configured 3880 Model 3s with Service Clearance
- Faster Microprocessing
- Service Information Messages (SIMs)
- Online SIM Alert Messages
- Faster Access to DASD by Moving DPS Array Into the Storage Control
- Field Upgradable to 3990 Model 3
- Remote Maintenance
- Writeable Diskette
- Reduced Power and Cooling Requirements

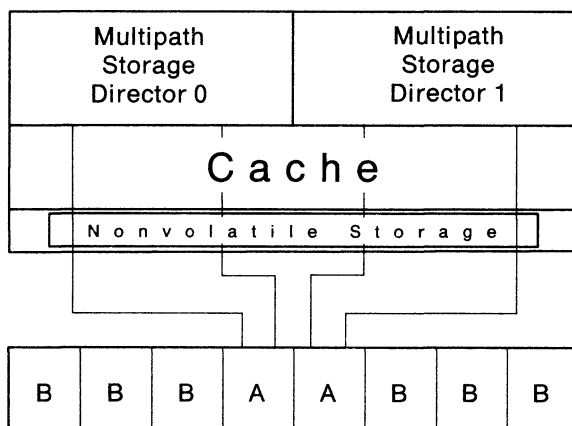
SLVCOTM2

Figure 6. Two Cross-Configured 3880 Model 3s and a Single 3990 Model 2

**A Dual-Framed 3880
Model 23 Storage Control**



**A Single 3990
Model 3 Storage Control in DLSE Mode**



3380 DASD AJ4/AK4

**3990 Model 3
Improvements**

- 32, 64, 128, or 256MB Cache
- 4MB Nonvolatile Storage (NVS)
- DASD Fast Write
- Dual Copy
- Cache Fast Write
- 4.5 Megabyte Per Second Cache Data Transfer Speed
- Dual Data Transfer and Logical Device Allegiance
- More Efficient Use of Cache Storage
- Improved Sequential Operations
- Retained Subsystem Status Across IMLs
- 3390 Attachment
- Dual line cord
- Nondisruptive DASD Installation/Removal
- DLS or DLSE Mode (Two or Four Data Transfer Paths to DASD)
- Multipath Storage Directors
- Concurrent Maintenance
- Independent Support Facilities
- Two Independent Storage Clusters
- Uses Less than Half the Floor Space of a Dual-Frame 3880 Model 23 with Service Clearance
- Faster Microprocessing
- Service Information Messages (SIMs)
- Online SIM Alert Messages
- Faster Access to DASD by Moving DPS Array into the Storage Control
- Field Upgradable to Larger Cache Sizes
- Remote Maintenance
- Writeable Diskette
- Reduced Power and Cooling Requirements

SLVC00TM3

Figure 7. Two 3880 Model 23s in Dual-Frame and a Single 3990 Model 3

Advantages of 3990 Storage Controls

The 3990 may be operated in either of two modes: Device Level Selection (DLS) or Device Level Selection Enhanced (DLSE). Both DLS and DLSE modes allow the 3990 to select the most appropriate and efficient path to any attached device. Fast write and dual copy operations can be used in either DLS or DLSE mode.

DLSE mode fundamentally enhances the performance and data availability characteristics over those of DLS mode.

Device Level Selection Mode of Operation

Device Level Selection (DLS) mode permits two single-path storage directors to access data in the DASD subsystem. In DLS mode, there is a one-to-one relationship between the storage director and the storage path.

DLS provides DASD strings with two paths to each device and simultaneous data transfer to any two devices in the 2-path string (except 3380 Model AA4 or A04).

When the 3990 attaches to 3380 AD4 or AE4, it operates in DLS (2-path) mode. When the 3990 is attached to 3380 AA4, the 3990 operates according to the dynamic path selection capability of the 3380 AA4. When the 3990 is attached to 3380 AJ4/AK4 DASD, and the logical subsystem operates in DLS mode, feature 9432 must be installed on the 3380 AJ4/AK4.

A storage path is a *single-path storage director* in DLS mode. For 3990 Models 2 and 3 in DLS mode, there are two independent logical DASD subsystems within this configuration (3990 Model 1 has only one logical DASD subsystem). As with two 3880s, the operating system addresses each of the four storage paths. Storage directors 0 and 2 address one set of devices, storage directors 1 and 3 address a different set of devices.

Notes:

1. For a description of the features that are required when attaching the different models of the 3380 DASD to the 3990, see "Chapter 6. Optional Features" on page 85.
2. 3390 Models do not attach to a 3990 in DLS mode.

Device Level Selection Enhanced Mode of Operation

Device Level Selection Enhanced (DLSE) mode permits two *multipath storage directors* to access data in the DASD subsystem. Each multipath storage director has two storage paths. In DLSE mode there is a one-to-one relationship between the storage director and the storage cluster, consisting of two storage paths.

DLSE provides DASD strings with four paths to each device and *simultaneous data transfer* to any four devices in the attached 4-path strings. DLSE allows any device to be selected over any of four paths, including those of the same head disk assembly (HDA), instead of either of two paths, as in DLS—resulting in a significant improvement in data throughput and overall subsystem performance.

DLSE is a logical extension of DLS. However, the physical configuration and the mode of operation are somewhat different. In the same way that the channel subsystem relieves an extended architecture (MVS/XA or MVS/ESA) processor of the path management task, so the 3990 DLSE subsystem relieves the channel subsystem (or operating system for non-XA or non-ESA system control programs)

of the need to select the storage path from the storage control to the DASD. This is done by pairing storage paths within a cluster into one multipath storage director. Thus, the host processor addresses the multipath storage director just as it does a 3880 single-path storage director. The multipath storage director allocates the storage path to use for data transfer operations.

Note: The dynamic path reconnection capability of the 3380 or 3390 is only available with XA or ESA system control programs. In a non-XA or non-ESA system, switching can only be done between the two storage paths within a multipath storage director.

When the 3990 Model 2 or 3 is attached to 3380 Models AJ4 or AK4, two modes of operation are possible: DLS or DLSE. DLS mode with 3990 Model 2 or 3 attached to Models AJ4 or AK4 works just as it does when the 3990 Model 2 or 3 is attached to Models AA4, AD4, or AE4. (The 3990 Model 1 operates in DLS mode only.)

When a 2-path and a 4-path subsystem is installed, the 3990 operates in DLSE mode.

DLSE mode is required only when the 3990 is attached to 3390s or 3380 Models AJ4/BJ4 or AK4/BK4 configured as 4-path strings. *DLSE mode permits concurrent and independent data transfer with any four devices within a 4-path DASD string.* Only the Models 2 and 3 can be configured as DLSE subsystems, which can have up to 64 device addresses in a single logical subsystem (twice as many addresses as in a 3880 subsystem).

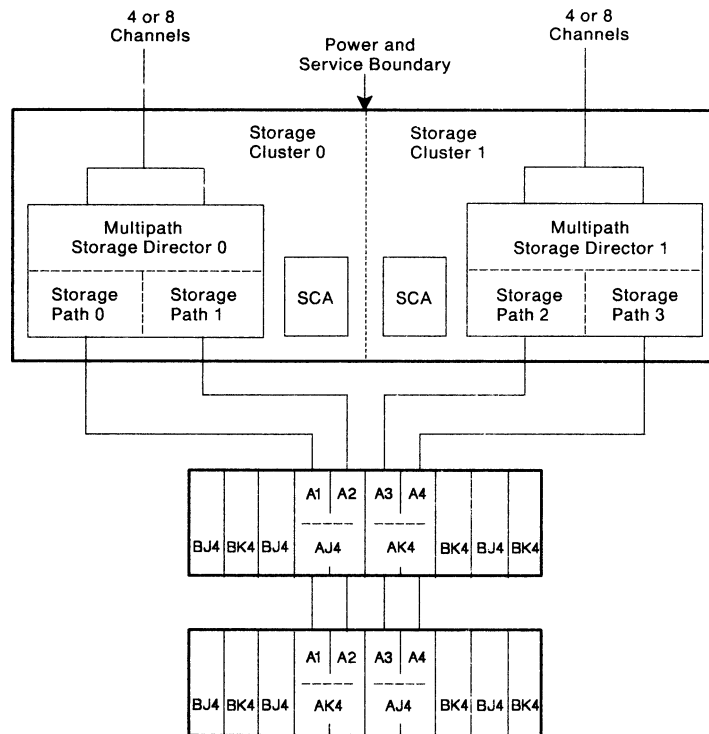
DLSE provides important performance advantages:

- Four independent and simultaneous data transfer paths to the same DASD strings from the 3990 Storage Control. Also, in an XA or ESA environment an I/O operation can be dynamically reconnected on any one of the four paths; in fact, with DLSE there are four completely independent paths to the DASD, substantially improving the availability of the data.
- The performance impact of the loss of one of the four paths is substantially less than the performance impact of the loss of a path in a DLS subsystem. A DLSE storage subsystem, even with one of its four paths disabled, performs better than two DLS subsystems with all of its paths available.
- The decision on which storage-control-to-DASD path to use is determined at the storage control level, eliminating possible delays if the selected path is busy.

Figure 8 shows the physical connections of a 3990 Model 2 or 3 Storage Control in DLSE mode attached to 3380 Models AJ4 or AK4 4-path strings. The 3990 Model 2 or 3 can attach one or two 4-path DASD strings or one 4-path DASD string and one or two 2-path DASD strings.

The DLSE configuration requires feature 9433 be installed on the 3380 AJ4 or AK4.

Figure 8 shows a 3990 Model 2 or 3 in DLSE mode.



Legend: SCA = Shared Control Array

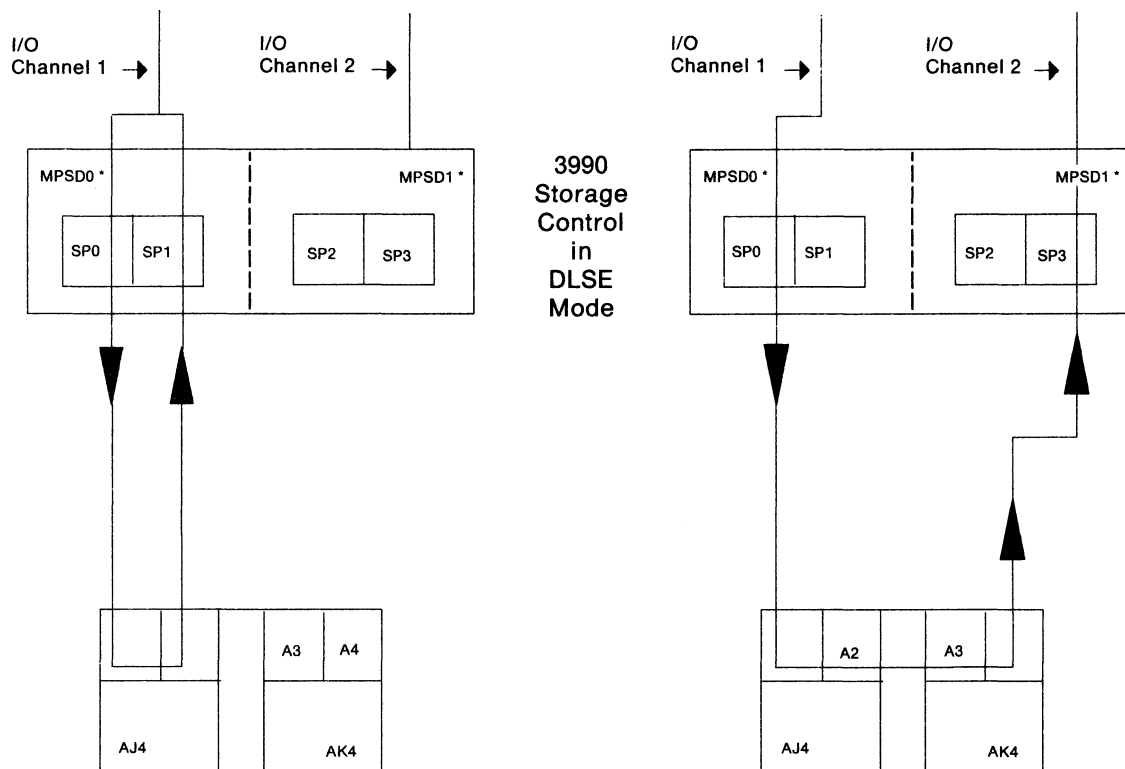
8/12/2013

Note: The figure shows a maximum 3380 configuration of four A-units and 12 B-units. A minimum 3380 configuration consists of two A-units in the same string.

Figure 8. 3990 Model 2 or Model 3 in DLSE Mode Attached to 3380 Models AJ4 or AK4 4-Path Strings

Figure 9 shows an example of an I/O operation of the 3990 in DLSE mode in S/370, 370/XA, and ESA/370 environments. In S/370, DLSE allows an I/O operation to complete on either of the storage paths in the multipath storage director that initiates the I/O operation. In 370/XA or ESA/370, DLSE allows an I/O operation to return over any of the available paths.

Note: For a description of the features that are required when attaching the different models of the 3390 or 3380 DASD to the 3990, see "Chapter 6. Optional Features" on page 85.



Note: S/370 Input/Output MUST start and complete on the same channel path.

Note: 370/XA or ESA/370 Input/Output may start and complete on any path.

* MPD = Multipath Storage Director

Figure 9. Input/Output Operation Example, 4-Path DASD (DLSE)

SLVCOTMA

Fast Write

There are two types of fast write operations: **DASD fast write** and **cache fast write**. Both fast write operations can improve performance for write hits or full-track format write operations. Fast write operations can be used in either DLS or DLSE mode.

DASD Fast Write

DASD fast write improves storage subsystem performance because immediate access to DASD is not required for write hits and full track format writes. DASD fast write stores data simultaneously in cache and in nonvolatile storage (NVS).

Using DASD fast write requires no modifications to access methods or applications. Access to DASD is not required to complete the DASD fast write operation for write hits. Because a copy of the data is put into the NVS, the storage director returns channel-end and device-end status together at the end of the data transfer to cache and NVS. This allows the program to continue processing without waiting for the data to be put on DASD. The data remains in cache and in NVS until the data is written to DASD to free space in the cache or NVS. Most write operations operate directly with the cache without going to the DASD, resulting in the same performance as a read hit operation.

Figure 10 shows the **DASD fast write** operation.

DASD Fast Write - Improved Data Performance

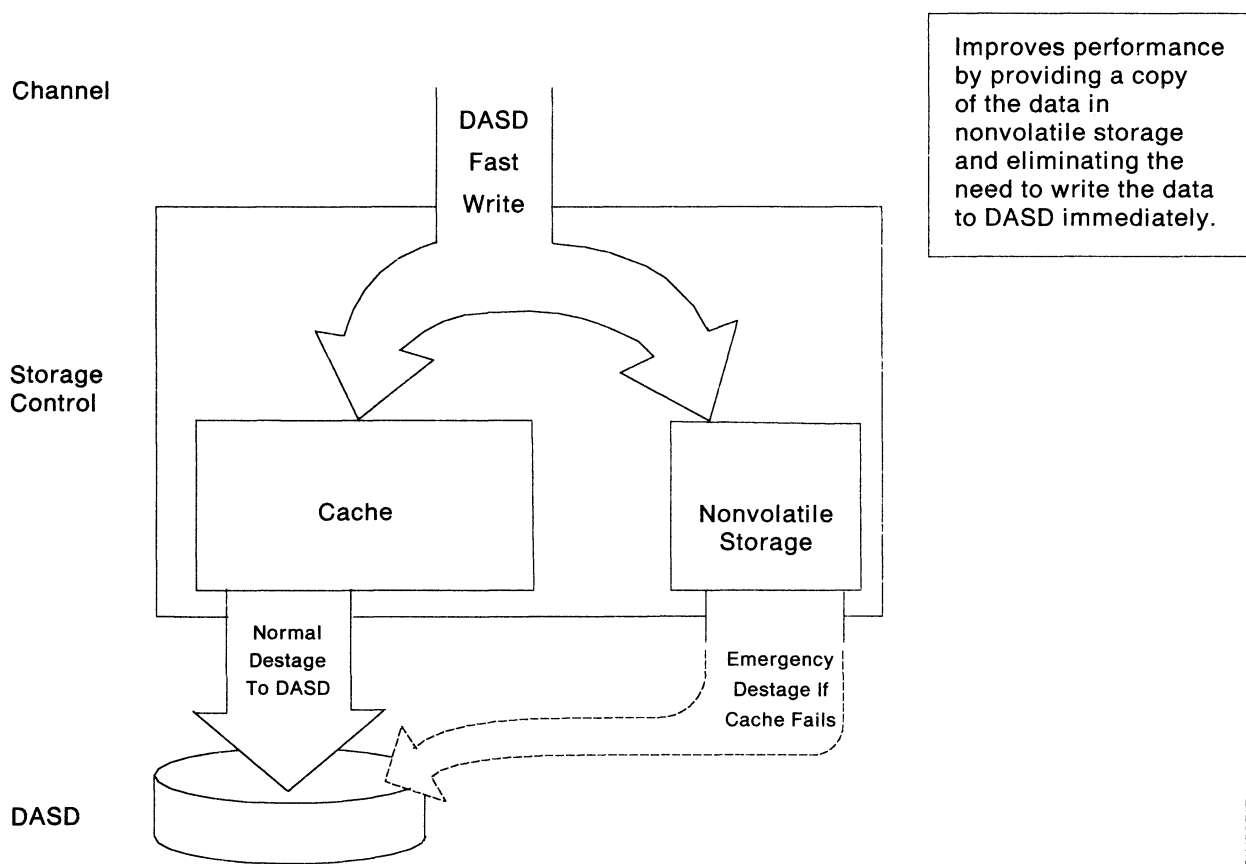


Figure 10. DASD Fast Write Operation

Cache Fast Write

Cache fast write is an option designed for use with special kinds of data, such as a DFSORT sort work file. When DFSORT sort work files are allocated on a 3990 Model 3, the files are kept temporarily in the cache. The data might not be written to the DASD. The channel program is complete when the data is written to cache. Cache fast write does not use the nonvolatile storage. DFSORT Release 10 and higher supports cache fast write.

Cache fast write operations are program specific. Each channel program invoking cache fast write uses a cache fast write ID and turns on a cache fast write indicator in the Define Extent command. If cache fast write operations are permitted, only the cache copy of the data track is updated during the write operation. If data is lost because of a cache or subsystem failure, the using job is notified and the job can be rerun if the user desires.

Figure 11 shows the **cache fast write** operation.

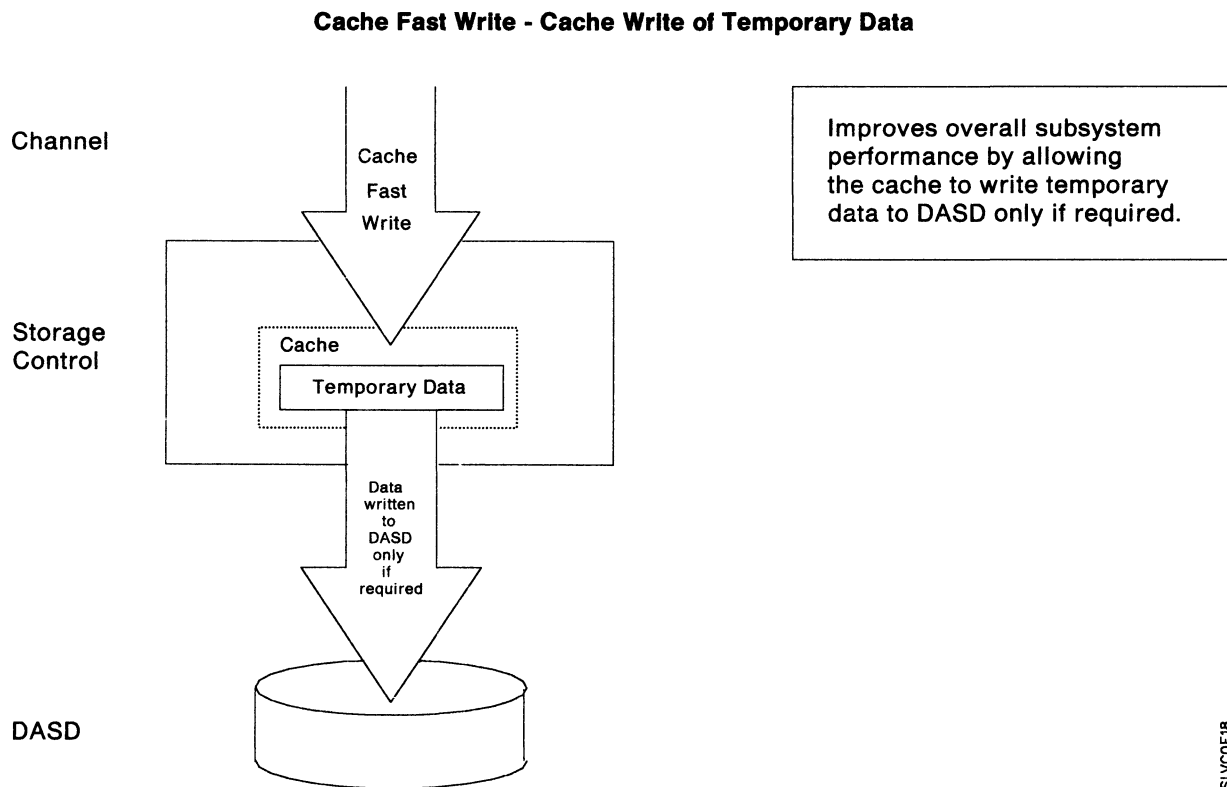


Figure 11. Cache Fast Write Operation

Dual Copy

Dual copy allows the Model 3 to create and maintain a duplicate copy of the data on a volume and store it on a different device in the same subsystem.

Dual copy provides important advantages:

- Protects a critical volume from a single device failure.
- Is transparent to the application program.
- Switches operations automatically to the secondary if a failure occurs on the primary. The switching is transparent to the application.
- Provides for near-continuous operations by providing a back-up copy if loss of access occurs to one of the dual copy devices.
- Ensures that the second copy is logically identical to the primary copy.
- The 3990 Model 3 creates the secondary copy of the dual copy volume within the subsystem, without host I/O for the copy operation in order to initially synchronize the two devices.

Dual copy improves the availability of data. The status of the dual copy operation is kept in the nonvolatile storage. Dual copy is activated or deactivated by a system utilities command. Using dual copy should require no modifications to access methods or applications. See "Chapter 5. Using Dual Copy, DASD Fast Write, and Cache Fast Write in MVS Environments" on page 63 for a further discussion about dual copy.

The two physical devices are a **duplex pair**; a primary device and a secondary device. The dual copy operation is managed by the subsystem. All I/O operations are directed to the primary volume. The 3990 Model 3 automatically updates both copies of the data. Data is accessed from the secondary device if the primary device is not available. The duplex pair is established by a system command. The secondary device is automatically synchronized with the primary when both devices are available to the subsystem.

Because the secondary device is offline, the processor knows of only one device—the primary device. After the processor writes the data to the primary device, channel-end and device-end are presented and the storage control later completes the write operation from the cache to the secondary device. The write operation to the secondary device is transparent to the host.

The DASD fast write capability and the dual copy capability can be combined to produce a **fast dual copy**. See "Chapter 5. Using Dual Copy, DASD Fast Write, and Cache Fast Write in MVS Environments" on page 63 for a discussion of considerations for using dual copy and dual copy performance.

Any two devices may be built into a duplex pair provided that:

- Both devices have the same track and cylinder format, and are the same size
- Both devices are within the same logical DASD subsystem.

In particular, 3380 Models A, D, and J can be made into a duplex pair; a 4-path 3380 AJ4 and a 2-path 3380 AA4 or AD4 can also be in a duplex pair. For all other DASD, both devices must be the same model and operate in the same mode, such as 3390 mode or 3390s in track compatibility mode. In particular, two 3390 Model 1s or two Model 2s in 3390 mode may be in a duplex pair; or two 3390 Model 1s or two Model 2s in 3380 track compatibility mode may be in a duplex pair.

IBM recommends that the primary device and the secondary device be in separate strings.

Figure 12 shows the dual copy operation.

Dual Copy - Improved Data Availability

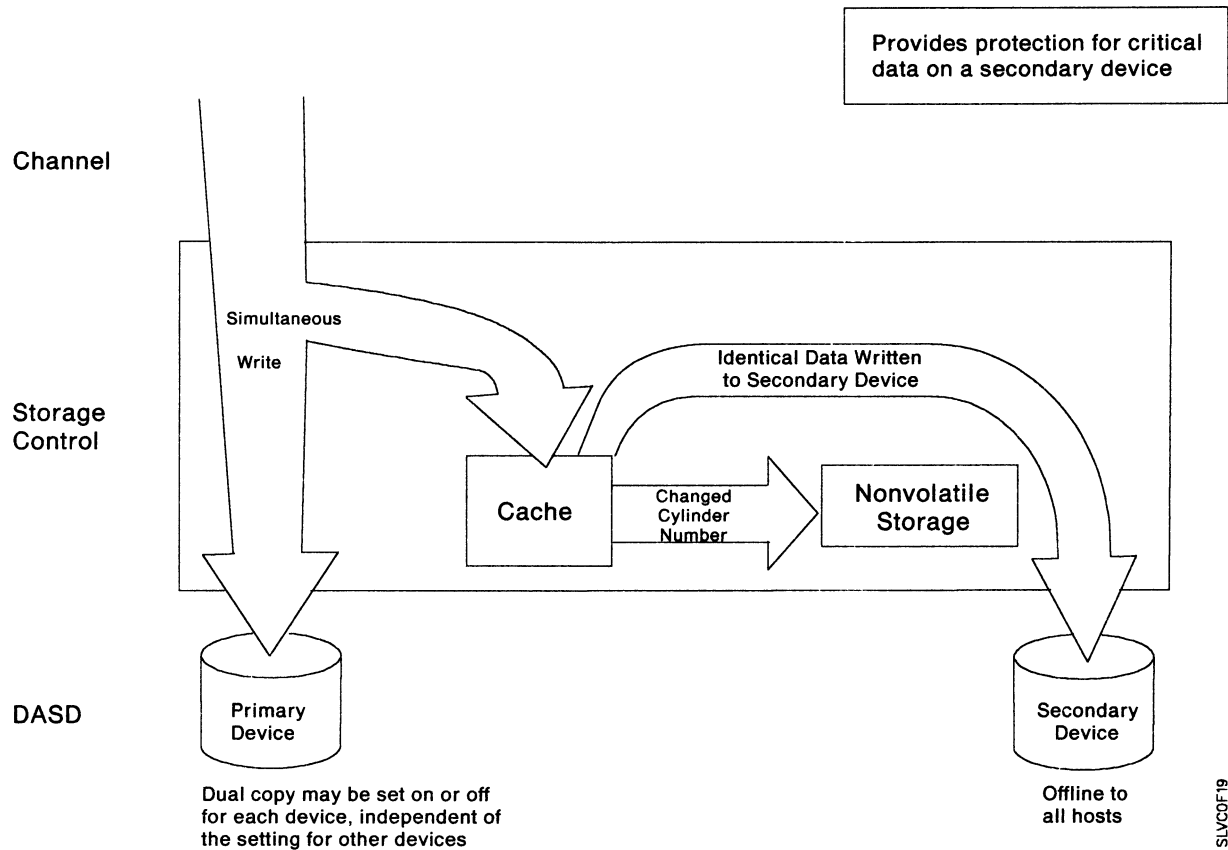


Figure 12. Dual Copy Operation

When a read error occurs on the primary device (permanent data check), the Model 3 will retrieve (stage) the data from the secondary device. The failure is reported by a console message so that the installation can take action to correct it later.

When a write request error occurs and the primary device has failed, the secondary device is made the operational device. All further write requests are directed to the secondary device as it now becomes the primary device. The failure is reported by a console message so that the installation can take action to correct it later.

For more information on the DASD fast write, cache fast write, or dual copy operation, see the *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*.

Branching Transfers, Dual Data Transfer, and Logical Device Allegiance in the 3990 Model 3

Branching Transfers

The 3990 Model 3 uses branching transfers to permit simultaneous transfer of data from the channel to DASD and the cache, and from the DASD to the channel and cache. Branching transfers provide a substantial reduction in storage control and device busy time for certain operations and an overall improvement in performance.

Dual Data Transfer

Dual data transfer allows up to four simultaneous and independent data transfer operations on lower ports and at the same time, up to four simultaneous and independent data transfer operations on the upper ports. Up to eight operations are permitted in the 3990 Model 3 (only two data transfer operations are permitted at any one time in the 3880 Model 23).

Logical Device Allegiance

Logical device allegiance permits the 3990 Model 3 to complete operations without selecting the physical device if the request does not have to access the physical device. Host access to cache data may be concurrent with stage and destage activities to the same volume.

Advantages

All three capabilities, branching transfers, dual data transfers, and logical device allegiance permit the 3990 Model 3 to perform more tasks simultaneously, resulting in greater performance for the user.

Nonvolatile Storage

The nonvolatile storage (NVS) provides random-access electronic storage for 4MB of data. The NVS has its own separate power region to facilitate concurrent maintenance. If power is lost to the 3990 before DASD fast write and/or dual copy operations have completed to DASD, a battery-backup system maintains power in NVS for up to 48 hours with a fully-charged battery to prevent data loss. When power is restored, the Model 3 destages any data in NVS to DASD.

3990 Storage Control Components

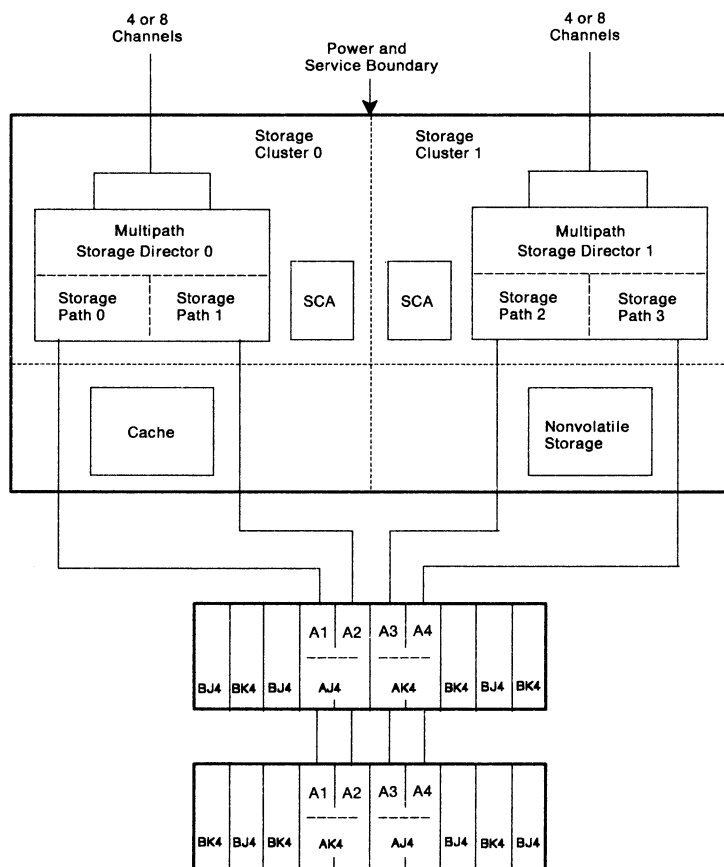
To improve data availability and serviceability, the major components of the 3990 are grouped into independent power and service regions. These regions include the storage clusters, the cache, and the NVS. The 3990 Model 1 has one storage cluster, the 3990 Models 2 and 3 have two. Each storage cluster contains channel attachments, storage directors, storage paths, a shared control array, and a support facility for maintenance. In a 3990 Model 3, the cache and the NVS are shared by all the storage paths in both clusters, but are physically and logically separate from the storage clusters.

A clear benefit of each of the clusters is that they are independent components. Each provides a separate power and service region and two separate paths to the DASD. Because DASD strings are attached with at least one path in each cluster, loss of power to one cluster, the cache, or the NVS, does not prevent access to data as processing can continue through the other storage cluster. Depending on

the 3990 configuration, two paths in DLSE mode or one path in DLS mode would remain available for access to the DASD should one of the storage clusters become unavailable. A service representative can take one cluster for maintenance activities without inhibiting DASD access through the remaining paths.

The 3990 Model 3 contains cache and nonvolatile storage (NVS); each is in a separate power and service region. Cache and NVS are accessed by both storage clusters. Cache and/or NVS service actions can take place while I/O activity continues through the storage clusters. A NVS service action can occur while the two storage clusters continue to operate with the cache and DASD.

Figure 13 shows the power and service boundaries of a Model 3 that allow operations to continue on the remaining paths to DASD should a cluster, the cache, or the NVS become unavailable.



Legend: SCA = Shared Control Array

slv20n36

Figure 13. Power and Service Boundaries of the 3990 Model 3

Channels

Each storage cluster can be attached to either four or eight channels. For 3990 Models 2 and 3, each channel may be connected to the two storage paths in the cluster. Therefore, a processor can have several paths to data, or each of several processors can have its own path (or paths) to data. For 3990 Model 1, each channel attaches to one storage path by being assigned to one storage cluster or the other.

Storage Directors

The storage directors interpret channel commands and control the storage paths, cache, NVS, and attached DASD.

In DLS mode, the storage director is called a single-path storage director and contains one storage path.

In DLSE mode, the storage director is called a multipath storage director and contains both storage paths in a cluster.

Each storage director has a unique channel address.

Storage Paths

Each storage cluster contains two storage paths. Each storage path separately connects to DASD. During channel-connected operations, the storage path is coupled with a particular channel.

In DLS mode, there is a one-to-one relationship between the storage director and the storage path. In 3990 subsystems, DLS mode defines the storage paths as single-path storage directors. Single-path storage directors are functionally equivalent to 3880 storage directors. Each single-path storage director provides, through a single channel address, a path to DASD. In DLS mode, there are two storage paths to each device in a 2-path string.

In DLSE mode, there is a one-to-one relationship between the storage director and the storage cluster (consisting of two storage paths). In 3990 subsystems, DLSE mode defines the storage clusters as multipath storage directors. The multipath storage director provides, through a single channel address, multipath access to DASD. Through one storage director address, the multipath storage director selects either storage path in the cluster for data transfer operations, thereby improving response time. In DLSE mode, there are four storage paths to each device in a 4-path string.

Shared Control Array

The shared control array (SCA) is electronic storage that contains information about the status of the 3990 Storage Control and the attached DASD. In the Model 2 or 3, identical information is kept in the shared control array of each cluster. In dual-frame configurations, the shared control array information is duplicated in the two clusters that are paired together.

The shared control array information formerly was kept in the dynamic path selection (DPS) array of the 3380 Models AA4, AD4, and AE4. Having this information within the 3990 permits the storage director to obtain data faster.

Support Facility

Each independent storage cluster has its own support facility.

A major reliability, availability, and serviceability (RAS) benefit of the 3990, the support facility and the independent service regions permit concurrent maintenance. The support facility also provides a remote maintenance support capability. Among other tasks, the support facility generates the service information messages, communicates with the other support facility (in the Models 2 and 3), runs the maintenance analysis programs (MAPs) and diagnostics, maintains logs for the storage cluster, and provides a nonvolatile microcode patch area on its writeable diskette.

The concurrent maintenance and the remote maintenance support capabilities of the support facility (in each independent storage cluster) in combination with the service information messages should meaningfully reduce the time it takes to diagnose and perform maintenance and migration activities.

Each storage cluster contains a microprocessor-controlled support facility that communicates with:

- A diskette drive
- The cluster storage paths
- A remote support adaptor
- The cluster power sequence controls (dual-frame capable only)
- The support facility in the other cluster (in Models 2 and 3)
- The operator panel
- The maintenance panel.

The support facility provides the following reliability, availability, and serviceability functions:

- Monitors power to the storage cluster (cache and NVS power in Model 3)
- Maintains a record of errors and reports those errors that exceed established thresholds for the storage control and 3390 DASD
- Detects, logs, and reports problems in the other cluster's support facility (in two-cluster storage controls)
- Provides maintenance analysis procedures (MAPs) for a service representative
- Permits support personnel to analyze certain machine conditions from a remote location through an external modem
- Disables paths determined to be defective
- Provides nonvolatile microcode update capability
- Provides check-1 error recovery for the storage paths.

Cache

Cache is high-density, electronic storage in the 3990 Model 3 that is shared by all storage paths. Frequently used data can be transferred to and from the cache and the channel at channel speeds. Access time between the cache and the channel is much faster than between the DASD and the channel because there are no DASD seek or rotational delays. A least-recently-used (LRU) algorithm (and other algorithms) keeps high-activity data in the cache.

Cache is divided into 16KB segments. When a track is brought into the cache, the appropriate number of slot segments (3 or 4) is allocated. Any slot segments not needed to contain required records are freed. Further, the 3990 Model 3 can logically relate separated segments in the cache and treat them as a unit of data.

Cache is in a separate power region from the storage clusters. If a storage cluster is offline, cache processing still continues through the other storage cluster.

Data transfers between the cache and the channel operate at the maximum speed of the channel, either 4.5 or 3.0MB per second. All data transfers requiring access to the physical DASD will occur at the DASD data transfer rate, 4.2MB per second for 3390s or 3.0MB per second for 3380s.

Branching data transfers, dual data transfers, and logical device allegiance provide a substantial improvement in storage control utilization, a reduction in device busy time and an overall improvement in performance.

Cache Operations

With the 3990 Model 3, there are three different caching operations:

- Subsystem caching (read only caching)
- DASD fast write
- Cache fast write.

Subsystem caching is a prerequisite for DASD fast write and cache fast write. With subsystem caching, only read operations benefit from cache. With DASD fast write or cache fast write, the performance benefits of caching read operations are extended to write operations.

Read Caching Operations

If a copy of the data is in the cache when the processor initiates a read request (read hit), the storage control transfers the desired data from the cache to the channel. If a copy of the data is not in the cache (read miss), the storage control sends the requested data directly to the channel from the DASD and, at the same time, moves that data (plus the rest of the data from that record to the end of the track) into the cache in anticipation of future use. Requests for following records on that track are read from the cache and are read hits.

Basic Write Caching Operations

If a copy of the data is in the cache when the channel initiates a write request (write hit), the storage control writes the data directly to the DASD and, at the same time, writes that data into the cache. The record in cache is updated because it may be referred to again. This is an example of branching data transfers. Writing to the DASD provides data integrity because copies in the cache and on the DASD are identical.

When the record being updated is not found in the cache, the condition is called a write miss. The record is written directly to the DASD, and is NOT written in the cache.

Fast Write Operations (With Either DASD Fast Write or Cache Fast Write)

The cache management algorithms for both fast write functions are the same. They differ in the way they are invoked, and in the fact that DASD fast write also uses the NVS to provide protection against power failures, but they both process read and write hits and misses the same way.

Most write operations are write hits because typical applications read a record before updating it, or the write operation itself formats the full track. This last operation is called a full-track format write. In this write operation, the new record is written, and the rest of the track is formatted for new data. Thus, there is no need to verify the data on the track before allowing the cache write. Full-track format writes are considered cache hits.

For either DASD fast write or cache fast write operations, if a copy of the data is not in the cache when the channel initiates a write request or the operation is not a format write the storage control writes the data to DASD and the cache. The remainder of the track is staged into the cache.

For DASD fast write, all writes from the channel to the cache are also stored in the NVS.

Both cache fast write and DASD fast write operations are explained in "Fast Write" on page 34.

Caching Algorithms

Several caching algorithms determine how the cache is managed for a specific type of I/O operation. Caching algorithms include normal, sequential, bypass-cache, and inhibit cache loading. These algorithms are parameters in the Define Extent command and are used for the remainder of the I/O operation (a channel program).

The cache uses normal caching algorithms unless directed otherwise by software. Normal caching means that the data is staged to the cache after being referred to in a read operation and remains in the cache until least-recently-used (LRU) algorithms permit the data to be overlaid by other data. Normal caching applies to read operations and both fast write operations. This is the most appropriate algorithm for random access methods, including the Basic Direct Access Method (BDAM), the keyed direct processing of Virtual Storage Access Method (VSAM) key-sequenced data sets (KSDS), the Partitioned Access Method (PAM), and the Overflow Sequential Access Method (OSAM).

Sequential access methods such as Basic Sequential Access Method (BSAM), Queued Sequential Access Method (QSAM), and sequential processing of VSAM entry sequenced data sets invoke sequential caching.

Sequential caching of a channel program is determined by a bit in the Define Extent command, just as it was in the 3880 Model 23. For the 3990 Model 3, the sequential algorithm is improved to provide more data in the cache and to keep the data in the cache longer. These improvements to sequential caching greatly increase hit ratios resulting in better performance.

During sequential caching, the 3990 Model 3 pre-stages data so that, usually, four tracks are in the cache.

Bypass-cache does not use the cache and the I/O operation goes directly to the DASD. The IBM program, Device Support Facilities (ICKDSF) uses bypass-cache to direct operations to DASD.

Inhibit cache loading is set by the Define Extent command and uses existing copies of tracks if they are in the cache, but does not load any new tracks into the cache. The purpose of this cache mode is to avoid the overhead that may be caused by track promotion when there is no benefit (that is, whenever these tracks are not likely to be accessed again in the near future).

Inhibit cache loading is not used by any of the MVS or VM access methods but Data Facility Data Set Services (DFDSS) uses inhibit cache loading.

Channel Commands

The 3990 command set supports the Count, Key, and Data (CKD) data format; and the CKD and Extended CKD (ECKD) command sets as used for 3390 and 3380 disk storage operations. The command sets are described in the *IBM 3990 Storage Control Reference*. For further discussion, see also *Introduction to Nonsynchronous Direct Access Storage Subsystem*.

The channel command set includes commands that support dynamic path selection, cache control, and 3990 statistical and status information.

Hardware and Data Availability Considerations

The 3990 Model 2 and Model 3 provide higher levels of hardware and data availability.

These new levels of hardware and data availability are provided by a number of fundamental technological enhancements:

- Up to four data paths to DASD
- Standard two-cluster design provides internal hardware redundancy
- Use of fewer components, cards and connections
- Improved caching technology and algorithms
- Independent power and service regions (the two clusters, the cache, and the nonvolatile storage)
- Concurrent maintenance
- Writeable diskette
- Dual line cord feature or dual-frame configurations
- Nondisruptive DASD installation and removal
- Control information shared by both clusters
- Enhanced error checking circuitry.

The 3990 Model 3 substantially improves DASD availability through use of dual copy.

The 3990 Model 3 continues several important availability capabilities used with the 3880 Model 23 storage control, among them:

- For cache, single, double, and triple bit error detection for all errors; single and double bit error correction for all errors, and triple bit error correction for most errors.
- Automatic reinitializing of the cache after a cache directory error.

These improvements result in a new level of hardware and data availability of DASD subsystems. The basic 3990 storage control technology, with its circuit redundancy, data path redundancy, and multiple power and service regions reduces the single points of failure in a 3990 Model 2 to three points, all associated with primary power distribution: a circuit breaker, line filter, and the power cord. These parts have such a low failure rate that many installations will be able to configure their DASD subsystems using a single 3990 Model 2. For those installations that desire to protect against the failure of even these highly reliable parts, the dual line cord feature is available. For 3990s with dual-frame capable power subsystems, the dual-frame configuration is available.

A properly configured 3990 Model 3 with the dual line cord feature, or a pair of 3990 Model 3s in a dual-frame configuration, eliminates all single points of failure that may prevent access to data on dual copy pairs— from processor channels to storage control to dual copy volumes. For maximum data availability, it is highly recommended that all 3990 Model 3 subsystems using dual copy or DASD fast write have either the dual line cord feature or be dual-frame configurations.

A single 3990 Model 3 with a dual-frame capable power subsystem has common circuitry between the cache and the nonvolatile storage. In the remote event of a failure of this component, temporary loss of access to data could occur until this hardware component is replaced. Of course, the data would always be protected by the nonvolatile storage. The dual-frame configuration, because of the way it cross-configures the cache and the nonvolatile storage of the two 3990 Models 3s, duplicates the common circuitry. 3990s with modular power subsystems do not have common circuitry between the cache and the nonvolatile storage.

When the dual line cord feature is installed on a 3990 Model 2 or 3 with the dual-frame capable power subsystem, the power subsystem is replaced with the modular power subsystem and the second line cord is added. 3990s with modular power subsystems cannot be dual-framed.

Note: The dual line cord feature cannot be installed on a 3990 Model 1 until the 3990 Model 1 is upgraded to a Model 2 or 3.

Figure 14 on page 46 shows a modular power subsystem with both a single and a dual line cord.

Modular Power Subsystem

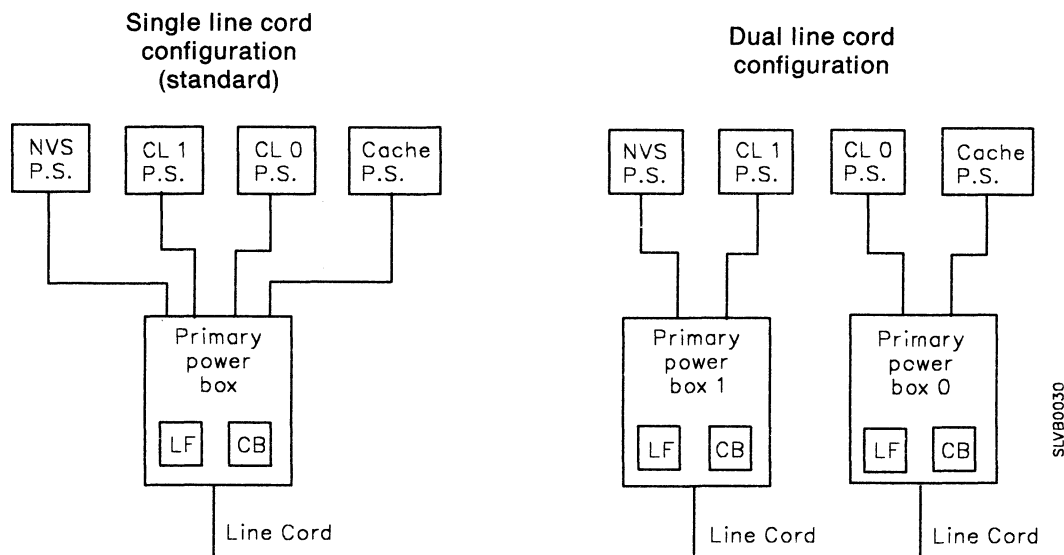


Figure 14. Modular Power Subsystem

For a 3990 Model 2 or 3 with the dual line cord feature, cluster 0 receives power from one primary power supply, and cluster 1 receives power from the other. In a 3990 Model 3, the cache receives power from the same primary power supply as cluster 0 and the NVS receives power from the primary power supply for cluster 1.

Processor Attachment

The 3990 attached to 3380 DASD can be attached to the following processors:

- 3090, 3090E, 3090J, 3090S,
- 3081, 3083, 3084
- 4381
- 9375, 9377.

3990s attached to 3390 DASD can be attached to the following processors:

- 3090
- 3090E
- 3090J
- 3090S
- 308x ¹

Note: ¹ – 308x processors require processor RPQs which increase channel speed to 4.5MB per second (the 3081 Model D does not support this upgrade).

For availability of required features, see your marketing representative. For environment-specific information on any support restrictions, software tools which support 3990 models, and processor support see:

- *Using IBM 3390 Direct Access Storage in an MVS Environment*
- *Using IBM 3390 Direct Access Storage in a VM Environment*
- *Using the IBM 3380 Direct Access Storage in an MVS Environment*
- *Using the IBM 3380 Direct Access Storage in a VM Environment*
- *Using the IBM 3380 Direct Access Storage in a VSE Environment.*

DASD Attachment

The 3990 must be configured with two storage paths in a subsystem in DLS mode and four storage paths in a subsystem in DLSE mode. Every DASD in the subsystem can be accessed by each storage path of the subsystem (except for 2-path DASD in DLSE intermixed configuration). Multiple paths in a subsystem provide alternate path capability from the processor for both data availability and performance.

The 3990 attaches IBM 3380 Models AA4 (except devices with serial numbers 10001 through 14999 for 60 Hz units or 10001 through 11818 for 50 Hz units), AD4, AE4, AJ4, and AK4. The 3990 does not attach to 3380 Model A04, Model CJ2. The 3990 also does not attach to 3380 Model AA4 that have been converted from Model A04 with serial numbers 80000 to 84999.

The 3990 Models 2 and 3 attach 3390 Models A1x and A2x; however, the 3380 Model AA4 may not be attached to a 3990 that has Feature 6120 installed, even if 3390s are not attached.

Table 4 shows which DASD models can be intermixed on the same string.

Table 4. IBM DASD Models that can be Intermixed on the Same String							
DASD Type	AA4	AD4	AE4	AJ4	AK4	3390 Model A1x	3390 Model A2x
3380 Model B04	Yes	—	—	—	—	—	—
3380 Model BD4	—	Yes	Yes	—	—	—	—
3380 Model BE4	—	Yes	Yes	—	—	—	—
3380 Model BJ4	—	—	—	Yes	Yes	—	—
3380 Model BK4	—	—	—	Yes	Yes	—	—
3390 Model B1x	—	—	—	—	—	Yes	Yes
3390 Model B2x	—	—	—	—	—	Yes	Yes
Note: Yes = Models can be intermixed on the same string. — = Models cannot be intermixed on the same string.							

Table 5 shows the combinations of 3390 models and 3380 models, and their capabilities (2-path or 4-path) that can be intermixed in the same logical DASD subsystem. If there are two strings of 3380s in the logical DASD subsystem, both strings must attach to the same storage paths, except when 2-path strings are intermixed with a 4-path string. In this case, one 2-path string attaches to two of the storage paths, and the other 2-path string attaches to the remaining two storage paths.

Table 5. DASD Models and Capabilities that can be Intermixed in the Same Subsystem						
Models and Capabilities	2-Path String of 3380 Model A	2-Path String of 3380 Models D and E	2-Path String of 3380 Models J and K	4-Path String of 3380 Models J and K	3390 Model A1x Strings	3390 Model A2x Strings
2-Path Strings of 3380 Model A	Yes	Yes	Yes	Yes	—	—
2-Path Strings of 3380 Models D and E	Yes	Yes	Yes	Yes	Yes	Yes
2-Path Strings of 3380 Models J and K	Yes	Yes	Yes	—	Yes	Yes
4-Path Strings of 3380 Models J and K	Yes	Yes	—	Yes	Yes	Yes
3390 Model A1x Strings	—	Yes	Yes	Yes	Yes	Yes
3390 Model A2x Strings	—	Yes	Yes	Yes	Yes	Yes
Note: Yes = Strings can be intermixed on the same logical DASD subsystem. — = Strings cannot be intermixed on the same logical DASD subsystem.						

Note: 3390 strings cannot be intermixed with 3380 Models AA4 in the same logical DASD subsystem. 3380 Models AJ4/AK4 cannot be in a 4-path and a 2-path configuration in the same logical DASD subsystem.

Chapter 3. 3990 Storage Control Models and Configurations

3990 Model 1

The 3990 Model 1 Storage Control attaches to 3380 DASDs in installations requiring only two storage paths to the DASD. In a 3990 Model 1, the two storage paths form one storage subsystem operating in DLS mode. In a 3990 Model 1, there is a one-to-one relationship between the storage path and the storage director.

The 3990 Model 1 has one storage cluster with two storage paths (storage directors) to DASD and a shared control array containing control and DPS information. Each storage director is an addressable storage path to DASD. Four channels can attach to the 3990 Model 1. One channel cannot attach to both storage directors. With the Four Channel Switch, Additional feature, eight channels can attach to the 3990 Model 1. The 3990 Model 1 has one support facility for maintenance activities.

The shared control array (SCA) contains information and status about the availability of the storage paths between channel and DASD and the status of the devices.

The 3990 Model 1 does not provide the dual line cord feature, dual-frame configurations, DLSE mode capability, nondisruptive DASD installation and removal, concurrent maintenance, cache, or attachment to 3390 DASD strings or 4-path 3380 DASD strings.

The 3990 Model 1 can be upgraded to a Model 2 and then to a Model 3 with cache and nonvolatile storage.

In the 3990 Model 1, there is no requirement that the channels be evenly divided between the storage directors and the two storage directors do not share channels. Each channel can attach to either storage director, but not to both.

The 3990 Model 1 cannot cross-connect to another Model 1. If cross-connection is desired, a single 3990 Model 2 provides cross-connection to all attached devices.

Figure 15 shows the major components of a 3990 Model 1.

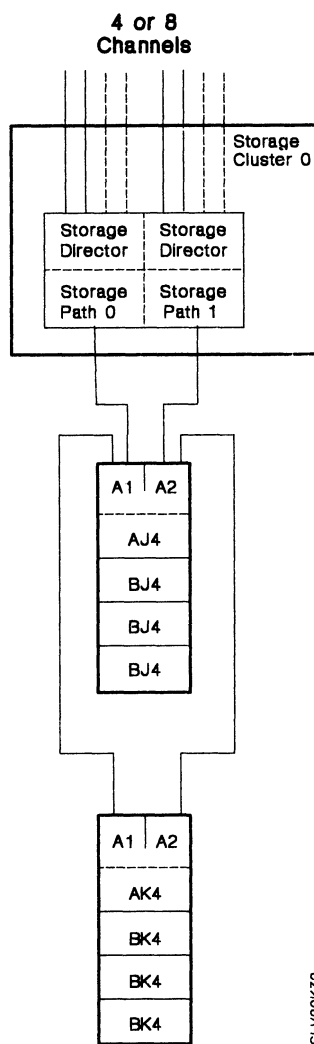


Figure 15. 3990 Model 1 in DLS Mode Attached to 2-Path DASD Strings

Notes:

1. A single 3990 Model 1 attaches to four channels. In Figure 15, the dotted lines show the extra channels added by the Four-Channel Switch, Additional feature.
2. For more examples of the different configurations of the 3990, see the *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*.
3. See "Chapter 6. Optional Features" on page 85 for descriptions of the optional features available on a 3990 Model 1.

3990 Models 2 and 3

The 3990 Models 2 and 3 operate in either DLS or DLSE modes. Each 3990 Model 2 or 3 has two storage clusters. Each storage cluster has two paths to the DASD. A 3990 Model 2 does not have cache or NVS. A 3990 Model 3 has either 32, 64, 128, or 256 megabytes of cache and 4 megabytes of nonvolatile storage.

In any 3990 in DLS mode, there is no requirement that all channels be connected to both storage directors within a cluster. In any 3990 in DLSE mode, there is one multipath storage director per cluster. Each channel accesses both storage paths within the multipath storage director. The installation has the flexibility to connect either a DLS single-path storage director or a DLSE multipath storage director to any of four or eight channels.

Figure 16 shows the maximum configuration of a 3990 Model 3 with two 4-path DASD strings.

The 3990 Model 3 DASD attachment is the same as the 3990 Model 2 and all storage paths attach to cache and to the nonvolatile storage.

For high data availability, the two storage paths in a subsystem in DLS mode are not in the same storage cluster. One storage path is in storage cluster 0 and the other is in cluster 1. In DLSE mode with four storage paths, two storage paths are in storage cluster 0 and the other two are in cluster 1.

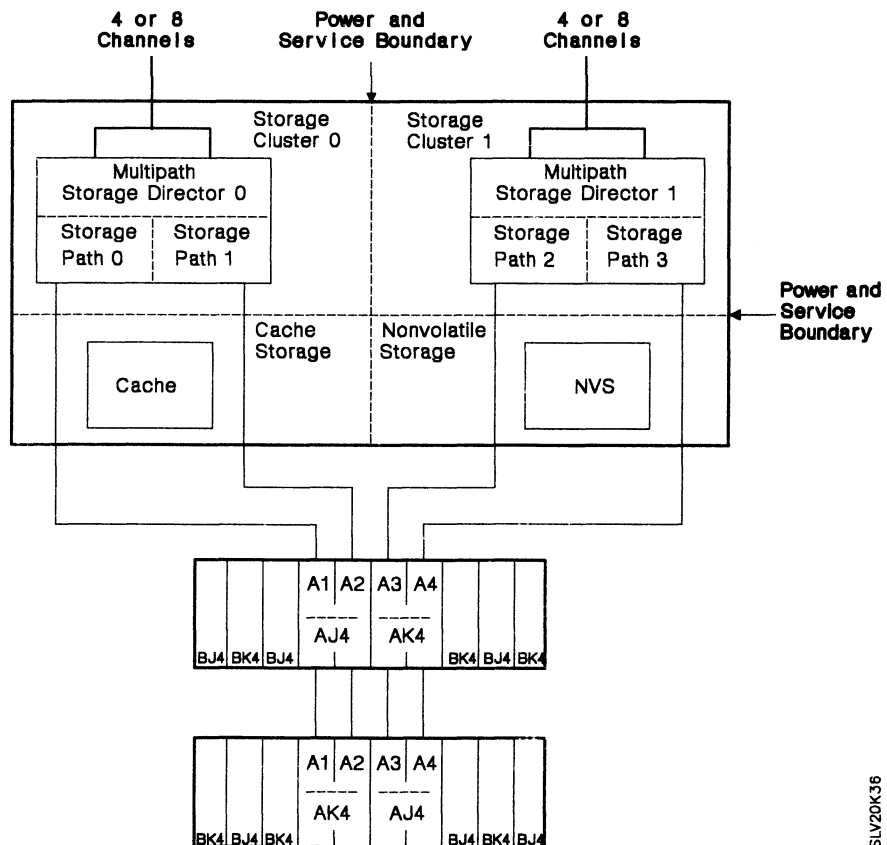


Figure 16. 3990 Model 3 in DLSE Mode Connected to 3380 Models AJ4 or AK4 4-Path Strings

Note: For more examples of the different configurations of the 3990, see the *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*.

Figure 17 shows the 3990 Model 2 with two logical DLS subsystems.

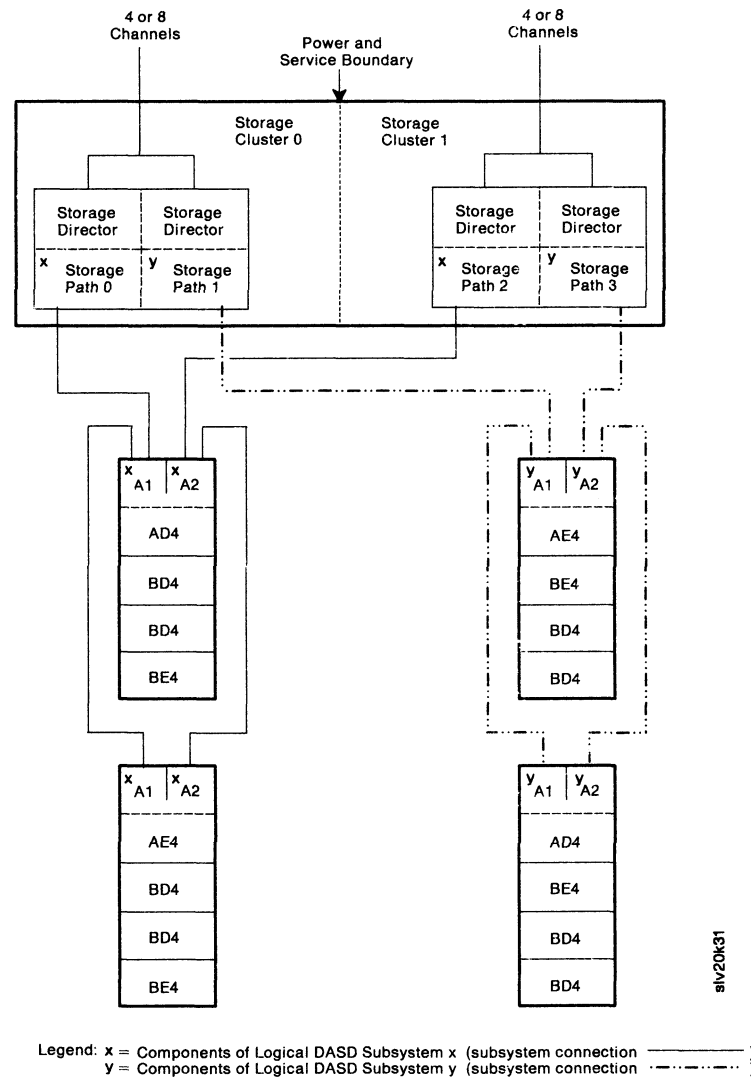


Figure 17. 3990 Model 2 in DLS Mode Attached to 2-Path DASD Strings

Each storage cluster in a 3990 Model 2 or 3 can attach to four channels for a total of eight channels in the logical DASD subsystem. If the Four Channel Switch, Additional feature is installed, each cluster can attach to eight channels, for a total of 16 channels in the logical DASD subsystem.

In DLSE mode, 2-path DASD strings can be intermixed with a 4-path DASD string. In this case, the 4-path DASD string can concurrently transfer data from any of four devices within the string. The 2-path DASD strings can concurrently transfer data from any of two devices within the string, except for 3380 AA4 strings.

Each storage cluster contains a shared control array (SCA). The SCA contains the control and status information about the storage paths. The SCA in one cluster is identical to the shared control array in the other associated cluster. Duplicating information in both clusters ensures that data is available through both of the clusters.

Cache is available in the 3990 Model 3 in four sizes: 32, 64, 128, or 256 megabytes. The size of the cache determines the specific model number.

Table 6. Model Designation by Cache Size	
3990 Model	Megabytes of Cache
G03	32
J03	64
L03	128
Q03	256

See "Chapter 6. Optional Features" on page 85 for descriptions of the optional features available on a 3990 Model 2 or 3.

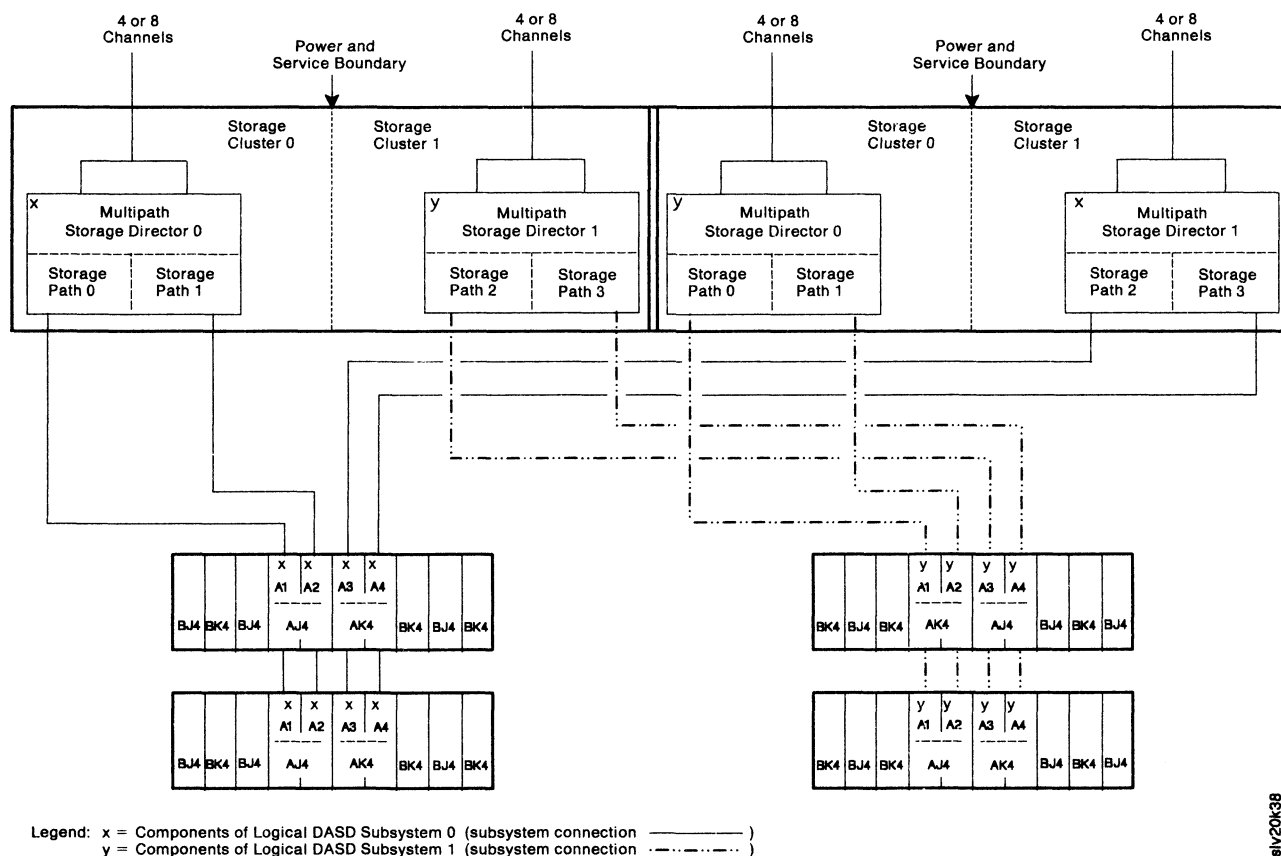
3990 Dual-Frame Configuration

A dual-frame configuration physically connects two 3990 Model 2s or two 3990 Model 3s with dual-frame capable power subsystems. The two frames of the 3990 Model 2s or 3s with dual-frame capable power subsystems are physically bolted together and the storage clusters, caches, and nonvolatile storages are interconnected.

Installing a dual-frame configuration requires that both models be 3990 Model 2s, or both be 3990 Model 3s with dual-frame capable power subsystems. Both 3990s in a dual-frame configuration must have the same channel switch configuration. Each 3990 Model 3 can have a different cache size.

Note: 3990s with modular power subsystems cannot be dual-framed. See "Dual Line Cord" on page 85 for a description of 3990 Models 2 and 3 with this feature.

In a dual-frame configuration, a DASD string is attached to one of the storage directors in each frame. In DLSE (4-path) mode, multipath storage director 0 in one storage control and multipath storage director 1 in the other storage control are connected to the same DASD string as shown in Figure 18.



slv20x38

Figure 18. Dual-Frame 3990 Model 2 in DLSE Mode Attached to 4-Path DASD Strings

Notes:

1. For more information on the dual-frame configuration, see "Hardware and Data Availability Considerations" on page 44.
2. For more examples of the different configurations of the 3990, see the *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*.
3. Feature 6120 is supported on 3990 Model 2 and 3 dual-frame configurations; however, both machines must have the feature installed.

Chapter 4. Programming Support

MVS Support for the 3990

The 3990 runs under MVS/ESA, MVS/XA, and MVS/370, with and without caching. MVS/ESA and MVS/XA support the 3990 extended functions DASD fast write and dual copy; and provide IDCAMS support for subsystem cache fast write with the appropriate level of Data Facility Product (DFP).

Warning: All systems that share a subsystem on which DASD fast write or dual copy is in use must run software that supports these 3990 extended functions.

Minimum MVS Program Levels for 3990 Basic Operations

Note: Different levels of software support are required when the 3990 is attached to 3390 DASD. For the most current listing of these 3390 software support levels, see *Using IBM 3390 Direct Access Storage in an MVS Environment*.

Although the information below lists only the minimum version and release levels required to support the 3990, you should always use the most current level available to take advantage of any product enhancements.

The minimum release level of MVS programs, **with the appropriate PTFs**, required to support 3990 basic operations is: (Basic caching includes caching as in a 3880 Models 13 and 23, and cache fast write operations of the 3990 Model 3.)

- MVS/ESA
 - MVS/System Product Version 3 Release 1.0
 - MVS/DFP Version 3 Release 1.0 *
 - Data Facility Data Set Services (DFDSS) Version 2 Release 4
 - Device Support Facility (ICKDSF) Release 9.0 with PTF UY90143
 - DFSORT Release 9.0 (with appropriate PTFs)
 - DFSORT Release 10 for cache fast write and Read Track support (with appropriate PTFs)
 - EREP Version 3 Release 3.2.
- * Although not MVS/ESA, MVS/DFP Version 2 Release 3.0 is also supported.
- MVS/XA
 - MVS/System Product Version 2 Release 1.2
 - MVS/XA DFP Version 2 Release 2.3 or Version 1 Release 1.3
 - Data Facility Data Set Services (DFDSS) Version 2 Release 2.0
 - Device Support Facility (ICKDSF) Release 9.0
 - DFSORT Release 9.0 (with appropriate PTFs)
 - DFSORT Release 10 for cache fast write and Read Track support (with appropriate PTFs)
 - EREP Version 3 Release 3.2.

- MVS/370
 - MVS/System Product Version 1 Release 3.5
 - MVS/370 Data Facility Product Version 1 Release 1.2
 - Data Facility Data Set Services Version 2 Release 2.0
 - Device Support Facility (ICKDSF) Release 9.0
 - DFSORT Release 9.0 (with appropriate PTFs)
 - DFSORT Release 10 for cache fast write and Read Track support (with appropriate PTFs)
 - EREP Version 3 Release 3.2.

This support includes:

- Service information message (SIM) processing

When a SIM is issued by the storage control, a SIM Alert message is issued to the host operator's console and the information is written to SYS1.LOGREC.
- IDCAMS support for remote support

The IDCAMS LISTDATA ACCESSCODE command provides an access code to permit a remote support session to be established. This support is available for all models of the 3990 and for the 3380 Model CJ2. The LISTDATA ACCESSCODE command is protected from unauthorized use by System Authorization Facility (SAF) checking.
- Reset Allegiance

Software can request termination of a device's allegiance to a channel path or path group which will only be successful if the:

 1. Device is busy or not ready, and
 2. Device is not reserved to another channel or path group.
- Reset Event

Software recognizes the resetting of DASD devices attached to a 3990 and responds by rebuilding the dynamic pathing arrays, rather than rebuilding them as part of error recovery.

MVS/ESA and MVS/XA Software Capabilities for 3990 Model 3 Extended Functions

Note: Different levels of software support are required when the 3990 is attached to 3390 DASD. For the most current listing of these 3390 software support levels, see *Using IBM 3390 Direct Access Storage in an MVS Environment*.

Extended function operations are provided for MVS/ESA and MVS/XA. MVS/DFP Version 3 Release 1.0 provides this support when running with MVS/SP Version 3 (for MVS/ESA) and MVS/SP Version 2 Release 2.3 (for MVS/XA). 3990 Extended function support is also provided for MVS/XA environments with Version 2 Release 4.0.

Extended function support includes:

- IDCAMS support of NVS, DASD fast write, cache fast write, and dual copy

Parameters are added to the IDCAMS SETCACHE command to enable/disable DASD fast write, cache fast write, and NVS, and to manage the dual copy capabilities.

- TSO execution of IDCAMS LISTDATA, SETCACHE, and BINDATA commands

LISTDATA and SETCACHE commands can be invoked as TSO commands in addition to being submitted as batch jobs. The normal TSO help functions are available for these commands.

- DEVSERV

The operator command Device Services (DEVSERV) has been modified to add information on caching status of volumes. This includes information on DASD fast write and dual copy for volumes attached to a 3990 Model 3 Storage Control.

- ISMF support

ISMF Volume Application supports caching subsystems by recognizing a subsystem identifier and using that identifier to select volumes for the volume list. In addition, caching status, cache fast write status, DASD fast write status, dual copy status, and subsystem identifier are provided for each volume in the volume list. With MVS/DFP Version 3 Release 1.0, ISMF can also be used to invoke SETCACHE functions and to request remote access authorization codes.

- SAF authorization

IDCAMS commands will use the System Authorization Facility to restrict invocation of the commands and certain parameters to only authorized users if RACF or an equivalent user routine is available. The IDCAMS BINDDATA and SETCACHE commands, as well as LISTDATA, require authorization.

In addition, users must have DASDVOL ALTER access authority for volumes that have different volsters than the target primary volume specified by the SETSECONDARY and REESTABLISHDUPLEX parameters of the SETCACHE command.

- SMF Type 22 records

System Management Facility includes records to indicate changes in the status of the 3990 Model 3 and attached devices. At each change of subsystem or device status, such as setting caching on or off or enabling DASD fast write, MVS creates an SMF Type 22 record. These records can be used to build an audit trail of all device/subsystem status changes.

- Asynchronous Operations Manager

The Asynchronous Operations Manager is a component of DFP that handles storage control I/O operations. In particular it issues the asynchronous commands and processes attention interrupts from the subsystem.

- DFSMS

DFSMS (part of MVS/DFP) provides for management of the 3990 cache resource at a data set level (available with MVS/DFP Version 3, Release 1.0 and MVS/SP Version 3, Release 1.0 only).

DFP Releases that Support the 3990 Models 1, 2, and 3

The following matrix shows the DFP releases that support the various 3990 models.

Table 7. DFP Releases that Support the 3990 Models 1, 2, and 3					
3990 Model	MVS/DFP 3.1.0	DFP/XA 2.3.0	DFP/XA 2.4.0	DFP/XA 1.1.3	DFP/370 1.1.2
3990 Model 1 and 2	Yes	Yes	Yes	Yes	Yes
3990 Model 3 basic caching	Yes	Yes	Yes	Yes	Yes
3990 Model 3 extended functions	Yes	No	Yes	No	No

VM Support for the 3990

The following text describes the VM support for the 3990. Although the information below lists only the minimum version and release levels required to support the 3990, you should always use the most current level available to take advantage of any product enhancements.

Note: Different levels of software support are required when the 3990 is attached to 3390 DASD. For the most current listing of these 3390 software support levels, see *Using IBM 3390 Direct Access Storage in a VM Environment*.

Minimum VM Program Levels for 3990 Model 1 and 2 Support

For VM support of the 3990 Model 1 and 3990 Model 2, the minimum release level of programs, **with the appropriate PTFs**, required to support 3990 Model 1 and 2 basic operations includes:

- VM/XA SP Release 1.0
- VM/XA SF Release 2.0
- VM/SP HPO Release 4.2
- VM/SP Release 4.0
- EREP Version 3 Release 3.2
- Device Support Facility (ICKDSF) Release 9.0.

VM Software Capabilities for 3990 Model 3 Basic Caching

Basic caching support for the 3990 Model 3 is provided by VM/SP HPO Release 5 with the appropriate PTFs and VM/XA SP Release 2 with the appropriate PTFs. VM/SP HPO also provides support for DASD fast write.

For information on the specific changes for VM/SP HPO Release 5 and 3990 operation, including how to use cache fast write for guest VM/SP HPO operating systems, see *VM/SP High Performance Option Release 5 CP Support for the IBM 3990 Storage Control Model 3*.

VSE Support for the 3990

The following text describes VSE programming support for the 3990.

Although the information below lists only the minimum version and release levels required to support the 3990, you should always use the most current level available to take advantage of any product enhancements.

Minimum VSE Program Levels for 3990 Models 1 and 2 Support

VSE support of the 3990 Model 1 and 3990 Model 2 is provided in VSE/SP Version 2, Release 1.7 and Version 3, Release 1.2. Included in the releases are:

- EREP Version 3 Release 3.2
- Device Support Facility (ICKDSF) Release 9.0.

See *Using the IBM 3380 Direct Access Storage in a VSE Environment* for program levels required for specific 3380 models.

Minimum VSE Program Levels for 3990 Model 3 Support

Programming support for the 3990 Model 3 basic caching operations is provided in VSE/SP 4.1.

TPF Support for the 3990

Transaction Processing Facility (TPF) support for the 3990 Models 2 and 3 is provided in an update to TPF 2.3 and TPF 3.1.

Chapter 5. Using Dual Copy, DASD Fast Write, and Cache Fast Write in MVS Environments

This chapter describes uses for dual copy, DASD fast write, and cache fast write in MVS environments.

Dual copy and DASD fast write are supported in MVS/ESA and MVS/XA. Cache fast write is supported for sort work files in DFSORT Release 10 or higher.

MVS/370 supports cache fast write for sort work files in DFSORT Release 10 or higher, but does not provide software support for dual copy or DASD fast write. IDCAMS commands to turn cache fast write on and off are not available in MVS/370. IDCAMS commands to turn cache fast write on and off are available in MVS/DFP Version 3 Release 1.0 and MVS/XA DFP Version 2 Release 4.0.

Using Dual Copy in MVS/ESA and MVS/XA Environments

The 3990 Model 3 dual copy function significantly improves data availability.

A properly configured single 3990 Model 3 with the dual line cord feature or a properly configured pair of 3990's in a dual-frame configuration eliminates all single points of failure which might prevent access to data on dual copy pairs—from processor channels to storage control to dual copy volumes. All path components are replicated. Failing operations are retried through another path, or to an alternate device, if necessary. Using dual copy may result in fewer application outages, fewer IPLs, and a reduced frequency of DASD recovery activity.

Using dual copy on important system and application volumes can provide significant relief in scheduling batch applications and routine backups. With the shrinking batch window and growing numbers of installations with 24-hour operations, many installations cannot backup as often as they would prefer. Some installations are copying key volumes while their applications are running (and updating the volumes). Using dual copy on key system and application volumes ensures that the subsystem maintains duplicate, up-to-date copies of these volumes. As a side benefit, installations may be able to reduce the frequency of full volume backups or Image Copies of data bases.

Note: Backups and log tapes cannot be completely eliminated because of the need to recover from application logic errors as well as multiple hardware failures.

Implementation of dual copy and DASD fast write is straightforward. No application changes are necessary as long as programs meet the following requirements:

- Do not perform Write Home Address Record Zero operations.
- Use only standard Record 0 format. MVS/ESA and MVS/XA access methods use the standard Record 0 format.
- Do not define duplicate Record IDs on a track.
- Alternate tracks must be assigned in the range provided by IBM. See the *IBM 3990 Storage Control Reference*.

Normally, using the dual copy function requires no special attention from the installation beyond having procedures in place to establish the dual copy pairs and perform recovery actions when required.

Using DFSMS to Manage Placement of Data Sets on Duplex Pairs

With DFSMS, use of storage classes determines those data sets that will be allocated to volumes with dual copy operations.

By creating a storage class of AVAILABILITY = CONTINUOUS, the Data Facility Storage Management Subsystem (DFSMS) will automatically allocate data sets specified with this level of availability to dual copy volumes specified in the DFSMS configuration. This allows the storage administrator the ability to manage (through the automatic class selection (ACS) routines) the placement of important high availability data sets automatically, by the proper specification of a DFSMS policy for the system.

3990 Model 3 Fast Dual Copy Performance Considerations

Dual copy is an availability-oriented function available in the 3990 Storage Control. Dual copy is intended to provide enhanced availability to an installation's important data. It can serve this purpose in two ways:

- By continuously providing a second copy of a volume, automatically kept synchronized by the storage control, and providing an automatic switch-over to the second copy in the event of a failure.
- By providing a hardware facility for nondisruptive migration of data from one volume to another like volume in the same subsystem.

There are a number of factors which contribute to overall subsystem response time of a 3990 Model 3 using dual copy with DASD fast write. These include the number of volumes that exist as dual copy pairs, the I/O activity to the subsystem as well as to the duplexed volumes themselves, the read/write ratios for the volumes in the subsystem, the percentage of total write activity directed at the dual copy pairs, block size and chaining, and presence of unusual channel programs. While we will provide you with general guidelines for performance with fast dual copy, it is important to note that performance in your installation is a function of all of the variables we will describe. A particular data set could exhibit different performance characteristics, depending on the particular configuration and its activity. We will discuss these items in more detail below.

Fast Dual Copy Operation

The presence of fast dual copy may have an impact on subsystem performance. This is due to the increased level of write activity required to maintain the second copy, as well as special care that must be taken with synchronization of key operations within the subsystem to maintain the two copies. When a fast dual copy write occurs, the device end sent to the channel (signalling completion of the I/O operation to the host system) follows immediately after the branching write of the updated record into cache and NVS. The device remains busy while the data is destaged to the primary device of the duplex pair (the secondary device is asynchronously updated at a later time). No further host system activity is allowed to the data on this volume until this update to the primary device is complete.

Fast Dual Copy Performance Measurements

Laboratory measurements were taken of 3380 Models J and K, and 3990 Model 3 fast dual copy. Benchmark scripts of TSO and IMS systems were used to obtain these measurements. The scripts were run at several different transaction rates. Measurement data were gathered with the 3990 Model 3 subsystem in basic caching, DASD fast write, and fast dual copy modes. For each type of data (TSO and IMS), there were two sets of fast dual copy runs. The only difference between these runs was the number of volumes using fast dual copy.

With these scripts we set the transaction rate and the transactions generated work (processor instructions and I/O's). Since this is an I/O measurement, the I/O rate and DASD subsystem response time are reported. The I/O rate generated in one set of runs might be slightly different in another set of runs. For example, at the lowest TSO transaction rate the basic caching run generated 119 I/O's per second, while the DASD fast write generated 123 I/O's per second. We considered these equal and called it approximately 120 I/O's per second.

It is important to remember that these measurements were taken in a laboratory environment. We had control of the variables and the luxury of being able to run the same workload over and over, while changing one variable at a time. In most real environments the workload is constantly changing. Your experience with the performance of fast dual copy will, therefore, vary from what we have seen. Consequently, it is very important that you review the guidelines for your data, develop a plan for implementing fast dual copy on a small number of volumes, measure the performance results, and proceed to add more fast dual copy volumes, as required, as long as you are meeting your performance and availability goals.

By changing the number of fast dual copy pairs, and thus the percentage of total subsystem writes directed at the fast dual copy pairs, these measurements give you an appreciation for the impact of varying one of the parameters affecting fast dual copy performance. In these measurements, the response time varied with the number of duplex volumes and the percentage of total subsystem write activity addressed to the duplexed pairs; as these factors increased, so did subsystem response time. A closer analysis shows that most of the increased response time was experienced by the duplex pairs themselves.

The performance values should be taken as indicators of relative performance characteristics. It is important to realize that some of the workload characteristics may not apply to other operating environments.

When evaluating DASD subsystem performance characteristics, it is also important to realize that improvements in DASD subsystem performance may not necessarily

result in system performance improvements. Other variables such as the processor cycle speeds, the main storage size, the operating system, and the communications network may limit the potential subsystem performance even if the DASD subsystem performance has been improved.

Other environments, configurations, and processors will experience different levels of performance. Accordingly, these figures do not constitute a performance guarantee or warranty.

TSO Fast Dual Copy

- System:
 - 3090-200E (128MB central with 256MB expanded storage)
- Software Environment:
 - TSO/E 1.4
 - JES 2.2.0
 - DFP 3.1.0
 - MVS/SP 3.1.0e.
- DASD subsystem:
 - 3990 Model G03 with 32MB of cache and 32 cached primary volumes (16 3380-J and 16 3380-K – plus devices for secondaries)
- DASD volumes groups:
 - TSO user, work, and catalog
 - ISPF library
 - System library
- Read to write ratio = 5:1
- Read hit ratio = 0.86
- DASD fast write hit ratio = 0.94

A series of runs were made with the 3990 in basic caching mode. The subsystem I/O rates ranged from approximately 120 to 380 requests per second. The runs were then repeated with DASD fast write active for all volumes in the subsystem. DASD fast write was able to improve subsystem response time by approximately 27% at the lowest I/O rate to approximately 23% at the highest I/O rate.

Subsystem Performance with Six Fast Dual Copy Pairs: Six volumes of data were selected for duplexing. These six volumes contained data sets that if access was lost to them it would cause an IPL or an interruption of TSO service. The read to write ratio for this data was 7.7:1 which represented 16% of all of the write activity in the measured subsystem. At the 120 to 200 plus I/O rate, the subsystem response was approximately 10% lower than basic caching. At approximately 380 I/O's per second, the subsystem response time was 5% lower than basic caching.

Subsystem Performance with Thirteen Fast Dual Copy Pairs: Seven additional volumes of data were then selected for duplexing. These seven volumes contained more write intensive data such that the read to write ratio for the thirteen volumes decreased to 4.3:1. These volumes represented 48% of all of the write activity in the measured subsystem. At the 120 to 200 plus I/O rate, the subsystem response time was approximately 10 to 11% higher than basic caching. At approximately 380 I/O's per second range, the subsystem response time was approximately 40% higher than basic caching. However, it is important to note that the highest subsystem response time measured was less than 16 milliseconds. This is lower than we could expect from noncached DASD.

IMS Fast Dual Copy

- System:
 - 3090-200S (128MB central with 256MB expanded storage)
- Software Environment:
 - IMS/VS 2.2 (full function)
 - MVS/SP 3.1.0e
 - JES 2.1.5
 - DFP 3.1.1.
- DASD subsystem:
 - 3990 Model Q03 with 256MB of cache and 32 cached primary volumes (16 3380-J and 16 3380-K – plus devices for secondaries)
- DASD volumes groups:
 - IMS database
- Read to write ratio = 7.2:1
- Read hit ratio = 0.72
- DASD fast write hit ratio = 0.99+

A series of runs were made with the 3990 in basic caching mode. The subsystem I/O rates ranged from approximately 130 to 380 requests per second. The runs were then repeated with DASD fast write active for all volumes in the subsystem. DASD fast write was able to improve subsystem response time by approximately 18% at the lowest I/O rate to approximately 12% at the highest I/O rate.

Subsystem Performance with Five Fast Dual Copy Pairs: Five volumes of data were selected for duplexing. These five volumes contained data that were updated quite often and considered critical to the IMS system. The read to write ratio for this data was 2.2:1 and represented 25% of all of the write activity in the measured subsystem. At the 130 to just under 200 I/O rate, the subsystem response time was approximately 10% lower than basic caching. From approximately 200 to 300 I/O's per second, the subsystem response time was approximately 7 to 8% lower than basic caching.

Subsystem Performance with ten Fast Dual Copy Pairs: Five additional volumes of data were selected for duplexing. These five additional volumes contained more write intensive data such that the read to write ratio for the ten volumes decreased to 1.85:1. These volumes represented 76% of all of the write activity in the measured subsystem. As the I/O rate increased from approximately 130 to the low 300 range, the subsystem response time went from approximately 4 to 20% higher than basic caching. At the 380 I/O's per second range, the subsystem response time was approximately 36% higher than basic caching. The highest subsystem response time measured was less than 24 milliseconds. This also is lower than we could expect to see from noncached DASD at these I/O rates.

Summary

In summary, the performance results cited showed that the overall DASD subsystem response time varies with the number and type of writes being duplexed. Fast dual copy subsystem performance can be much better than or not quite as good as basic caching depending on many variables.

Keep in mind that fast dual copy is primarily an availability feature for the 3990 Model 3 subsystem. Therefore, when preparing to implement fast dual copy, we suggest that you build a prioritized list of the data sets that may need dual copy protection. You should review system data sets, such as control data sets, as well as application data sets. The candidate list in "Potential Fast Dual Copy Candidates" on page 70 may give you some indications of the types of data sets you want to consider. Next, you need to gather RMF device activity reports for the volumes and subsystems containing these data sets. Use this information to evaluate your candidates for fast dual copy, using the guidelines described below, and modify your prioritized list of candidates appropriately. Proceed to establish fast dual copy pairs, starting with one or two volumes per subsystem, then gradually adding more.

Rather than just looking at the device or subsystem level performance of fast dual copy, it is very important to look at the system level impact—in terms of transaction throughput and transaction response time. What may appear to be a significant change in a particular device's response time, may or may not have a visible impact on total system performance. Even if there is measurable impact, the value of the second copy must be considered in making the final evaluation of the appropriateness of fast dual copy.

Fast Dual Copy Selection Recommendations

For best 3990 Model 3 subsystem performance when using dual copy, we recommend that you:

- Use fast dual copy, that is, set DASD fast write on to all of your dual copy pairs.
- Wherever possible, use DLSE DASD devices for your pairs.
- Limit the number of writes duplexed—start out by duplexing no more than 20 to 25% of the write activity in any given subsystem.
- Initially, plan on no more than three or four fast dual copy pairs within one logical DASD subsystem.
- When additional fast dual copy pairs are desired, fast dual copy volumes can be added one at a time. To ensure that performance levels are satisfactory, monitor subsystem performance and application response times after adding each volume.
- As much as is practicable, spread your fast dual copy pairs across all of your 3990 Model 3's. By doing so, you minimize the impact of having a larger number of fast dual copy pairs on one subsystem. Remember that both devices in a duplex pair must be in the same logical DASD subsystem.

The following items need to be considered as you plan your fast dual copy implementation, as they could impact fast dual copy performance. You need to evaluate your fast dual copy candidates for these characteristics. When you use fast dual copy for such data sets or volumes, evaluate your performance after setting fast dual copy on and before adding additional volumes.

- Complex channel programs, particularly those that cause writes across several tracks. One such example is JES2 CHECKPOINT.
- Data sets with high levels of sharing across tasks or systems. Such data sets tend to exhibit more queueing at the system level, and may be more sensitive to any increase in pend time resulting from the use of dual copy.
- High I/O rate sequential writes.
- Data sets with high I/O rates that get repeated write hits to individual tracks, such as JES SPOOL.
- If you need a redundant copy of the IMS OLDS or WADS, use the IMS facility.
- An already heavily loaded DASD subsystem (typically over 400 SIO's per second for IMS workloads). Consider placement of fast dual copy pairs across several subsystems.

Potential Fast Dual Copy Candidates

The following data sets, data bases, or catalogs might be evaluated as candidates for fast dual copy because of their importance to the system. However, many of the factors described above vary from one installation to another. Depending on how well these data sets conform to the guidelines described here, in your environment they may be good candidates for fast dual copy.

The following examples describe current recovery scenarios for several important classes of volumes that are candidates for dual copy. Note that a number of these data sets, such as RACF, have software duplicate copy capability. In such cases, you should evaluate the characteristics, both operational and performance, for both techniques and select the one most appropriate for your environment. The volume classes include:

- MVS system data sets
- RACF control data sets
- DFHSM control data sets
- Integrated catalog facility catalogs
- Critical system volumes (such as load libraries and IMS libraries)
- IMS data base volumes
- CICS application volumes
- Various user volumes.

MVS System Data Sets

Examples of data sets that are critical to the availability of an MVS system include:

- SYS1.LINKLIB
- SYS1.LPALIB
- SYS1.PARMLIB
- SYS1.PROCLIB
- SYS1.NUCLEUS
- Integrated catalog facility master catalog.

Using dual copy may prevent a system IPL caused by a hardware failure on a volume containing one of these data sets.

Without dual copy, an I/O error on the SYSRES volume may cause a system IPL. Installation response is to IPL again with the alternate SYSRES. Depending on recovery/restart requirements, the outage could be from 10 minutes to an hour or more. By using dual copy on SYSRES, IPLs due to DASD hardware failure can be reduced. The result is fewer end user outages and less recovery/restart activity.

Integrated catalog facility master catalog considerations are described later.

RACF Control Data Set

The Resource Access Control Facility (RACF) control data set is another critical system resource. An I/O error on this data set may lead to operations intervention.

RACF provides a facility for maintaining a software duplicate of information in the RACF control data set. If an error occurs on the primary copy, the installation may switch to the duplicate copy.

The installation may wish to evaluate the use of the 3990 Model 3 dual copy function for the RACF control data set.

DFHSM Control Data Sets

DFHSM is a started task whose automatic and command operations require access on demand to its control data set. DFHSM provides facilities to maintain backup versions of its control data sets and journal for up-to-date recovery. However, recovery procedures are manual and take time.

Installation management should review the impact of access time and physical I/O errors to their operations. Consideration for evaluating DASD fast write and dual copy include:

- Use of DFHSM inline backup for critical data sets
- Use of DFHSM to manage backup version for critical applications
- Backup window interference with scheduling critical operations.

While DFHSM provides several capabilities for logging updates and backing up data sets, you may wish to use the 3990 Model 3's dual copy function for DFHSM control data sets, the migration control data set, the backup control data set, and the offline control data set.

Integrated Catalog Facility Catalogs

The effects of a hardware failure on an integrated catalog facility catalog volume are different for master catalogs and user catalogs. If the master catalog fails, a system IPL is required. If a user catalog fails, application data is unavailable for allocation.

Integrated Catalog Facility Master Catalog: Without the 3990 Model 3 dual copy function, hardware failure of the master catalog volume may lead to a system IPL using the MVS alternate master catalog facility. During the nucleus initialization program (NIP) processing, the operator is prompted for an alternate nucleus member describing the alternate master catalog. Common practice today is to limit update activity to the master catalog by having it contain just the critical system data sets and pointers to user catalogs. Thus, if the master catalog has no changes, the alternate can be invoked with no difficulty. Using dual copy may reduce the IPLs caused by DASD hardware failures.

User Catalogs: User catalogs typically have a fairly high level of change activity. For this reason, many installations backup their more critical user catalogs frequently. Without dual copy, recovery typically involves:

1. Creating a new catalog entry
2. Importing the most recent backup copy of the catalog
3. Updating the catalog entries made since the backup copy.

Recovery may involve using volume table of contents (VTOCs), VSAM Volume Data Sets (VVDs), and System Management Facility (SMF) records, and can take anywhere from hours to days. Use of a tool such as the program offering Integrated Catalog Forward Recovery Utility (program number 6798-DXQ) may substantially reduce the manual effort required.

Dual copy may be desirable where there is a need to reduce the number of catalog backups, or where even the reduced recovery time provided by the integrated catalog facility recovery utility is not sufficient to satisfy user availability requirements.

Other Critical System Data Sets

There are a number of other critical system data sets, such as:

- Program libraries
- ISPF data sets
- LINKLIST data sets
- SMP data sets
- Other system data sets with low levels of update activity
- Other system data sets for which there are facilities in place to log changes or otherwise provide backup copies.

All of these may still be candidates for dual copy because dual copy insulates the operating system or applications from the effects of a DASD hardware failure.

Full Function IMS Data Bases

Without dual copy, an I/O error on an IMS data base volume can bring the application down and the installation must invoke data base recovery. Depending on the frequency of image copy and update activity, the recovery process could take several hours. The basic steps in IMS recovery include:

1. Locate the latest image copy tapes
2. Locate a spare volume
3. Restore all volumes of the data base
4. Find all log data sets
5. Generate JCL for the recovery
6. Run change accumulation (if appropriate)
7. Apply the log data sets or change accumulation data sets
8. Restart the application.

Data Base Recovery Control (DBRC) may be used to facilitate recovery. It can be used to identify the proper image copy tapes and log data sets, perform change accumulation, and generate the JCL for the recovery process.

Dual copy can reduce application outages caused by hardware failure and the resulting recovery activity.

CICS/DL1 Data Bases

Considerations are very similar to full function IMS data bases.

CICS VSAM Data Sets: Without dual copy, considerations here are quite similar to IMS, with some differences in the recovery process:

1. Locate the backup tapes
2. Locate a spare volume
3. Import the backup tape
4. Find the log tapes
5. Generate JCL for the application's recovery routine
6. Run this recovery routine to apply the log tapes
7. Restart the application.

If there is no user-written recovery routine, invoke the recovery procedures the installation has defined.

Dual copy can significantly reduce application outages caused by hardware failure and the resulting recovery activity.

Other Data Base Applications

Evaluate each data base for applicability of dual copy. An analysis of data bases with high-availability requirements will identify data set or volume candidates and the value of using dual copy for those candidates.

User Volumes

User volumes are typically backed up on a regular schedule determined by operations and the end user groups. A hardware failure may result in the restore of the most recent backup tape. Recovery of updates or new data sets created after the most recent backup tape was created is typically an end user responsibility. The impact of such a failure varies greatly, depending on the user's capability to recreate the lost data.

Using an archival/backup system such as DFHSM can substantially reduce the recovery effort. Typically, the latest backup tapes are applied, and then the restore function is used to recover data sets created or updated since the last backup, up to the last DFHSM backup cycle. Again, further recovery action is typically up to the user.

Dual copy may be used routinely on frequently-updated user volumes, on volumes containing data that is very difficult to recreate, or on application volumes with high-availability requirements.

DASD Fast Write in MVS/ESA and MVS/XA Environments

DASD fast write improves the performance of data sets that have a significant amount of write activity. Good DASD fast write candidates include catalogs, JES CHECKPOINT, SPOOL, other control data sets, and many logging functions with moderate levels of activity. Data bases or VSAM data sets may also be good candidates.

Typically, the write hit ratios observed have been very high for two reasons:

- Most applications read a record before updating it, so the number of truly random write misses is quite small. (Full-track format writes are considered write hits.)
- When writing a sequential output file, a format write is used for the entire track, which is considered a write hit by the 3990 Model 3 (no stage is required since the entire track is rewritten).

Some applications may not show a performance improvement with DASD fast write. As in the case of normal caching, update writes where the transfer size is more than half a track may not benefit from DASD fast write. Data for which the write hit ratio is less than 85% might not benefit from DASD fast write. Also, applications such as very-high-activity logging functions that write half a track or more at a time with very high request rates might not perform better than noncached DASD. (This is true even if there is a high write hit ratio.) Such applications must be individually evaluated.

MVS DASD Fast Write Data Selection Guidelines

Table 8 shows data set candidates for DASD fast write operations.

Table 8. MVS DASD Fast Write Data Set Selection Guidelines				
Data Set Categories	System and General Data Sets	TSO	IMS	CICS
Good DASD Fast Write candidates	DFHSM control data sets, other control data sets, SPOOL volumes, JES Checkpoint, and work volumes (in moderate number per 3990 Model 3)	Libraries, ISPF, TSO user data, Catalogs	Moderate-sized data bases, IMS message queues, scratch pad area, image copy, RECON, WADS	System Log, Intra-partition transient data, aux temporary storage, aux trace. Larger application data bases with high update activity (recommended with the larger cache sizes)
Potential good DASD Fast Write candidates (See the following note.)	Other system-type data sets, lookup tables or dictionaries, large sequential data sets, custom systems, and data sets with high update activity		Larger full-function data bases with high update activity (recommended with the larger cache sizes)	Larger application data bases with high update activity (recommended with the larger cache sizes)

Note: Custom-designed systems or other software might have good cache candidates that you can evaluate by the criteria given above. For example, data sets having read hit and DASD fast write hit ratios of 85% or better and a read-to-write ratio of 2:1 are good candidates for a 3990 Model 3.

The performance of DASD fast write in typical TSO and IMS data base applications has been measured using data collected from representative large TSO and IMS data base systems. In such environments, DASD fast write performance improved subsystem response time and increased the number of accesses per second.

In many installations, it is desirable to use DASD fast write for applications with very high levels of write activity. Such applications require analysis, and probably testing on the 3990 Model 3. It is difficult to develop guidelines for these, due to the varying load on the storage control and the I/O rate to the write data set. The suggested procedure is to place this high write activity application on the 3990 Model 3 first, and add more workload gradually. In general, the higher the hit ratio of the additional data, the more easily the additional load may be absorbed with a good response time.

DASD Fast Write Performance Measurements

This section shows performance measurements of 3880 Model AJ4s and AK4s and 3990 Model 3 DASD fast write extended function. The different operating measurements were obtained through measurements of normal production environments which represent a variety of workloads, configurations, and transaction mixes. The performance values should be taken as indications of relative performances characteristics. It is important to realize that some of the workload characteristics may not apply to other operating environments.

When evaluating DASD subsystem performance characteristics, it is also important to realize that improvements in DASD subsystem performance may not necessarily result in improvements in system performance. Other variables such as the processor cycle speeds, the main storage size, the operating system, and the communications network may limit the potential system performance even if the DASD subsystem performance has been improved.

Other environments, configurations, and processors will experience different levels of performance. Accordingly, these figures do not constitute a performance guarantee or warranty.

The performance information presented here was gathered at five different data processing installations utilizing the IBM 3990 Model 3 Extended Function, DASD fast write.

- One 3990 Model 3 with 53 cached DASD volumes (TSO — computer industry)
- One 3990 Model 3 with 32 cached DASD volumes (TSO — computer industry)
- One 3990 Model 3 with 52 cached DASD volumes (TSO — communications industry)
- One 3990 Model 3 with 37 cached DASD volumes (Online/TSO — banking industry)
- One 3990 Model 3 with 36 cached DASD volumes (IMS — aerospace industry).

Summary of DASD fast write performance results:

At all installations, when DASD fast write (DFW) was enabled to the volumes, average 3990 Model 3 response times dropped significantly. DASD models used in the measurements were connected to their 3990 Model 3s in device level selection enhanced (DLSE) mode. The activity rates (SIOs per second) and the read hit ratios (H/R) remained nearly constant during the DFW versus non-DFW comparisons at each installation. Therefore, we can assume that the response time improvement can be attributed entirely to the DASD fast write capability of the 3990 Model 3. The performance improvement in the DASD subsystem was provided

entirely by DASD fast write, not by additional cache. (Example 2 also shows an additional improvement resulting by adding additional cache.)

Certain volume groups generally exhibited response time improvements more significant than the 3990 Model 3 averages. These are the TSO user volumes, JES checkpoint/SPOOL and work volumes. In the IMS environment, it was the WADS volumes that displayed significant subsystem response time reductions. Volumes with a consistently low read-to-write ratio (less than two) received the greatest performance gain from DASD fast write.

Example 1 – DASD Fast Write (TSO)

The following shows the environment and measurements obtained in a computer industry installation.

- Software Environment:
 - TSO/E 2.1
 - JES2 2.2.0
 - DFP 3.1.1
 - MVS/ESA 3.1.1
 - MVS/XA 2.2.3 (on 3090-400).
- Systems:
 - 3090-600S (256MB main with 1024MB expanded storage)
 - 3090-600E (256MB main with 256MB expanded storage)
 - 3090 400E (256MB main storage)
 - 3090-400 (128MB main with 512MB expanded storage).
- DASD subsystem:
 - 3990 Model 3 with 53 cached volumes (41 3380-K and 12 3380-J)
 - 42 DFHSM managed TSO storage volumes
 - 1 DFHSM control volume
 - 3 JES SPOOL volumes
 - 1 JES checkpoint volume
 - 3 work volumes
 - 2 system volumes
 - 1 IGS (graphics) volume.
- The following chart shows the subsystem performance results:

Function	Cache Size	SIO/Sec	Response Time	Read Hit Ratio	Fast Write Hit Ratio	Read/Write Hit Ratio
Base Caching	128MB	318	13.5 ms	0.949	0.000	3.90
DASD Fast Write Change	128MB	322 1%	8.5 ms -37%	0.964	0.940	3.83

- Conclusions and Observations:

In addition to observing outstanding subsystem level performance improvements using DASD fast write (37% better response time), some significant improvements were noted for individual volumes.

For example, using only basic caching support, the JES2 SPOOL volume provided an average response time of 34.6 milliseconds. When DASD fast write was added for the volume, average response time for the SPOOL volume dropped to nine milliseconds, a reduction of 74%.

Example 2 – DASD Fast Write (TSO)

The following shows the environment and measurements obtained in a computer industry installation.

- Software Environment:
 - TSO/E 1.4
 - JES3 2.2.1
 - DFP 3.1.1
 - MVS/XA 2.2.3.
- Systems:
 - 3090-600E (256MB main with 1024MB expanded storage)
 - 3090-400 (128MB main with 512MB expanded storage)
 - 3090-400 (128MB main with 512MB expanded storage)
 - 3084-QX (128MB main storage).
- DASD subsystem:
 - 3990 Model 3 with 32 cached volumes (16 3380-K and 16 3380-J)
 - 16 DFHSM managed TSO storage volumes
 - 1 DFHSM control volume
 - 1 JES SPOOL volume
 - 4 work volumes
 - 5 engineering design volumes
 - 2 system volumes
 - 1 catalog volume
 - 1 VSAM volume
 - 1 spare volume.
- Subsystem performance results:

Function	Cache Size	SIO/Sec	Response Time	Read Hit Ratio	Fast Write Hit Ratio	Read/Write Hit Ratio
Base Caching	64MB	278	9.66 ms	0.958	0.000	5.55
DASD Fast Write Change	64MB	278 0%	6.79 ms –30%	0.950	0.934	6.00
DASD Fast Write Change	128MB	296 7%	5.78 ms –40%	0.961	0.956	6.38

- Conclusions and Observations:

In this installation, the cache size was also increased from the base of 64MB to 128MB. Notice that the read hit ratio increased very little, while the write hit ratio had a larger increase. The result was an additional improvement in average response time.

Also note that the SPOOL volume response the SPOOL volume response time improved 59% with DASD fast write (DFW) enabled.

Example 3 – DASD Fast Write (TSO)

The following shows the environment and measurements obtained in a communications industry installation.

- Software Environment:
 - TSO/E 2.1
 - JES3 2.2.1
 - DFP 3.1.0
 - MVS/ESA 3.1.1.
- Systems:
 - 3090-600E (LPAR) (64MB main with 64MB expanded storage)
 - 3090-400E (LPAR) (220MB main with 96MB expanded storage)
- DASD subsystem:
 - 3990 Model 3 with 52 cached 3380-K volumes
 - 17 DFHSM managed TSO storage volumes
 - 13 DFHSM Level 1 migration volumes
 - 12 application volumes
 - 4 system volumes
 - 6 other volumes
- Subsystem performance results:

Function	Cache Size	SIO/Sec	Response Time	Read Hit Ratio	Fast Write Hit Ratio	Read/Write Hit Ratio
Base Caching	64MB	262	12.3 ms	0.922	0.000	3.32
DASD Fast Write Change	64MB	284 8%	7.42 ms -40%	0.932	0.948	3.69

- Conclusions and Observations:

Subsystem level performance improved 40% using DASD fast write. The improvement was spread fairly evenly across most of the volume groups.

Example 4 – DASD Fast Write (Online/TSO)

The following shows the environment and measurements obtained in a banking industry installation.

- Software Environment:
 - TSO/E 2.1
 - CICS 1.7
 - JES2 2.2.0
 - DFP 3.1.1
 - MVS/ESA 3.1.1.
- Systems:
 - 3090-400E (128MB main with 512MB expanded storage)
 - 3090-400E (LPAR) (48MB main with 256MB expanded storage)
- DASD subsystem:
 - 3990 Model 3 with 37 cached volumes (27 3380-K and 10 3380-J)
 - 7 ATM database volumes
 - 5 RMDS archive volumes
 - 1 JES checkpoint volume
 - 6 JES SPOOL volumes
 - 6 work volumes
 - 8 batch master file volumes
 - 2 batch report output volumes
 - 1 RACF volume
 - 1 started task volume
- Subsystem performance results:

Function	Cache Size	SIO/Sec	Response Time	Read Hit Ratio	Fast Write Hit Ratio	Read/Write Hit Ratio
Base Caching	128MB	275	12.8 ms	0.898	0.000	3.53
DASD Fast Write Change	128MB	281 2%	7.54 ms -41%	0.910	0.910	3.14

- Conclusions and Observations:

In addition to observing outstanding subsystem level performance improvements using DASD fast write, some significant improvements were noted for individual volumes. For example:

- Using only base caching support, the JES2 SPOOL volume provided an average response time of 24.7 milliseconds. When DASD fast write was added for the volume, average response time for the SPOOL volume dropped to 7.75 milliseconds, a reduction of 69%.
- Large response time improvements measured were 68% for the JES checkpoint and 59% for the work volumes.

Example 5 – DASD Fast Write (IMS)

The following shows the environment and measurements obtained in an aerospace industry installation.

- Software Environment:
 - IMS/VS 2.2 (full function)
 - JES 2.1.5
 - DFP 3.1.0
 - MVS/ESA 3.1.1.
- System:
 - 3090-600E partitioned (128MB main with 192MB expanded storage)
- DASD subsystem:
 - 3990 Model 3 with 36 cached 3380-K volumes
 - 32 IMS database volumes
 - 4 IMS system volumes
- Subsystem performance results:

Function	Cache Size	SIO/Sec	Response Time	Read Hit Ratio	Fast Write Hit Ratio	Read/Write Hit Ratio
Base Caching	128MB	263	11.5 ms	0.800	0.000	4.27
DASD Fast Write Change	128MB	257 –2%	8.92 ms –22%	0.786	0.979	4.08

- Conclusions and Observations:

Subsystem level performance improved 22% using DASD fast write. The average response time for the four IMS system volumes decreased from 14.5 to 8.6 milliseconds, an improvement of 41% (with each WADS volume in excess of 50%).

All five performance measurements cited show response time improvements at equal workloads, indicating the performance leverage attributed to DASD fast write. However, performance improvements may also allow the 3990 Model 3 with DASD fast write to support longer strings and more data, while providing performance equal to or better than users are receiving today, and allowing 3990 Model 3 users to configure lower cost DASD subsystems.

Guidelines for Using Cache Fast Write with DFSORT in MVS Environments

Cache fast write is supported for sort work files in DFSORT Release 10 in all MVS environments. We recommend that you use the following guidelines in implementing DFSORT with cache fast write:

- Use cache fast write for small to intermediate-size sorts. Large sorts (for example, sorts that require more sort work space than is available in the cache) should be directed to DASD volumes that are not eligible for cache fast write because the volumes are not attached to a 3990 Model 3 or are attached to a 3990 Model 3 with:
 - device caching set off for the volume, or
 - caching set off for the subsystem, or
 - cache fast write set off for the subsystem.
- Estimate approximate sort work space requirements by multiplying the input data set size by 1.7. See *DFSORT Application Programming Guide* for details.
- Ensure that the estimated total concurrent sort work space in bytes does not exceed the available cache. If other data will also be cached, reduce the total concurrent sort work space appropriately. (See the recommended process below.) If possible, schedule larger sort runs at times when the cache workload is lighter.
- Follow the existing guidelines on placement of sort work file volumes:
 - Isolate the sort work volumes from other high activity system data sets
 - Place the sort work volumes on DASD strings dedicated to sort activity.
- Sort performance will probably not be improved by using more than three or four sort work files.
- Following are recommended guidelines for selecting and migrating data to a 3990 Model 3 that is expected to have a significant sort workload:
 - Migrate the sort volumes to the 3990 Model 3 before caching the other data
 - Measure the performance
 - Gradually add more caching volumes to the subsystem
 - Monitor both the performance of the sort volumes and the performance of the added volumes
 - Continue adding more workload as long as the subsystem performance is satisfactory.

MVS Cache and Cache Fast Write Data Set Selection Guidelines

Table 9 shows data set candidates for basic caching and cache fast write operations.

Table 9. MVS Cache and Cache Fast Write Data Selection Guidelines				
Data Set Categories	System and General Data Sets	TSO	IMS	CICS
Best cache candidates	Partitioned data sets, PROCLIB, control data sets, catalogs, RACF and DFHSM control data sets	Libraries, ISPF, TSO user data, catalogs	ACB, ADF work data base, program libraries	Program libraries
Good cache candidates	Sequential input data sets		Data base indexes, moderate-sized data bases (primarily inquiry), IMS short message queue	Application data sets
Potential good cache candidates (See the following note.)	Other system-type data sets, lookup tables or dictionaries, custom systems		Image copy, RECON data set, larger data bases (with larger cache sizes), scratch pad area, long message queue, QBLKS	
Cache Fast Write candidates	SORTWK			

Note: Custom-designed systems or other software might have good cache candidates that you can evaluate by the criteria given above. For example, data sets having a read hit ratio of 65% or better and a read-to-write ratio of 2:1 are good candidates for a 3990 Model 3.

Chapter 6. Optional Features

Features Available for 3990 Storage Controls

The following optional features are available on the 3990 Models:

- Four Channel Switch, Additional (all models of the 3990)
- Remote/Local Capability (all models of the 3990)
- Dual Line Cord (3990 Models 2 and 3 only)
- 3390 Attachment (3990 Models 2 and 3 only).

Four Channel Switch, Additional

The Four Channel Switch, Additional feature 8172 permits four additional channels to be attached to each storage cluster.

Remote/Local Capability

The Remote/Local Capability feature 7149 permits the operator to enable or disable the channels either at the 3990 operator panel or at a remote switch panel. This feature should be ordered once for a 3990 Model 1 and twice for a 3990 Model 2 or 3990 Model 3, regardless of the number of channels.

Dual Line Cord

The Dual Line Cord feature is available for both already installed 3990 Models 2 and 3 as well as for new 3990 Models 2 and 3 shipped from the factory. See Table 10 for which feature to install. With this feature, a single-frame 3990 provides all the data availability and reliability that was previously available only in a 3990 dual-frame configuration.

Table 10. Dual Line Cord Feature Numbers

Feature Number	3990 Model	Type Of Power Subsystem	Factory Install	Field Install
9535	2 or 3	modular	yes	no
5352	2 or 3	modular	no	yes
4353	2	dual-frame capable	no	yes
5353	3	dual-frame capable	no	yes

3390 Attachment

The 3390 Attachment feature 6120 for IBM 3990 Models 2 and 3 allows attachment of the IBM 3390 in both 3390 mode and in 3380 track compatibility mode. The 3390 operates in four-path mode at a data transfer rate of 4.2 megabytes per second. On 3990 Model 3s attached to IBM 3090s with 4.5 megabytes per second channels, cache transfers will be 4.5 megabytes per second. 3990 Models 2 and 3 with feature 6120 installed do not support IBM 3380 Models AA4, A04, and B04, or 3380

Model CJ2; however, all other 3380 models can be attached to a 3990 along with 3390s.

IBM 3990 Models 2 and 3 with feature 6120 installed support the following mutually exclusive configurations of 3390 and 3380 A-units.

1. One 3390 Model A14, A18, A24, or A28 AND one or two (in any combination) 3380 Models AD4, AE4, AJ4*, or AK4* (* with specify 9432)
2. One 3390 Model A14, A18, A24, or A28 AND two (in any combination) 3380 Models AJ4** or AK4** (** with specify 9433)
3. One or two 3390 Models A14, A18, A24, or A28 AND no 3380 Models
4. One to four strings of the following in any combination: 3380 AD, AE, AJ*, AK* (*with feature 9432)
5. One or two 3380 Models AD4 or A4E AND two 3380 Models AJ4** or AK4** (**with feature 9433) in any combination.

Feature 6120 is supported on 3990 Model 2 and 3 dual-frame configurations; however, both machines must have the feature installed.

Features Available for 3380 Direct Access Storage Devices

One of two features is required on all 3380 Model AJ4/AK4 that attach to a 3990. These features specify whether the DASD string operates in 2-path or 4-path mode.

3380 AJ4/AK4 2-Path String Attachment

The 3380 AJ4/AK4 2-Path String Attachment feature 9432 must be installed on all 2-path 3380 Model AJ4 or AK4s before the strings are attached to a 3990 operating in DLS mode. Feature 9432 allows the 3380 AJ4 or AK4 to operate in a 2-path string. DLS mode permits two single-path storage directors to access data in the DASD subsystem.

3380 AJ4/AK4 4-Path String Attachment

The 3380 AJ4/AK4 4-Path String Attachment feature 9433 must be installed on all 3380 Models AJ4 or AK4 before the strings are attached to a 3990 Model 2 or 3 operating in DLSE mode. Feature 9433 allows the 3380 AJ4 or AK4 to operate as a 4-path string. DLSE mode permits two multipath storage directors to access data in the DASD subsystem.

Engineering Changes Required for Other 3380 Models

The 3990 does not require any features to attach to 3380 Models AD4 or AE4. However, the 3380 Models AD4 or AE4 do require the 3380 AD4/AE4 Installation/Maintenance Update Engineering Change EC465359 ECA131 to attach to the 3990.

The 3990 does not require any features to attach to 3380 Model AA4. However, the 3380 Model AA4 does require the 3380 AA4 Installation/Maintenance Update Engineering Change EC444944 ECA130 to attach to the 3990.

3390 Features

The 3390 does not require any features to attach to the 3990 Storage Control.

MB	Megabyte
MPSD	Multipath storage director
MVSCP	MVS configuration program
NVS	Nonvolatile storage
OLTEP	Online Test Executive Program
OS	Operating system
PAM	Partitioned access method
PFK	Program function keyboard
PLPA	Pageable link pack area
PTF	Program temporary fix
R0	Record zero
RACF	Resource Access Control Facility
RMF	Resource Measurement Facility
RPS	Rotational position sensing
RRDS	Relative record data set
RTM/SF	Realtime Monitor/Systems Facility
SAM	Sequential access method
SCA	Shared control array
SCDS	Source control data set
SF	Support facility
SIM	Service information message
SLR	Service Level Reporter

SMF	System Management Facilities
SML	MVS Storage Management Library
SMS	Storage Management Subsystem
SQL/DS	Structured Query Language/Data System
SSID	Subsystem identifier
TCO	Triple Capacity Option
TPF	Transaction Processing Facility
TSO	Time sharing option
TTR	Track record
UCB	Unit control block
UCW	Unit control word
VIO	Virtual input/output
VMMAP	VM Monitor Analysis Program
VMPPF	VM Performance Planning Facility
VMPRF	VM Performance Reporting Facility
VSAM	Virtual storage access method
VTAM	Virtual telecommunications access method
VTOC	Volume table of contents
VVDS	VSAM volume data set
WTO	Write to operator
XA	Extended architecture

4-path string. A series of physically connected DASD units in which the head of string provides four data transfer paths that can operate simultaneously. A 3390

4-path string requires one A-unit, while two 3380 Model AJ4/AK4 units are required for a 3380 4-path string.

Short Title	Full Title	Order Number	Contents
Operator's Library: OS/VS2 MVS System Commands	<i>Operator's Library: OS/VS2 MVS System Commands</i>	GC28-1031	Describes syntax and use of MVS/370 commands
System: Data Administration	<i>MVS/370 System: Data Administration</i>	GC26-4056	Describes MVS/370 DFP, and how to modify and extend the data management capabilities of the operating system
System Generation Reference	<i>MVS/370 System Generation Reference</i>	GC26-4063	Describes how to do a sysgen, iogen, or edtgen
VSAM Administration: Macro Instruction Reference	<i>MVS/370 VSAM Administration: Macro Instruction Reference</i>	GC26-4074	Describes using VSAM macro instructions for VSAM data sets
MVS Storage Management Library			
MVS SML: Configuring Storage Subsystems	<i>MVS Storage Management Library: Configuring Storage Subsystems</i>	SC26-4409	Describes evaluating hardware configurations, developing capacity plans, and performance, availability and space utilization considerations
MVS SML: Managing Data Sets	<i>MVS Storage Management Library: Managing Data Sets</i>	SC26-4408	Describes managing data sets, catalogs and control data sets, establishing and enforcing data set policy, and data set security
MVS SML: Managing Storage Pools	<i>MVS Storage Management Library: Managing Storage Pools</i>	SC26-4407	Describes storage requirements for groups of data sets, designing storage pools, making transition to pooled storage, and maintaining and monitoring storage pools
MVS/ESA Storage Management Library			
Focus on Storage Management	<i>MVS/ESA Storage Management Library: Focus on Storage Management</i>	SC26-4655	Introduces efficient storage management concepts; written for management
Leading an Effective Storage Administration Storage Group	<i>MVS/ESA Storage Management Library: Leading an Effective Storage Administration Storage Group</i>	SC26-4658	Setting up a storage administration group, providing service, and planning for future storage management
Managing Data Sets and Objects	<i>MVS/ESA Storage Management Library: Managing Data Sets and Objects</i>	SC26-4657	Communicating with user groups, managing active and inactive data sets, managing objects, managing catalogs and control data sets, establishing and enforcing data set and object policies, and providing data set and object security
Managing Storage Pools	<i>MVS/ESA Storage Management Library: Managing Storage Pools</i>	SC26-4656	Determining data storage requirements, designing storage pools, making the transition to pooled storage, and maintaining storage pools
Project Planning for Migration	<i>MVS/ESA Storage Management Library: Project Planning for Migration</i>	GV26-1022	This PC diskette provides samples for customers to create their own project plans, service level agreements, and user's guides
Storage Management Reader's Guide	<i>MVS/ESA Storage Management Library: Storage Management Reader's Guide</i>	SC26-4658	Listing of book titles, order numbers, and brief descriptions of storage-related libraries
Storage Management Subsystem Migration Planning Guide	<i>MVS/ESA Storage Management Library: Storage Management Subsystem Migration Planning Guide</i>	SC26-4659	Defining the principal phases and tasks required to plan and implement a migration to the Storage Management Subsystem (SMS), a component of MVS/DFP V3
Performance, Availability, and Tuning Information			
Component Failure Impact Analysis—An Availability Management Technique	<i>Component Failure Impact Analysis—An Availability Management Technique</i>	GC20-1865	Planning for hardware availability through configuration
MVS/XA: JES3 User Modifications and Macros	<i>MVS/Extended Architecture: JES3 User Modifications and Macros</i>	SC23-0060	Describes tailoring JES3 exits and macros
Cache RMF Program Description and Operation	<i>Cache RMF Reporter Program Description and Operation</i>	SH20-6295	Provides detailed information on Cache RMF Reporter

Short Title	Full Title	Order Number	Contents
Operations	<i>VM/SP Operator's Guide</i>	SC19-6202	Discussion of VM/SP operator commands and facilities, including DDR, MONITOR, VARY, and DMKFMT
	<i>VM/SP HPO Operator's Guide (Release 4.2)</i>	ST00-1898	Discussion of VM/SP HPO operator commands and facilities, including DDR, MONITOR, VARY, and DMKFMT
	<i>VM/SP HPO Release 5 CP Support for the IBM 3990 Storage Control Model 3</i>	GC23-0430	Describes the support of VM/SP HPO Release 5 for the 3990 Model 3.
	<i>VM/XA SF Real System Operation</i>	GC23-0139	Discussion of VM/XA SF operator commands and facilities, including DDR, MONITOR, VARY, and DMKFMT
	<i>VM/XA SP Real System Operation</i>	SC23-0371	Discussion of VM/XA SP operator commands and facilities, including DDR, MONITOR, VARY, and DMKFMT
	<i>VM/XA SF Virtual Machine Operation</i>	GC23-0138	Discussion of VM/XA SF operator commands and facilities
Planning	<i>VM/SP Planning Guide and Reference</i>	SC19-6201	Discussion of VM/SP hardware and software planning, system design, and system definition
	<i>VM/SP HPO Planning Guide and Reference</i>	SC19-6223	Discussion of VM/SP HPO hardware and software planning, system design, and system definition
	<i>VM/XA SF Virtual Machine Planning</i>	SC23-0167	Discussion of VM/XA SF hardware and software planning, system design, and system definition
	<i>VM/XA SP Planning and Administration</i>	GC23-0378	Discussion of VM/XA SP hardware and software planning, system design, and system definition
	<i>VM/XA Systems Facility Planning Guide</i>	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
VM/Integrated System Information			
VM/IS Planning For Your System	<i>VM/Integrated System Planning For Your System</i>	SC24-5337	Provides pre-installation planning instructions for VM/IS and should be read before <i>VM/IS Installing Your System</i>
VM/IS Installing Your System	<i>VM/Integrated System Installing Your System</i>	SC24-5341	Provides step-by-step instructions for installing VM/IS; use this book in conjunction with <i>VM/IS Planning For Your System</i>
VM/IS Managing Your System	<i>VM/Integrated System Managing Your System</i>	SC24-5338	Provides operation and administration instructions for VM/Integrated System
VM/IS Reporting System Problems	<i>VM/Integrated System Reporting System Problems</i>	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	<i>VM/Integrated System Learning to Use Your System: Error and Information Messages</i>	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
VM Performance and Monitor Information			
VM/RTM Program Description/Operations Manual	<i>VM/370 Real Time Monitor Program Description/Operations Manual</i>	SH20-2337	Description of RTM functions, commands, and reports for performance monitoring
RTM/SF Program Description/Operations Manual	<i>VM/XA Realtime Monitor/Systems Facility Program Description/Operations Manual</i>	SH26-7000	Description of RTM/SF functions, commands, and reports for performance monitoring

Short Title	Full Title	Order Number	Contents
VSE/AF System Management Guide	<i>VSE/Advanced Functions System Management Guide</i>	SC33-6352	Description of system management tasks under VSE/AF
VSE/AF Operation	<i>VSE/Advanced Functions Operation</i>	SC33-6353	Various operational considerations including the PRINTLOG function
VSE/POWER Application Programming	<i>VSE/POWER Application Programming</i>	SC33-6374	Information on using the system's accounting function
VSE/POWER Installation and Operation Guide	<i>VSE/POWER Installation and Operation Guide</i>	SC33-6371	Description of VSE/POWER input and output scheduling functions
VSE/SP Administration	<i>VSE/System Package Administration</i>	SC33-6405	Use of Interactive Interface dialogs for tasks including storage management
VSE/SP Installation	<i>VSE/System Package Installation</i>	SC33-6404	Assistance with regenerating VSE system files and libraries
VSE/SP Operation	<i>VSE/System Package Operation</i>	SC33-6406	Use of Interactive Interface dialogs for tasks including storage management
VSE/VSAM Programmer's Reference	<i>VSE/VSAM Programmer's Reference</i>	SC33-6435	Detailed information on VSE/VSAM and converting non-VSAM files
VSE Reference Information			
DITTO Program Reference and Operations Manual	<i>Data Interfile Transfer, Testing and Operations Utility for VSE and VM Program Reference and Operations Manual</i>	SH19-6104	Detailed information on using VSE/DITTO
DL1/DOS/VS Data Base Administration	<i>DL1/DOS/VS Data Base Administration</i>	SH24-5011	Assistance with moving DL/1 data bases
DL1/DOS/VS Resource Definition and Utilities	<i>DL1/DOS/VS Resource Definition and Utilities</i>	SH24-5021	Assistance with moving DL/1 data bases
VSE Fast Copy Data Set Installation Reference	<i>DOS/VSE Fast Copy Data Set Installation Reference</i>	SC33-6082	Specific information on using the VSE/Fast Copy utility
VSE/PT Program Description/Operations Manual	<i>VSE/Performance Tool Program Description/Operations Manual</i>	SH20-2171	Instructions for installation and use of optional VSE performance monitoring product
VM Running Guest Operating Systems	<i>VM Running Guest Operating Systems</i>	GC19-6212	Discussion of guest operating systems under VM
VMMAP User's Guide and Reference Manual	<i>Virtual Machine Monitor Analysis Program User's Guide and Reference</i>	SC34-2166	Description of VMMAP commands and reports

Index

A

- access method services (IDCAMS)
 - support for extended functions 58
 - support for remote support sessions 58
 - TSO execution of commands 59
- additional channels 85
- array, shared control 40
- assumptions, performance 76
- attachment of direct access storage devices 47
- availability (RAS) features 18

B

- basic write caching operations 42
- blue channel interface cables 17
- branching transfers 38
- bypass-cache 44

C

- cache
 - algorithms 43
 - description 42
 - size of 53
 - 3990 Model 3 cache size designations 53
- channel
 - commands 44
 - features 85
- channel interface cables 17
- characteristics
 - channel 40
 - dual copy 36
 - fast write 34
 - nonvolatile storage 38
 - shared control array 40
 - storage cluster 38
 - storage director 40
 - storage paths 40
 - support facility 41
- cluster, storage (see storage) 38
- commands, channel 44
- concurrent media maintenance 19
- configurations
 - dual-frame 54
 - 3990 Model 1 49
 - 3990 Model 2 51
 - 3990 Model 3 51
- connecting to 3380 AJ4 and AK4 32
- control array, shared 40

D

- DASD
 - See direct access storage device attachment
- DASD attachment 47
- DASD fast write 34
- data transfer, dual 38, 42
- device level selection
 - description 30
- device level selection enhanced
 - description 30
 - performance advantages 31
- DFSMS, user data sets 64
- direct access storage device attachment
 - string length allowed 47
 - 3380 models 47
 - 3390 models 47
- director, storage 40
- diskette, writeable 20
- DLS
 - See device level selection
- DLSE
 - See device level selection enhanced
- dual copy
 - description 36
 - overview 18
 - using 63
- dual data transfer 38, 42
- dual line cord feature 85
- dual-frame configuration 54

E

- ease of migration improvements 21
- environmental improvements 24

F

- fast dual copy 36
- fast write
 - cache fast write 35
 - DASD fast write 34
 - description 34
 - overview 17
- features
 - dual line cord feature 85
 - four channel switch, additional 85
 - on 3990 Model 1 49
 - on 3990 Model 2 or 3 53
 - optional 85
 - remote/local capability 85
 - 3390 attachment feature 85

writeable diskette 20

Numerics

2-path mode 40

3380 attachment

intermixing strings 48

models that attach to 3990 47

string length allowed 47

3390 attachment

intermixing strings 48

models that attach to 3990 47

string length allowed 47

3390 attachment feature 85

370 input/output operations 33

370/XA input/output operations 33

3990 Model 1

available features 49

upgrading to 3990 Model 2 49

with DLS subsystem 50

3990 Model 2

description of 51

with DLS subsystems 52

with DLSE subsystems 52

3990 Model 3

cache size and designations 53

description of 51

with DLS subsystems 52

with DLSE subsystems 52

4-path mode 40

Order No. GA32-0098-1

This manual is part of a library that serves as a reference source for systems 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: Copies of IBM publications are not stocked at the location to which this form is addressed. Please 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.

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If you wish a reply, give your name, company, mailing address, and date:

If we require further information, may we contact you by phone?

____ Tel. No. _____

What is your occupation? _____

Number of latest Newsletter associated with this publication: _____

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 them directly to the address in the Edition Notice on the back of the front cover or title page.)

Reader's Comment Form

Fold and Tape

Please Do Not Staple

Fold and Tape



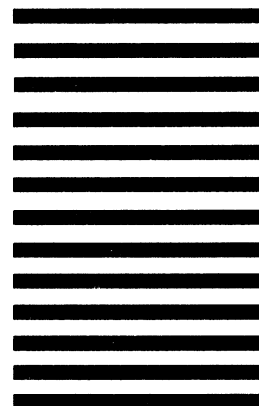
NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department 61C
P.O. Box 27155
Tucson, AZ, USA 85775-4401



Fold and Tape

Please Do Not Staple

Fold and Tape





Storage Subsystem Library

3990 Publications

IBM 3990 Storage Control Introduction	GA32-0098
IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide	GA32-0100
IBM 3990 Storage Control Reference	GA32-0099
Cache Device Administration	GC35-0101
IBM 3990 Operations Study Guide	GA32-0131
IBM 3990 Operations and Recovery Reference	GA32-0133
Introduction to Nonsynchronous Direct Access Storage Subsystems	GC26-4519

3390 Publications

IBM 3390 Direct Access Storage Introduction	GC26-4573
Using IBM 3390 Direct Access Storage in an MVS Environment	SC26-4574
Using IBM 3390 Direct Access Storage in a VM Environment	SC26-4575
IBM 3390 Direct Access Storage Reference Summary	GX26-4577

Shared Publications

Master Bibliography, Index, and Glossary	GC26-4496
Maintaining IBM Storage Subsystem Media	GC26-4495

3380 Publications

IBM 3380 Direct Access Storage Introduction	GC26-4491
Using the IBM 3380 Direct Access Storage in an MVS Environment	GC26-4492
Using the IBM 3380 Direct Access Storage in a VM Environment	GC26-4493
Using the IBM 3380 Direct Access Storage in a VSE Environment	GC26-4494
IBM 3380 Direct Access Storage Reference Summary	GX26-1678

GA32-0098-1

